

A New Era of Measurable Effects?
Essays on Political Communication in the New Media Age

Andrew M. Guess

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2016

ProQuest Number: 3742941

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 3742941

Published by ProQuest LLC (2015). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

© 2015

Andrew M. Guess

All rights reserved

ABSTRACT

A New Era of Measurable Effects?

Essays on Political Communication in the New Media Age

Andrew M. Guess

In this dissertation, I explore the ways in which traditional processes of opinion formation, media exposure, and mobilization operate in a networked, fragmented, and high-choice environment. From a methodological standpoint, one of the advantages of this shift toward Internet-mediated activity is the potential for enhanced measurement. In my dissertation, I take advantage of the data trail left by individuals in order to learn about political behavior and media effects online. Combining this measurement strategy with field experiments conducted in naturalistic online environments, I am able to shed light on how longstanding concerns in political science manifest themselves in the present-day media landscape. The overarching theme is that, thanks to advances in both research design and technology, many well-articulated concerns about the impact of the Internet on politics and public life can now be subjected to rigorous scrutiny. As I show here, the most dire predictions—about people’s tendency to cocoon themselves into ideological echo chambers or opt for low-cost “slacktivism” over more meaningful contributions to collective action—appear to lack strong support. But it is also clear that results clearly depend on the structural features of a particular medium: Twitter enables peer effects and the mutual reinforcement of viewpoints, while the high-choice environment of the Web may inherently lead to moderation.

Contents

List of Tables	ii
List of Figures	iii
Introduction	1
1 Measuring Online Media Exposure	7
2 When Treatments Are Tweets	39
3 Online Media Choice and Moderation	74
Bibliography	112
Appendices	120

List of Tables

1.1	Overview of measures	24
1.2	Determinants of misreporting of media exposure	30
1.3	Predictors of news reception	36
2.1	Top environmental Twitter accounts	52
2.2	Design and outcomes, Study 1	54
2.3	Design and outcomes, Study 2	57
2.4	Direct message effects, Study 1	59
2.5	Tweet button effectiveness, Study 1	61
2.6	Peer effects of tweet encouragement, Study 1	62
2.7	Direct message effects, Study 2	64
2.8	Tweet button effectiveness, Study 2	65
2.9	Peer effects of tweet encouragement, Study 2	67
3.1	YouGov Pulse sample demographics	83
3.2	Sample media slant scores	84
3.3	Effectiveness of treatment encouragement	92
3.4	Moderation of individuals' media diet partisanship	95
3.5	Results from YouGov Pulse experiment	104
3.6	Views toward for-profit colleges	106
3.7	Views toward for-profit colleges (earlier study)	107

List of Figures

1.1	Screen shot of colored links	18
1.2	Survey flow for Study 1	20
1.3	Results, Studies 1 and 3	27
1.4	Results, Study 2	28
1.5	News recall by story	33
1.6	Predictors of news recall	34
2.1	Illustration of Twitter network	50
2.2	Daily tweet volume of organization	51
2.3	Heterogeneous effects by gender and status	69
3.1	YouGov panel recruitment email	82
3.2	Density of site visits by ideological slant	86
3.3	Experimental study timeline	89
3.4	Individuals' media diet partisanship	94
3.5	Post-treatment media balance?	96
3.6	More comparisons	98
3.7	Results from YouGov Pulse experiment	102

Acknowledgments

First things first: Andrew McKenna originally encouraged me to go to graduate school. I have yet to regret that decision. And before I went, I had two shining examples before me: My dad, George, who discovered grad school could be both a means to a career and an end in itself, and Nicolas Blarel, who was serious about political science before I knew who Condoleezza Rice was.

This project has benefited—immeasurably—from the comments and suggestions of others. With apologies to those I’ve inevitably left out, then, I thank Noah Buckley-Farlee, Alex Coppock, Kevin Elliott, Andrew Gelman, Bryan Gervais, Don Green, Daniel Hopkins, Trish Kirkland, Lucas Leemann, Diana Mutz, Jonathan Nagler, Christopher Ojeda, Justin Phillips, Markus Prior, Jaime Settle, Gaurav Sood, Dannagal Young, and Adam Zelizer. For insight and encouragement along the way, I am also grateful to Farid Ben Amor, Al Fang, Neelan Sircar, and Jeremy Singer-Vine.

Since my first semester at Columbia, Bob Shapiro has been an unfailingly supportive advisor who saw the value in my work even when I didn’t. I’m glad I listened, and I continue to be inspired by his ability to meld pathbreaking quantitative research with strong normative convictions. Whether in the classroom—or more often outside of it—Jeff Lax helped me understand the practice of political science.

Most importantly, he gave me just the right mix of encouragement and real talk. Too often we hear only one or the other.

Very, very special thanks are due to Doug Rivers, Brian Law, and Matthieu Devlin at YouGov for facilitating access to the Pulse panel data and for their patience and forbearance during the process of data collection. Their assistance to me went above and beyond what could normally be expected from a survey research firm, and they took my somewhat oddball ideas seriously.

It may not be a popular thing to admit, but the truth is that two of the chapters were greatly improved thanks to the magic of peer review and the thorough comments of five anonymous reviewers and several supportive journal editors.

Finally, I've been extremely fortunate to benefit from the support of institutions. Columbia University and the Department of Political Science provided extensive research and administrative support for my project. Some of the analysis in this dissertation was likewise made possible by Columbia's Hotfoot High Performance Computing (HPC) Cluster. The eagerness of the League of Conservation Voters to include new experimental designs in its programs made Chapter 2 possible. And in signaling its support, the American Association for Public Opinion Research showed me that other people might be paying attention to my research.

To Mama, Dada, and Marty. Thanks for sticking with me throughout. Mama, thank you for believing in me. Dada, your war stories and general wisdom helped me through the tough parts. Marty, you made it out before me but you showed me it could be done!

Introduction

In this dissertation, I explore the ways in which traditional processes of opinion formation, media exposure, and mobilization operate in a networked, fragmented, and high-choice environment. These concerns are not trivial: Extrapolating from current trends, most Americans will soon receive most of their news online rather than from television or newspapers. And beyond serving as a source of information, the Internet is increasingly where politically relevant phenomena now take place—from micro-level behavior such as deliberation and volunteering to elite-level campaign strategies.

From a methodological standpoint, one of the advantages of this shift toward Internet-mediated activity is the potential for enhanced measurement. Much has been said about “big data,” the massive amounts of information about everyday consumer behavior that accumulates online, either in public waiting to be collected or on proprietary servers. In my dissertation papers, I take advantage of this data trail left by individuals in order to learn about political behavior and media effects online. In so doing, I not only advance the field by helping to orient it toward emerging political and social habits, but I address several vexing problems with the existing literature on these topics.

Another advantage of focusing on the Internet is that the nature of the medium

offers intriguing opportunities for innovative research designs. For some research, the use of online surveys may be a matter of cost or convenience; when studying online media, however, administering survey instruments via the Internet is the approach that maximizes realism. Since survey tools deployed online are embedded within a subject's existing informational environment, ecological validity can feasibly be achieved. Thus, by situating the measurement where the behavior of interest actually occurs, I can improve the study of how individuals' self-constructed media environments interact with their beliefs and actions.

Combining these measurement strategies with field experiments conducted in naturalistic online environments, I am able to shed light on how longstanding concerns in political science manifest themselves in the modern, fragmented media landscape. I begin by confirming the need for new measurement approaches: As many previous studies have suggested in other contexts, survey evidence of media exposure is plagued with overreporting and leads to inflated estimates of the overall political news audience online. I then demonstrate two ways of resolving this problem. Next, turning to social media, I (along with two graduate-student co-authors) employ behavioral outcome measures to study the effectiveness of strategies for mobilizing grassroots activists on Twitter. This set of experiments not only shows how to encourage political participation in a new context, but it suggests that network effects—in which subjects pass on campaign messages to their own followers—can be more powerful than appeals from the original source. Finally, applying some of the methods I introduce in the first paper, I conduct a test of the partisan selective exposure hypothesis online.

The overarching theme of these papers is that, thanks to advances in both research design and technology, many well-articulated concerns about the impact of the Internet on politics and public life can now be subjected to rigorous scrutiny. As I have done here and continue to do in my ongoing work, the most dire predictions—

about people’s tendency to cocoon themselves into ideological echo chambers or opt for low-cost “slacktivism” over more meaningful contributions to collective action—appear to lack strong support. But it is also clear that results clearly depend on the structural features of a particular medium: Twitter enables peer effects and the mutual reinforcement of viewpoints, while the high-choice environment of the Web may inherently lead to moderation. This dissertation strives to reorient research on the Internet and politics toward a more nuanced appreciation of the interaction between political behavior, information, and media structure.

Below, I summarize each paper in turn.

Paper 1

My first dissertation paper is an introduction and validation of the technique I propose for producing behavioral measures of individuals’ online exposure to political media. Via a series of experiments, I show that survey-based self-report methods lead to substantial overreporting of exposure to individual online media sources (including both large sites such as CNN.com and smaller, partisan outlets). However, questions in which respondents are asked to list the sites they visit in an open-ended format produce the least amount of overreporting. Still, if possible to obtain, behavioral measures are clearly preferable for measuring exposure. I demonstrate two methods for doing so: indirectly, by taking advantage of the way that web browsers display visited and unvisited links, and directly, by asking subjects to install a piece of software that scans their web histories against a fixed list of political sites.

Abstract

Self-reported measures of media exposure are plagued with error and questions about validity. Since they are essential to studying media effects, a substantial

literature has explored the shortcomings of these measures, tested proxies, and proposed refinements. But lacking an objective baseline, such investigations can only make relative comparisons. By focusing specifically on recent Internet activity stored by web browsers, this paper's methodology captures individuals' actual consumption of political media. Using experiments embedded within an online survey, I test three different measures of media exposure and compare them to the actual exposure. I find that open-ended survey prompts reduce overreporting and generate an accurate picture of the overall audience for online news. I also show that they predict news recall at least as well as general knowledge. Together, these results demonstrate that some ways of asking questions about media use are better than others. I conclude with a discussion of survey-based exposure measures for online political information and the applicability of this paper's direct method of exposure measurement for future studies.

Paper 2

For the second dissertation paper, we designed and implemented field experiments on the social network Twitter to study the effects of different online mobilization strategies. Aside from producing novel results on the relative effectiveness of tweets and private direct messages, we also developed innovative design enhancements that allowed us to avoid the many pitfalls of experimenting over a public network.

Abstract

This study rigorously compares the effectiveness of online mobilization appeals via two randomized field experiments conducted over the social microblogging service Twitter. In the process, we demonstrate a methodological innovation designed to capture social effects by exogenously inducing network behavior. In both experiments, we find that direct, private messages to followers of a nonprofit ad-

vocacy organization's Twitter account are highly effective at increasing support for an online petition. Surprisingly, public tweets have no effect at all. We additionally randomize the private messages to prime subjects with either a "follower" or an "organizer" identity but find no evidence that this affects the likelihood of signing the petition. Finally, in the second experiment, followers of subjects induced to tweet a link to the petition are more likely to sign it—evidence of a campaign gone "viral." In presenting these results, we contribute to a nascent body of experimental literature exploring political behavior in online social media.

Paper 3

My third dissertation paper applies the measurement approach of the first by leveraging—for the first time—data on the actual browsing behavior of an online panel of respondents maintained by the survey research firm YouGov. In the first place, this approach allows me to demonstrate the power of direct behavioral measures of media exposure on a larger sample. Then, in a companion study, I deploy two online field experiments exploring subjects' open-ended search strategies for learning about political topics. In addition to overcoming the measurement problem, this approach allows for valid causal inference on real-world outcomes—a major difficulty for studies of media effects, online or not. Studying these processes online is important for identifying how media choice and information interact in an era of selective exposure.

Abstract

Despite decades of frequent hypothesizing and inconclusive testing, evidence for partisan selective exposure remains elusive. Research on this venerable phenomenon tends to be either observational (using behavioral measures of exposure) or experimental (in which researchers manipulate the content of treatments in a laboratory

setting). This study advances the literature on several fronts. First, I focus on Internet media, already the dominant source of political information for younger Americans and quickly becoming so for the rest of the population. This has theoretical implications for how a fragmented, high-choice environment moderates traditional media effects. Second, in collaboration with an online polling firm, I collect unique and unprecedented data tracking the real-time browsing behavior of a panel of Internet survey respondents. Third, I supplement this observational portrait with two online field-experimental studies to test the real-world behavior of smaller but similar subject pools. In both studies, I employ a novel measurement strategy that allows me to capture trace data on individuals' actual consumption of political media. This allows me to directly study how open-ended search strategies for political information vary by partisan affiliation. In doing so, I uncover consistent evidence of moderation rather than selective exposure. I also find mixed evidence on the downstream effects of media consumption on attitudes, opinions, and knowledge.

Measure for Measure: An Experimental Test of Online Political Media Exposure

Survey-based measures of media exposure are notoriously imperfect. Respondents tend to inflate their news consumption in self-reports, which can paint a misleading picture of the overall news audience and bias estimates of regression parameters in studies of media effects. Recent research has estimated that for network news, this measurement error can overstate aggregate exposure by at least 200 percent (Prior 2009a).

Such problems, and persistent findings since the 1940s of “minimal effects” in observational studies, once led some scholars to abandon the media-effects enterprise entirely (Lazarsfeld, Berelson and Gaudet 1944; Graber 1997). Nearly 50 years after the canonical studies of the Columbia School, Bartels described this state of affairs as “one of the most notable embarrassments of modern social science” (1993).

Yet research on media effects continues. Since the pioneering experiments of Iyengar and Kinder (1987), scholars have worked to identify how exposure to campaign content or news coverage in the lab affects opinion change, candidate eval-

uations, vote choice, and other outcomes. Others have since studied these effects in the field (e.g., Gerber et al. 2011). But even though researchers in such designs control the content and timing of treatments, measuring subjects' exposure can still be necessary for the analysis—for example, to estimate complier average causal effects (CACE). At the same time, survey-based measures, in overcoming problems of external validity, remain a vital tool for disentangling the role of the media in real-world studies of political behavior.

Thus far, efforts to validate existing survey-based measures have suffered from a persistent problem: lacking a measure of true exposure (e.g., Prior 2009*b*; Romantan et al. 2008). I address this shortcoming by demonstrating the usefulness of a simple yet powerful new method of recovering subjects' actual history of exposure. The approach necessitates a shift in focus from the more traditional domain of television and newspapers to the Internet. In so doing, this paper provides the first evidence of the accuracy of measures designed specifically to capture online political news consumption.

Such measures will only increase in importance as more Americans use the Internet as their primary source for news and information about the political world. A recent Pew report found that 39% of respondents went online or used a mobile device to catch up on news "yesterday," and this proportion increased to 50% when social networks, email, and podcasts were included (Pew 2012). Given continuing generational gaps in news consumption habits, online sources will likely overtake more traditional media in the near future. But so far, the literature on media exposure has lagged behind.

The measure of actual exposure is constructed from subjects' browsing history. Web browsers, such as Firefox and Google Chrome, continually collect information on which pages users visit and when they do so. These are saved locally and remain secure. However, it is possible with users' permission to obtain data on which of a

predetermined list of sites has been visited within a given period of time. This can be matched to specific question wordings in surveys. In this paper, I demonstrate how to generate the data both indirectly and directly: first by taking advantage of how browsers display visited and unvisited links to users, and second by using a browser plug-in that collects the information automatically.

I find that among the three self-reported measures of media exposure tested, asking respondents to list the websites they have visited to learn about politics in an open-ended format produces the best results. In the aggregate, open-ended questions recover an accurate picture of the overall online audience for news, and at the individual level they cause substantially less overreporting. I also show in a validation exercise that exposure measures constructed from open-ended responses predict the reception of recent political news stories better than general political knowledge. However, there is still a worrying level of absolute misreporting across all question types, reinforcing the arguments of some scholars that direct measures—such as the one demonstrated in this paper—should be used when possible for studying individuals' political behavior. However, the open-ended survey measure appears suitable for use in capturing aggregate exposure.

The paper proceeds as follows. First, I briefly review the critiques of existing exposure measures and discuss the special challenges of studying online media. Second, I outline three distinct experimental designs intended to test and benchmark the survey-based measures of online media exposure. Third, I present basic results comparing the performance of the measures against each other and against the true measure on several different metrics. Fourth, I analyze characteristics of both individuals and sites that predict misreporting in surveys. Fifth, I test the survey-based exposure measures against a measure of general political knowledge in predicting news recall. A final section then concludes with a discussion of how both open-ended survey questions and the direct measures I introduce can improve

research.

Improvements vs. Proxies

Scholars have proposed a number of workarounds to improve existing exposure measures. One class of solutions aims to incentivize accuracy among respondents, either by offering them money or giving them more time (Prior and Lupia 2008). Another attempts to anchor respondents' expectations by providing them with basic population-level summaries of behavioral frequencies, such as how often most Americans watch network news broadcasts (Burton and Blair 1991; Prior 2009b).

These proposed solutions may not be reliable, however, and reporting population-level data requires this information to be known in advance. Especially in the case of questions probing subjects' new-media consumption habits, this may not always be the case, or at least the population frequencies in question may not be stable over time.

Other scholars have concluded that substitutes should be used instead. Most prominently, Price and Zaller (1993) have proposed and validated general political knowledge questions as a proxy for news media exposure. However, as Althaus and Tewksbury (2007) note in their report on media use measures in the American National Election Studies (ANES), such knowledge questions must be periodically aligned with contemporary realities and thus may make longitudinal comparisons difficult.

The final approach is to forego exposure measures entirely in favor of *direct* behavioral measures. A notable recent attempt to do so relied on a mobile phone-based ambient recording system deployed by a media research firm in Chicago and New York (LaCour 2013). By matching digital audio signatures with actual broadcasts in subjects' media markets, the innovative study design enabled the collection of accurate, continuously updated records of exposure to specific programs on tele-

vision and radio. Unfortunately, the likelihood of another such study in the near future is low: Issues of cost aside, the company, Integrated Media Measurement Incorporated, has been acquired and its study discontinued.

In sum, while the problems of relying on self-reports are manifest and well-known, the proposed alternatives have shortcomings of their own. Improving on existing measures should continue to be a priority, with special attention to the shifting media landscape in which studies are being conducted.

Survey-Based Measures of Online Media Exposure

One of the main recommendations of Althaus and Tewksbury's report for the ANES is to adapt existing measures to more closely match respondents' experiences with newer forms of media. Such efforts will have to address the challenges discussed by Mutz and Young (2011, 1023-1024) in their recent survey of public opinion and communication research:

Indeed, asking the average citizen whether he or she watches, reads, or listens to "news" these days is the classic example of a bad survey question because the very definition of what constitutes "news" is in flux. Because scholars have yet to come to grips with all of these recent changes, we know little about where people are getting their exposure to political information and argument, and whether the source makes any difference.

An important feature of the online news environment that differs markedly from the traditional realms of television, radio, and newspapers is the sheer quantity of available sources. The Internet offers potentially thousands of avenues for learning about the political developments of the day, from mainstream news sites and aggregation portals to blogs and other online publications targeted to narrower audiences. A corollary to this proliferation of sources is that outlets often cater to a specific range of the ideological spectrum (Mullainathan and Shleifer 2005).

But if the number of potential sources of exposure presents an additional difficulty to the measurement task, the way people interact with political news online may actually make it easier. The bulk of Internet news consumers primarily visit a small handful of mainstream websites—many but not all of which are associated with established newspapers or cable channels (Gentzkow and Shapiro 2011; Hindman 2008). Moreover, while TV viewership may be considered at least in part a passive activity, accessing political information online is a more purposeful activity that might lead to more accurate recall (Mutz and Young 2011).

Given these attributes of the online media environment, the studies presented here allow for a wide range of specific online sources to be measured while focusing as much as possible on subjects' deliberate—as opposed to incidental—exposure to political content.

Types of Survey Questions

Based on a review of the survey response and media exposure literatures, I have identified three different types of questions that can be adapted for capturing online media habits for political news.

First, *check-all questions* allow respondents to check off any number of options from a given list of information sources, from none to all. This is a common format: For example, the National Annenberg Election Survey (NAES), in its 2008 online panel, allowed multiple selections from lists in its exposure questions on television. Dilliplane, Goldman and Mutz (2013) validate this “program list technique” as a measure for exposure to political information on television. The ANES also adopted versions of this question format in its 2012 Evaluations of Government and Society Study and pre-election questionnaire.

Second, *open-ended questions* allow respondents to give their answers without restrictions on length or the choice set (Sudman, Bradburn and Schwarz 1996,

225). The Pew Research Center has used this format in its Internet & American Life Project surveys; its 2010 post-election survey asked respondents which web sites they used for campaign and election news, coding up to three responses.

And third, *forced-choice, yes/no questions* (Smyth et al. 2006) require respondents to choose either “Yes” or “No” to each question in a sequence, in this case a list of online political news sources. “Don’t Know” and not answering are not options with this type of question.

Given these questions, I derived and pre-registered (Humphreys, Sanchez de la Sierra and van der Windt 2013)¹ two hypotheses on their relative performance as benchmarked against the measure of true exposure, which is described further in the next section.

First, Sudman, Bradburn and Schwarz (1996) argue that open-ended questions are useful for self-reporting of behavioral frequencies, an application closely related to the measurement of media exposure (especially frequent media exposure, as it is often framed in question wording). One advantage of this type of question is that since respondents are not presented with a pre-selected list of choices, it reduces errors in judgment caused by the familiarity of a given item: “Reliance on accessibility or familiarity, however, also leaves a respondent open to error. Although true frequency can increase familiarity, so can factors like ease of perceiving an item, expectation induced by context, and probably many other variables” (Tourangeau, Rips and Rasinski 2000, 142-143). This leads to my expectation that open-ended prompts will counteract respondents’ tendency to overreport their exposure to online political media as compared to the check-all format:

¹Before analysis began, I uploaded a pre-registration document to my public academic website: <http://polisci.columbia.edu/files/polisci/u227/PreregistrationDesignDocumentGuess2013a.pdf>

H1: Open-Ended vs. Check-All

Respondents to open-ended prompts will be less likely to overreport their exposure to online political media than respondents asked to select from specific outlets in a list.

Unlike open-ended questions, both the forced-choice and check-all question types present respondents with a list of choices. However, the task faced by respondents differs substantially. Check-all questions can encourage satisficing (Krosnick 1991): respondents can leave all items unchecked or select only the first option they can justify. But forced-choice questions require a response one way or another for each item, which should make satisficing harder (Sudman and Bradburn 2012). It should also result in more “Yes” responses, “both because respondents process throughout the list and because they more deeply process each individual response option, making them more likely to think of reasons the options apply” (Smyth et al. 2006):

H2: Forced-Choice vs. Check-All

Respondents to a successive series of forced-choice, yes-no prompts will be more likely to overreport their exposure to online political media than respondents asked to select from specific outlets in a list.

Research Design and Data

To test these hypotheses, and to more generally determine the accuracy of these three survey-based measures, I deployed a series of survey experiments on Amazon.com’s Mechanical Turk online marketplace (MTurk).² In addition to a pilot study in 2012, the surveys were in the field in four separate batches from November 22, 2013 to January 11, 2014.

²The MTurkR package by Leeper (2013) proved useful for this step.

MTurk is an online marketplace for posting modular, low-cost “human intelligence tasks” (HITs). Among other applications, it enables quick recruitment of subjects comprising an opt-in online sample in many ways more representative of the general population than samples found in published lab experiments. For example, Berinsky, Huber and Lenz (2012) found that while MTurk samples are younger than other adult convenience samples, the subjects are older than those found in convenience samples composed of students at college campuses.³

Data collected in this way can be susceptible to the weaknesses of crowdsourcing in general, namely habitual users and “spammers” who race through surveys or submit meaningless responses. For this paper, I took two steps to reduce the possibility of data contamination. First, at the front end, participation was restricted to registered users who are residents of the United States, age 18 or older, and whose approval rate (as determined by task requesters) for completing HITs was 95% or greater. Second, I included a “screener” question at the beginning of the survey to induce attentiveness. Following one of the recommendations of Berinsky, Margolis and Sances (2013), respondents could only continue with the survey after successfully completing the screener question. Since the authors found that “shirkers” differ from other respondents in meaningful ways, I opted for this “training” approach and did not drop any respondents.⁴

In the survey experiments, respondents were randomly assigned to one of the

³MTurk non-probability samples are decidedly skewed along several dimensions. Using a nationally representative face-to-face survey as a benchmark (the 2008 ANES), Berinsky, Huber and Lenz (2012) reported that in their investigation, MTurk respondents were more likely to be female (60.1% versus 54.9%), better educated, have lower income, and be white (95% versus 86.8%). Subjects recruited by MTurk are also far less likely to be religious (41.8% said they have no religion, compared to 20.1% in the ANES) or married (39% versus 50.1%), somewhat more liberal or Democratic-leaning, and less interested in politics.

⁴The screener question asked respondents for their favorite color, but only those who read all of the text saw the following instructions: “To demonstrate that you’ve read this much, just go ahead and select both red and green among the alternatives below, no matter what your favorite color is. Yes, ignore the question below and select both of those options” (see Berinsky, Margolis and Sances 2013). An error message was displayed until respondents correctly answered the question, after which they were allowed to continue.

three question types for measuring online media exposure. The questions were designed to hold as much as possible constant between conditions, including the available choices and the time period respondents were asked to think about (30 days). Question wordings were as follows.⁵

Check-all condition. This question, adapted from Annenberg, read: “Which of these websites have you visited or used in the past 30 days for news, if any? Select ALL answers that apply.” Respondents were shown a list of 27 sites, including the main network and cable news homepages, news websites such as NYTimes.com, niche sites such as Politico, and a balanced selection of smaller partisan websites. Respondents were also given the option to check “Other (please specify),” which included a text box.

Open-ended condition. Subjects were given a large blank text box with the following instructions: “Please list any websites or blogs that you have visited in the past 30 days for news. **Take some time to ensure that you think of all the sites you have visited.**” Respondents were encouraged to “take some time” and “think” about all the online sources they have visited in order to coax a full response in the absence of pre-selected choices. (The word “blogs” was also added to ensure that subjects did not skew their responses toward large or mainstream news sites.)

Forced-choice condition. Respondents were shown an ordered list of websites—the same as in the check-all condition—with the choices of “Yes” and “No” given next to each and the question “Have you visited this site?” Above the list was this sentence: “Which of these websites have you visited or used in the past 30 days for news, if any?”

To determine the list of choices given to respondents in the first and third conditions, sites were first selected on the basis of total visits as determined by a pilot administered in the summer of 2012. This was then augmented to ensure that sites

⁵See Appendix A1 for full wordings and available responses for all three questions.

for the major cable and network news organizations, newspapers, and large online portals (AOL, Yahoo!) were covered. Finally, a balanced selection of sites for partisan and nonpartisan political enthusiasts was added (e.g., ThinkProgress, Politico, Drudge Report).⁶

Study 1: The Link Classification Technique

In Study 1, respondents were asked to complete a *Link Classification Task* after answering one of the randomly assigned survey questions. This task allowed for the creation of the measure of actual online media exposure.

To do this, I presented subjects with an extended list of “masked” URLs of political news sites and blogs. Each link was displayed with exactly the same generic text, indicating nothing about the identity or content of the corresponding URL. Web browsers display visited and unvisited links in different colors; the survey exploited this feature by asking all respondents to check off each line containing a link displayed as visited (i.e., purple in most cases) by their browsers. Subjects were told that the “linked websites are all related to political news.”

As illustrated in Figure 1.1, respondents saw a long page filled with hyperlinks labeled solely with the word “LINK”—a sea of purple and blue. To the left of each line on the page, there was a single check box. Following the instructions, any line containing a purple (“visited”) link was to be checked. In this way, the survey collected a snapshot of subjects’ recent browser history for those specific sites.

The list of sites included in the *Link Classification Task*—and thus the universe of sources with which the survey-based measures can be benchmarked—was assembled as follows.⁷ First, adapting a strategy used by Gentzkow and Shapiro

⁶The social link aggregator Reddit was also included with the expectation that the site is popular among respondents recruited via MTurk.

⁷See Appendix A2 for a full list of the sites included. The order of the links displayed was randomly shuffled for each respondent.

Remember, please check **ALL** rows containing any links shown in **PURPLE**. Leave all other rows unchecked.

- [LINK LINK](#)
- [LINK LINK](#)
- [LINK LINK](#)
- [LINK LINK](#)
- [LINK LINK LINK](#)
- [LINK LINK](#)
- [LINK LINK](#)
- [LINK LINK](#)
- [LINK LINK](#)
- [LINK LINK](#)
- [LINK](#)
- [LINK LINK LINK](#)
- [LINK](#)
- [LINK](#)
- [LINK LINK](#)

Figure 1.1: A screen shot of how the *Link Classification Task* looked to survey respondents. Visited links were purple, and unvisited links were blue.

(2011), I obtained comScore data on total unique visitors per month to U.S. websites, averaged from July to August 2013 for users at home and at work. I first took all sites with at least 300 unique visitors per month on average in the “General News” category and all sites with at least 50 unique visitors per month on average in the “Politics” category (which was smaller). I then added any sites in the larger “News/Information” category with at least 1,500 unique visitors that were not already included. Next, I winnowed down this list by removing sites that were primarily local in focus, non-news-related, or not based in the U.S. (with the exception of sites like BBC and *The Guardian* that cover American politics extensively). I ensured that the sites of the top 10 newspapers by total average circulation (both print and digital as of March 2013, according to the Alliance for Audited Media) were included. Finally, I crawled several directories of partisan political blogs to

complete the list.⁸

The final result was a list of 155 links, along with a limited number of variations to ensure that the sites were properly captured by the technique.⁹ Each line of the *Link Classification Task* consists of one to three separate URLs pointing to the same site. URLs were coded to point to main pages in order to capture purposeful behavior rather than incidental visits to specific pages.

This method of constructing the true measure of online media exposure relies on the assumption that subjects, when presented with generically labeled hyperlinks to sites on the Internet, will truthfully distinguish between those marked as visited and those marked as unvisited by their browser software. This assumption is plausible since the task, divorced from any substantive context involving politics or media, is mechanical and straightforward. However, it does need to be tested.

A validity check using hard-coded color displays demonstrated that the method is generally accurate. In Study 1, a fourth treatment condition randomly presented any of the three survey questions, followed by a version of the *Link Classification Task* in which all the link colors were fixed in advance unbeknownst to subjects. Thus, there was a predetermined “correct” distribution of responses. Figure 1.2 illustrates Study 1’s design.

I randomly selected five of the 155 links to be hard-coded as visited.¹⁰ The proportion of the 132 subjects in the validity condition who correctly marked each visited link ranges from 89.4% to 97%, with an average of 95% accuracy. The low

⁸I used the following listings as additional sources for sites to include: the top 25 in Alexa’s Politics > News and Media and News > Weblogs categories; Evan Carmichael’s Top 50 Political Blogs of 2009; Rightwing News’ 100 Most Popular Conservative Websites for 2013; The Free Republic’s Top 100 Conservative Political Websites of 2007; and the List of Political Blogs (<http://politicalbloglistings.blogspot.com>).

⁹For example, a user might visit Yahoo! by typing either “www.yahoo.com” or “yahoo.com”; both versions were included for sites where applicable.

¹⁰This sparseness was by design: In the 2012 pilot, the mean number of sites checked as “visited” was 3.4.

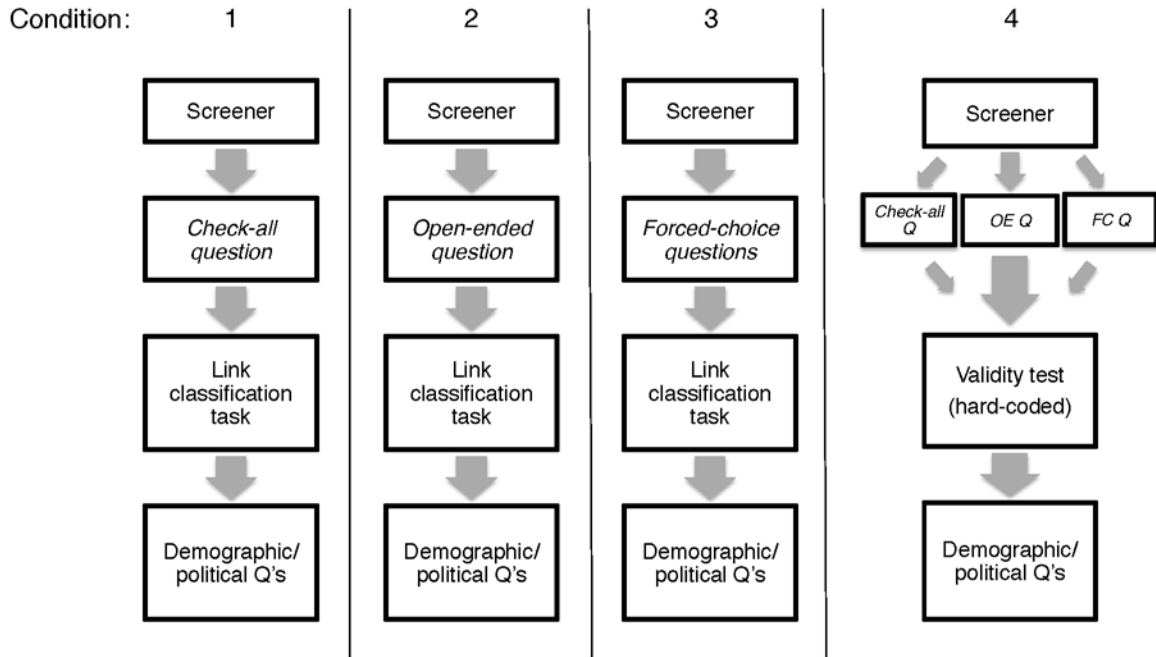


Figure 1.2: Survey flow for Study 1. Subjects were randomly assigned to treatment conditions 1-4.

end of the range corresponds to the fact that I marked the second but not the first of two links in row 93 as visited. The high end of the range corresponds to the purple link in the first row, consistent with satisficing. The proportion of subjects who correctly left each unvisited link unmarked ranges from 99.2% (in only two cases) to 100% (in the other 148 cases). Overall, this test verifies that the method of asking subjects to manually check visited (i.e. purple) links provides an accurate measure of browsing history.

Study 2: Browser Widget

In Study 2, the *Link Classification Task* was replaced with an automated method of collecting the same information: a software plug-in for Google Chrome web browsers. The other change from Study 1 was that there were only three conditions, one for each survey question type; the validity condition was not needed in this

design.

Since Chrome has powerful capabilities for querying users' web histories necessary for this kind of application, I only allowed subjects using that browser to enter the study.¹¹ Google Chrome is currently the second most widely used web browser in the United States, behind Microsoft Internet Explorer.¹² But among the MTurk respondent pool, Chrome is by far the most popular: 62.9% of respondents in Study 1 used the browser to take the survey.

I implemented the design by asking respondents at the end of the survey to follow a link to install the browser extension. Since requiring workers to install software was a violation of MTurk policies when the study was conducted, this step was left as optional for bonus credit. In the final step of the survey, respondents were given a unique identifier code and told to run the extension by clicking the button added to their browser's toolbar. When this was done, a small window appeared asking for the code. Once users submitted the code, the plug-in immediately scanned their web history, encoded the data, and sent it to the survey database. Subjects were encouraged to uninstall the extension once it finished.

The advantages of this approach are considerable: If the browser extension is installed and run properly, there will be no measurement error in retrieving a subject's exposure data. Perhaps even more attractive is the ability to precisely control the time period for which the browser history is queried. The *Link Classification Method* is susceptible to variations in the frequency with which subjects clear their browser histories, which could range from daily to never. With the browser extension, both the software and the survey questions were calibrated to the same time period—the past 30 days.

¹¹The browser extension was written in JavaScript using the Google Chrome API. Munson, Lee and Resnick (2013) used a similar approach in a recent study of media slant in users' online reading habits.

¹²29.99% of users compared to 37.5% for IE, according to StatCounter web analytics for January-December 2013. Source: <http://gs.statcounter.com/#browser-US-monthly-201301-201312-bar>

A final advantage of this approach is that the software can be adjusted for different search criteria. In the present design, the code was written to capture intentional behavior: For instance, a visit to a specific page on NYTimes.com but not the main page itself was *not* counted as exposure. This most closely matches the way exposure history was captured in the *Link Classification Task*. However, future designs could easily change this rule and incorporate any number of permissible URL variations via regular expressions.

From the respondent's perspective, the process is likely less tedious and time-consuming than scrolling through a list of 155 links. However, there are some downsides. First, some users may feel uncomfortable with installing third-party software that accesses their web history. I addressed this issue by making clear that the plug-in only searches against a fixed list of sites (identical to the one used for the *Link Classification Task*) which respondents could view before continuing. Second, and relatedly, this method necessarily cut down the subject pool. In addition to non-Chrome users, survey respondents who did not wish to install a browser extension were dropped from the analysis.

As a result, the effective response rate for this design was 14.6%, for a final sample size of 85 (check-all condition: N=30; open-ended condition: N=21; forced-choice condition: N=34). Although this attrition raises the possibility of bias, I show in Appendix A4 that demographic and political characteristics of the respondents were comparable between the full sample and the subset that installed the browser extension.

Study 3: Testing News Recall

While the previous two designs focused on accuracy, Study 3 was intended to investigate how well the survey-based exposure measures predict news recall as compared to a measure of general political knowledge (Price and Zaller 1993), thus

serving as an important test of validity.

Like Study 2, Study 3 kept only the three main treatment conditions. The other difference was the inclusion of three questions testing the recall of recent political news stories. These were modeled on a series of questions included in the 1989 ANES Pilot Study and analyzed by Price and Zaller (1993) in their original validation of general knowledge questions as proxies for news reception.

The wordings were as follows, each beginning with a yes-or-no question and followed by an open-ended question which was then coded:

1. “Do you remember any recent stories about a new policy in Colorado taking effect on Jan. 1, 2014? [If yes:] Do you recall what that policy involves?”¹³
2. “Have you seen or heard any news stories about the *Duck Dynasty* star Phil Robertson? [If yes:] Do you know why he was recently in the news?”¹⁴
3. “Do you remember hearing about certain federal government benefits that expired at the end of 2013? [If yes:] Do you recall what kind of benefits they were?”¹⁵

The general knowledge questions were included in all versions of the design.¹⁶

Measures

In the analyses that follow, I make use of several measures of accuracy. I construct measures capturing individuals’ tendency to overreport, underreport, and misre-

¹³Correct answer: legalization of marijuana for recreational use

¹⁴Correct answer: inflammatory comments about gays and African Americans

¹⁵Correct answer: expiration of Emergency Unemployment Compensation program

¹⁶I asked four questions: (1) “Who is the Chief Justice of the U.S. Supreme Court?”; (2) “For how many years is a United States senator elected – that is, how many years are there in one full term of office for a U.S. senator?” (open-ended); (3) “Who is the Speaker of the House of Representatives?”; (4) “On which of the following does the U.S. federal government spend the least money?” with available choices of foreign aid (the correct answer), Medicare, national defense, and Social Security.

port true exposure in either direction. Table 1.1 summarizes how each measure is constructed.

Exposure Data		Measures		
Actual	Reported	Overreporting	Underreporting	Misreporting
0	0	0	0	0
0	1	1	0	1
1	0	0	1	1
1	1	0	0	0

Table 1.1: Overview of the accuracy measures used in the analysis. For a given individual, each line represents one of the possible combinations of actual and reported exposure to a site.

To code responses to the open-ended question, I first examined the results, inductively generating a running list of mentioned websites (along with spelling variations) included in the *Link Classification Task*. I separately collected common shorthand designations for various sites in the list. Then, I coded exposure by iteratively running text searches for the resulting strings (disregarding case).

For each respondent, these measures are summed over all the possible websites they could have reported visiting given the treatment condition (27 for the check-all and forced-choice conditions, 155 for the open-ended prompt). Thus, in the table above, the measures of overreporting, underreporting, and misreporting would be created by summing down the rows. In the same way, I also generated raw counts of individuals' total number of actual and reported site visits.

Finally, at the site level, I created measures of total visits (both actual and reported) per site across all respondents in a given condition.

Results

For the main results using the *Link Classification Task*, I pool across Studies 1 and 3, keeping observations from treatment groups 1-3 (N = 1,112). For these observations

the design is identical.¹⁷

The total number of visits per site across all conditions—according to the measure of actual exposure—tracks the findings of Gentzkow and Shapiro (2011), who showed that the bulk of traffic to online political news sites is taken up by a relative few mainstream outlets.¹⁸ While this would never be true in a representative sample, Reddit, a social news aggregator, was by far the most visited site among the MTurk respondents, with 558 total visits.¹⁹ The next-most-visited site was CNN.com, with 293 visitors in the sample, followed by The Huffington Post with 146 and BuzzFeed with 131. The counts drop off fairly quickly, leaving a “long tail” of smaller, more narrowly targeted outlets such as Politico (37) and Talking Points Memo (14).

Figure 1.3 breaks down these total counts by treatment condition, showing both actual and reported exposure to online news sources summed across respondents. The left graph shows reported versus actual exposure for those assigned to the open-ended prompt. There does not seem to be an obvious pattern of either over- or underreporting exposure, at least in the aggregate. The middle graph focuses on the check-all question, which seems to induce some amount of overreporting in general. The same is true of the forced-choice, yes/no question as shown on the right, which displays at least as much overreporting as the check-all question. Moreover, as the 95% confidence intervals around the Loess curves illustrate, the counts of reported and actual exposure in the open-ended condition are statistically indistinguishable from each other. This is not the case in the check-all and yes/no conditions, illustrating a general tendency toward overreporting.²⁰

¹⁷Balance checks and summary statistics are reported in Table A1 of Appendix A3.

¹⁸These descriptive findings are shown in Figure A1 of Appendix A3.

¹⁹That there is substantial overlap between the Reddit and MTurk communities will not surprise anyone who has spent time in either one.

²⁰Figure A2 in Appendix A3 plots the differences between the actual and reported counts for all three conditions on the same graph.

Overall, the yes/no question seems to induce somewhat more overreporting on average. One peculiarity of these graphs is that Yahoo! News, Google News, and The Huffington Post exhibit a fair degree of overreporting in all three conditions. This is possibly because some respondents visited individual article pages rather than the main page but still counted this as exposure, an artifact of the *Link Classification Task* design.²¹

²¹Conversely, in a pilot version of this study, YouTube was the most visited site according to the *Link Classification Task* yet was rarely reported as a news source. Since this likely reflected respondents' conception of YouTube as an entertainment and not a news resource, it was kept out of Studies 1-3 altogether.






Figure 1.3: Number of respondents reporting exposure to each site (actual and reported) for each treatment condition in Studies 1 and 3. Each graph is in reverse order by true exposure, with 27 each for the check-all and yes/no conditions and the 31 top sites in the open-ended condition.




Figure 1.4: Number of respondents reporting exposure to each site (actual and reported) for each treatment condition in Study 2. Each graph is in reverse order by true exposure, with 27 each for the check-all and yes/no conditions and the 33 top sites in the open-ended condition. The magnitudes on the x-axis are lower than in Figure 1.3 because of the sample size.

Such idiosyncrasies raise the possibility that site-level characteristics systematically promote over- or underreporting. In Appendix A5, I explore whether a site's overall audience or partisan leaning (as measured by comScore data) have an effect on self-reported exposure. I find some evidence that audience size, but not partisan lean, is associated with overreporting.

Next, I report the site-level totals from Study 2, which used the browser extension in place of the *Link Classification Task*, in Figure 1.4. The results are striking: The same general pattern holds as in Figure 1.3. The measure of exposure from open-ended responses yields totals that are statistically indistinguishable from those of the direct measure (left panel), while both the check-all and yes/no questions generate overreporting. While it is unclear whether the latter causes even more overreporting, it does seem to produce the most variability in reported exposure. As in Studies 1 and 3, there is less noise in the actual measure of exposure than in the self-reported measures—in other words, error seems to be reduced regardless of whether link classification or the browser software was used to collect the data.

Individual Determinants of Misreporting

Despite the aggregate site-level results reported above, it is possible that they obscure important effects on the individual level. For example, does the open-ended prompt produce a more accurate overall picture of online news exposure because respondents are induced to report more accurately about all sources? Or is it perhaps because the question causes respondents to misreport in two different directions that cancel out?

To investigate, I report regressions in Table 1.2 modeling the determinants of individuals' propensity to overreport, underreport, and misreport exposure to online media. In all models, the check-all condition is the base case. To facilitate ease of interpretation, I use OLS with Huber-White robust standard errors to report the

Table 1.2: OLS regressions with robust standard errors in parentheses (N=1112).

	<i>Dependent variable:</i>					
	Overreporting		Underreporting		Misreporting	
	(1)	(2)	(3)	(4)	(5)	(6)
Open-Ended	-1.61*** (0.17)	-1.53*** (0.17)	1.47*** (0.38)	1.41*** (0.37)	-0.14 (0.41)	-0.11 (0.39)
Yes/no	1.82*** (0.27)	1.82*** (0.26)	0.002 (0.13)	0.002 (0.15)	1.82*** (0.28)	1.82*** (0.28)
Age		-0.01 (0.01)		0.01 (0.01)		-0.001 (0.01)
Hawaii/Pacific		1.34 (1.55)		0.84** (0.38)		2.19 (1.54)
Other Race		0.52 (0.57)		-1.58 (1.40)		-1.06 (1.47)
Native		-2.54*** (0.76)		-2.42 (1.98)		-4.96** (2.08)
Black		0.66* (0.34)		-0.01 (0.33)		0.65 (0.44)
Asian		-0.28 (0.32)		0.12 (0.27)		-0.16 (0.40)
Hispanic		-0.65* (0.34)		2.22 (2.03)		1.57 (2.03)
Female		-0.23 (0.19)		0.08 (0.25)		-0.15 (0.30)
Income		-0.01 (0.04)		0.07 (0.06)		0.06 (0.07)
Education		0.10 (0.07)		0.24* (0.12)		0.34** (0.14)
Ideology		-0.07** (0.03)		-0.01 (0.02)		-0.07** (0.04)
Party ID		0.02 (0.07)		0.10 (0.15)		0.11 (0.16)
Knowledge		0.42*** (0.09)		0.10 (0.14)		0.51*** (0.16)
Attention		-0.08 (0.07)		-0.001 (0.01)		-0.08 (0.08)
Computer		-0.36 (0.24)		0.39* (0.23)		0.03 (0.31)
Cleared		0.99*** (0.20)		-0.37 (0.24)		0.63** (0.30)
Purple		0.63*** (0.23)		-0.51 (0.59)		0.12 (0.62)
Constant	3.33*** (0.15)	2.33*** (0.58)	0.64*** (0.08)	-1.43 (1.17)	3.97*** (0.16)	0.89 (1.29)
Adj. R^2	0.17	0.22	0.02	0.03	0.02	0.04

Note:

*p<0.1; **p<0.05; ***p<0.01

findings.²²

I include dummies for the treatment conditions, in addition to a full range of demographic and political controls. *Knowledge* was constructed using the questions testing general political knowledge, and is coded from 0-4. *Attention* was an attempt to measure respondents' interest in politics and attention devoted to political news online.²³ I also include several variables that attempt to correct for any technological issues that could have inadvertently caused measurement error with the *Link Classification Task*. *Computer* is a dummy variable for subjects' response to the question, "Are you completing this survey from the computer/web browser combination you typically use to read news online, including about American politics?" This was intended to isolate the effect of respondents who reported exposure to media that they used another computer or browsing device to access. *Cleared* is a variable coded from the question, "Have you recently cleared the web history of the browser you are currently using to complete this survey (i.e. within the last 30 days)?" This captures the possibility that some users will have nearly empty web histories that could skew the actual exposure measure. And finally, *Purple* coded respondents' open-ended comments at the end of the survey, where they were asked if they encountered any problems with the survey such as not seeing any purple links in the *Link Classification Task*. This was intended to capture any other possible technological issues that could have caused the *Link Classification Task* to fail.

Several results stand out. First, the only models with any real explanatory power are those that predict overreporting. This makes sense since overreporting is the

²²Table A2 in Appendix A3 shows the same results with both quasi-Poisson and negative binomial models as robustness checks since the dependent variable is overdispersed count data. All substantive effects discussed in the text also hold in both models, while additional significant findings are not consistent between the models.

²³"How interested are you in information about what's going on in government and politics?" and "When you watch or read news on the Internet, how much attention do you usually pay to news about politics?"

most well-known problem with self-reported exposure measures. As expected, and confirming both preregistered hypotheses, the open-ended condition causes significantly less overreporting than the check-all question, and the yes/no question causes more. In terms of magnitudes, the coefficient on the open-ended condition, -1.53 , implies that the absolute level of overreporting (2.33 additional sites on average for the check-all condition) is reduced by two-thirds. The yes/no condition, by contrast, increases overreporting by an even greater magnitude. Models 3 and 4 suggest that open-ended questions also encourage more underreporting, which also makes sense given the cost of recall. This is also a likely reason why the open-ended question, despite significantly less overreporting, does not perform any better than the check-all question when it comes to misreporting in general.

It appears that a number of respondents had either recently cleared their browser histories or experienced some trouble displaying the colored links. In particular, those with recently cleared histories were far more likely to “overreport,” although that result is an artifact of measure construction—the actual exposure for these respondents was artificially low.

A few demographic characteristics stand out as significant explanatory factors. Echoing a well-known finding in self-reports of voting (Belli, Traugott and Beckmann 2001), respondents who were more knowledgeable about politics were more likely to overreport. In addition, respondents with more education were more likely to misreport exposure, a finding that seems to be driven more by underreporting. Finally, more conservative respondents were less likely to overreport exposure than more liberal respondents.

Predicting News Reception

If a particular measure of media exposure is to be useful beyond a mere accounting of “who visited what website when,” it should be able to predict respondents’

ability to report political events discussed in the news. In a classic treatment of news reception, Price and Zaller (1993) propose and validate general knowledge as a predictor of news recall, arguing that it performs better as an “indicator of a general propensity for learning about news events” (138). In Study 3 (N=700), I incorporated questions about recent “news events,” structured in the same fashion as the questions examined in the original paper. This allowed me to investigate whether any of the particular survey-based measures perform as well as, or better than, general knowledge.




Figure 1.5: News recall broken down by individual story and predictor. The x-axis plots, from left to right, general knowledge, reported media exposure, and actual media exposure as measured by the *Link Classification Task*. The Colorado story is in green, the *Duck Dynasty* story is in pink, and the unemployment benefits story is in brown.

The three events I asked about—Colorado’s legalization of recreational marijuana, *Duck Dynasty* star Phil Robertson’s controversial remarks about gays and African Americans, and the expiration of extended unemployment benefits—varied in terms of their overall resonance. To illustrate how general knowledge, reported media exposure, and actual media exposure relate to how well respondents were able to recall these stories, I recreated Figure 1 from Price and Zaller’s article in Figure 1.5 above. The x-axes show the level of general knowledge, reported exposure, and actual exposure; for the latter two, I summed across sites for each respondent and recoded the variables to be on the same scale as knowledge, from 0 to 4. The y-

axes plot the percentage of respondents in each category who were able to correctly answer the question about each story.

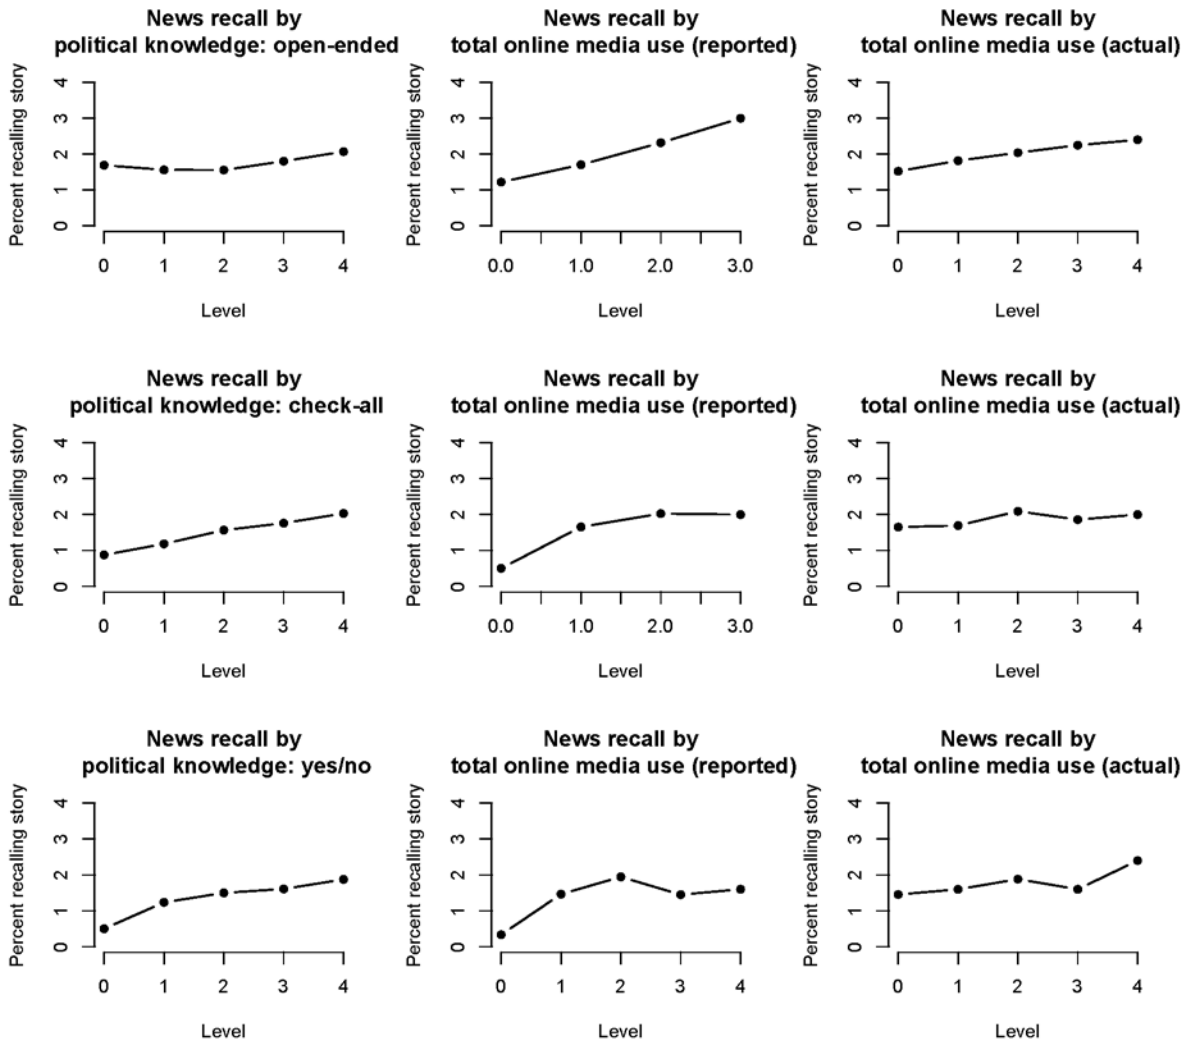


Figure 1.6: News recall (composite measure) broken down by predictor. The x-axis plots, from left to right, general knowledge, reported media exposure, and actual media exposure as measured by the *Link Classification Task*. The y-axis plots the average number of stories recalled, out of 3. Each row represents a different treatment condition.

The predictive power of general knowledge is clear: As knowledge increases (from 0 to 4), the percentage of respondents who recall the marijuana story increases from nearly 52% to nearly 75%, and the percentage who recall the *Duck Dynasty* story increases from 48% to 74%. The unemployment issue is less well

known, but recall of that story increases as well, from under 19% to over half. Reported media use also shows a clear association with recall, especially with the marijuana story, but the relationship is not monotonic. Moreover, for *Duck Dynasty* there is no positive association at all. Meanwhile, the tally of actual media use has a positive but more modest association with news recall.

To investigate the performance of the exposure measures, I first created an index of overall news reception by summing up the indicators for each news story. Figure 1.6 plots this index against the same predictors, with each treatment group on a different row.²⁴ From this picture, reported open-ended exposure (top middle graph) clearly shows the strongest association with news recall—stronger than either general knowledge or *actual* use. Table 1.3 regresses this index on total reported exposure and general political knowledge separately within each treatment group. In the middle column, the coefficient on *Exposure* is statistically significant and larger than that for *Knowledge*: For respondents randomly assigned to the open-ended prompt, the total number of websites they reported being exposed to predicts news reception over and above the level of general political knowledge. The exposure measure in the check-all condition is also significant, but its effect is more modest. Also notable is that for those assigned to the open-ended group, the *Attention* measure behaves as we would expect, reducing overall news reception.

Discussion

The results presented here will be of considerable use to scholars seeking to improve research designs in both observational and experimental settings—especially given that data collection online is increasingly popular for both practical and substantive reasons (Iyengar 2010). Taken together, the findings of the three studies

²⁴Appendix A3 contains a version of this graph in Figure A3 with the recall measures for each story kept separate.

Table 1.3: OLS with Huber-White robust standard errors in parentheses.

	<i>DV = News reception (0-3)</i>		
	Check-all	Open-ended	Yes/no
	(1)	(2)	(3)
Exposure	0.05** (0.02)	0.14*** (0.03)	0.01 (0.02)
Knowledge	0.21*** (0.06)	0.01 (0.05)	0.14* (0.07)
Attention	-0.11 (0.07)	-0.14*** (0.05)	-0.14* (0.07)
Age	0.001 (0.01)	-0.001 (0.01)	0.01* (0.01)
Pacific		-0.01 (0.15)	0.91*** (0.30)
Other	0.73*** (0.26)	-0.08 (0.42)	-0.43 (0.48)
Native	-0.88*** (0.22)		
Black	-0.52** (0.22)	-0.12 (0.34)	-0.03 (0.22)
Asian	0.15 (0.24)	-0.12 (0.19)	-0.18 (0.23)
Hispanic	-0.31 (0.19)	-0.12 (0.27)	0.14 (0.30)
Female	0.12 (0.14)	0.13 (0.12)	-0.10 (0.12)
Income	-0.03 (0.02)	0.03 (0.02)	-0.04 (0.03)
Education	0.02 (0.04)	0.05 (0.05)	0.12** (0.05)
Ideology	0.05 (0.04)	-0.09*** (0.04)	0.002 (0.003)
Party ID	0.06 (0.05)	0.08* (0.05)	-0.04 (0.04)
Computer	0.30* (0.16)	0.08 (0.16)	-0.47*** (0.14)
Cleared	-0.003 (0.12)	0.09 (0.12)	0.21* (0.13)
Purple	-0.22 (0.14)	0.04 (0.14)	0.02 (0.13)
Constant	0.67 (0.51)	1.45*** (0.37)	1.29*** (0.43)
<i>N</i>	216	253	231

Note:

*p<0.1; **p<0.05; ***p<0.01

converge on the conclusion that among survey-based measures of online media exposure, open-ended questions perform the best. Evaluated against baseline measures of true exposure constructed in two different ways, they produce an accurate snapshot of aggregate online media use. They also encourage significantly less overreporting of exposure to specific websites, although this is partially balanced out by more underreporting. As expected, forced-choice, yes/no questions cause even more overreporting than check-all questions, and both result in an inflated picture of the overall audience for online political media.

Holding actual exposure constant, the kinds of sources that tend to get reported most are those with the greatest reach. This is likely both because those sources are more likely to be listed as a choice to be selected in closed-ended questions and because they are more likely to be easily retrieved from memory.

On the individual level, more knowledgeable respondents are more likely to overreport exposure to political media, and there is some evidence that more conservative respondents are less likely to overreport. Without further research on a representative sample, it is unclear whether these findings generalize to the population, but they raise the possibility that there are demographic and political correlates of misreporting behavior on surveys of media exposure.

In addition to performing better in terms of overreporting and aggregate exposure, open-ended questions predict reception of recent political news stories better than Price and Zaller's knowledge measure, at least in the sample analyzed here.

In general, however, it is important to remember that the absolute level of misreporting is high, relative to overall exposure, regardless of the survey-based method used. Fortunately, the methods demonstrated here point a way forward. In applications where direct behavioral measures of Internet-based activity are needed, online surveys can be equipped to present respondents with a version of the *Link Classification Task* tailored to a particular universe of sites. As demonstrated in the

validity test, this procedure generates nearly error-free data on actual online exposure without the privacy concerns of a software extension. If deployed as part of a representative sample online, the method could shed light on exactly what types of people are more or less likely to report using a particular source.

Classifying links based on appearance is a relatively straightforward and un-intrusive way of obtaining data on subjects' online browsing behavior for a given domain. Its main disadvantage is inconvenience to respondents (which increases linearly with the number of URLs included in the list). In research designs where using this procedure is infeasible, researchers have two options. First, they can opt for survey-based measures: open-ended responses that are coded, analyzed, then possibly adjusted for known biases documented in representative data. This is an especially attractive option for studies not focusing entirely or at all on Internet-based media, since open-ended queries are endlessly adaptable.

Second, researchers can consider using the browser extension method. While Study 2 illustrated the tradeoff between measurement error and response rate with this design in the survey context, it may be better suited for more controlled settings. In laboratory experiments or computer-administered online surveys in which researchers have some degree of control over the equipment, this may actually be the least intrusive option. And most compellingly, this approach could be harnessed to tap into more refined measures of actual browsing behavior, such as logging not just whether and when users have visited a given URL, but how many times. This "dosage" measure would help separate substantively important habitual media habits from one-off visits, which the current method cannot distinguish between (Mutz and Young 2011).

When Treatments Are Tweets: A Network Mobilization Experiment Over Twitter

WITH ALEXANDER COPPOCK AND JOHN TERNOVSKI

While much enthusiasm about the Internet focuses on its ability to foster informal and decentralized forms of organization (Shirky 2008; Benkler 2006; Bennett and Segerberg 2012), traditional groups have long recognized its potential for recruitment and mobilization (Obar, Zube and Lampe 2012). This is especially true in the realm of politics: With meaningful political behavior now commonplace online, campaigns have added email and Facebook appeals to their arsenal of tactics (Krueger 2006; Gaby and Caren 2012). Nonpartisan and advocacy organizations have similarly turned to social media to engage their supporters.

This study examines the effects of an online mobilization campaign via what we believe to be the first randomized field experiments conducted on the social network Twitter. Our design allows us to identify the effects of both private messaging and “natural” network behavior directed toward supporters of a nonprofit

advocacy organization, the League of Conservation Voters (LCV). The two field experiments we present here follow nearly identical designs and lead us to draw very similar conclusions about political mobilization over Twitter.

Our primary manipulation in both experiments exposes some subjects to a public tweet only and others to one of two private direct messages. The wording of the direct messages primes subjects with one of two identities based on previous research demonstrating the behavioral consequences of associating political actions with a particular self-concept (Bryan et al. 2011). With this manipulation, we investigate whether the passivity associated with a “followers” label or the higher level of commitment associated with an “organizers” label has an impact on our two outcomes: signing an online petition and tweeting (or retweeting) the petition link.

Our secondary manipulation encourages a random subset of petition signers to tweet the petition to their own followers. This has the consequence of randomly assigning the followers of petition signers to be exposed to tweets directing them to the petition. This design allows us to explore network effects while avoiding homophily concerns (McPherson, Smith-Lovin and Cook 2001).

We find that private direct messages on Twitter are highly effective tools for generating online petition signatures, a common advocacy goal. In addition, we find no evidence that the type of identity primed affects the likelihood of signing an online petition. However, those assigned to the “follower” condition are more likely in both experiments to tweet a link to the petition to their own followers. Results from our secondary manipulation are mixed: In our second experiment but not the first, we find evidence of network effects among followers of the organization’s own followers. Finally and most surprisingly, no one who was exposed to only the public tweet either signed the petition or tweeted the link to their own followers.

In addition to providing practical guidance for organizations with dedicated

follower networks, these results suggest that the advantages of personal appeals identified in face-to-face campaigns can carry over into the virtual world (Rosenstone and Hansen 1993; Gerber and Green 2000). They also show that invoking identities associated with different levels of commitment to a cause can affect one's propensity to comply with a simple request. This evidence has potential implications for the continuing debate on whether online campaigns of the type studied here merely promote "slacktivism"—token, low-cost emblems of support incapable of sustaining meaningful collective action (Morozov 2009; Gladwell 2010).

This paper proceeds as follows. In the following section, we review recent theoretical arguments on whether social media can facilitate or hinder collective action. We then outline the potential pitfalls of analyzing experimental manipulations over social networks like Twitter. In the next section, we provide an overview of the universe of subjects—the follower network of a large nonprofit advocacy organization—and place it in context. Then, we describe the research design and analytic strategy of the two experiments. The three subsequent sections report the results of both experiments in addition to the heterogeneous effects of treatment by account type. The final section concludes with a discussion.

Participation, Collective Action, and Social Media

Scholars have sought to understand how lowered communication costs and the proliferation of online social networks have altered the traditional logic of collective action (Olson 1965). One set of responses focuses on the scope of action: The relatively low-effort individual contributions involved in online campaigns necessarily limit the value of resulting public goods (Gladwell 2010; Shulman 2009). In contrast to this pessimistic assessment, some communication scholars—basing their insights on grassroots anti-globalization campaigns or, more recently, events leading up to the Arab Spring—delineate how new communication technologies can

foster real-world spontaneous collective action in the absence of formal organization or leadership (Bennett and Segerberg 2012). These theories extend Olson's classic work to networked, often digitally mediated contexts.

Little theoretical work thus far explicitly invokes Twitter (see Marwick and boyd 2011 for an exception), but many of the ideas carry over from discussions of viral email campaigns. For example, Bimber, Flanagin and Stohl (2005) associate the act of contributing to a collective good with a transition across the boundary from private to public. However, because the new technologies often blur the line between public and private, they find that "boundary crossing in connection with public goods takes on forms not so readily recognizable in the theoretical terms of free riding, selective incentives, and organization" (p. 378). In this alternative framework, forwarding a petition (or, perhaps, retweeting a message) is a "nearly costless request" to make private or semiprivate information—that is, the fact that someone supports a campaign—public.¹

Another perspective, championed by Benkler (2006) and Bennett and Segerberg (2012), emphasizes how networks enable the co-production of collective goods, fostered via individual self-expression. Prototypical examples in this tradition include the open-source software movement and Wikipedia, relatively decentralized networks of contributors who help to create and maintain free public resources. In the political realm, the movement against the Stop Online Piracy Act (SOPA) or for net neutrality might fall in this category of "connective action," although the type of mobilization studied here falls more clearly under the traditional umbrella of "organizationally brokered networks" (Bennett and Segerberg 2012, p. 756). As we demonstrate in these experiments, networked participation can arise from centrally organized campaigns as well.

¹Sharing private information with third parties is now an ingrained part of online behavior, with potential benefits for both traditional advocacy groups and networked movements.

Classic treatments of participation in politics understandably focus on factors affecting the likelihood that different types of citizens will vote, volunteer for a candidate, or contribute in some other way on behalf of a political cause. This study necessarily approaches the topic differently, focusing on people who have already self-selected as followers of a prominent environmental advocacy organization. In effect, we condition on the usual determinants of participation as presented in the Civic Voluntarism Model (Verba, Schlozman and Brady 1995): resources, motivation, and intensive engagement with a political issue. Accordingly, our experiments focus on the model's remaining ingredient, "networks of recruitment" (p. 3).² Although Verba, Schlozman and Brady assumed that such networks would consist of commonly studied social ties such as friends, family, and co-workers, we extend this notion to include connections based on affinity.

Recruitment networks are a useful concept for studying mobilization in an online context. Aside from capturing the possibility of "viral" or network patterns in campaign activity, they allow us to compare the effectiveness of appeals that originate from the organization itself with those from non-affiliated but like-minded supporters of its environmental mission. This distinction between direct appeals and peer effects is important for at least two reasons. The first relates to Olson's original observation about the "noticeability" of individual contributions, which implies that as organizations become large, shirking becomes unobservable to other members and free-riding inevitable (barring selective incentives, coercion, or other inducements to cooperate). Peer effects facilitated via transparent social networks, by contrast, are one possible way in which online organizing could overcome the problem of "noticeability" and mitigate the incentives to free-ride in large groups (Lupia and Sin 2003).

²The organization-centered design also addresses any concern about our lack of covariates for these traditional predictors of participation, although we analyze differential effects by factors that we were able to capture, gender and organizational status.

Second, the distinction matters because social ties may reflect group membership. This insight arises from the Elaborated Social Identity Model, which was constructed to explain how group membership can induce collective action (Drury et al. 2005). The theory presupposes the existence of a grassroots in-group and a powerful out-group; any action taken by the in-group against the out-group that appears to succeed is “experienced as joyful and exhilarating” (Barr and Drury 2009, p. 245). This model has been applied to induce voter mobilization by labeling targets as “voters” rather than simply people who vote (Bryan et al. 2011). “Donor” and “activist” identities have also been associated with increased charitable donations and activism, respectively (Aaker and Akutsu 2009), though no published large-scale field experiments have evaluated the impact of the “donor,” “organizer,” or “activist” identity labels on those outcomes.

With our design, elaborated below, we simultaneously address several of the theoretical debates raised in the literature. First, we measure the effects of our manipulations on two primary outcomes: filling out (or “signing”) an online petition, and tweeting (or retweeting) a link to the petition to one’s own followers. The former is recognizable as a contribution to a public good, one traditionally valued as part of the political organizer’s toolbox (Karpf 2010). The latter is a somewhat more ambiguous—but arguably less costly—action that, if repeated by many other members, could lead to increased public awareness of the campaign (and its magnitude of support). An important difference between the two outcomes is that tweeting more directly captures whether the “noticeability” of the behavior leads to enhanced effectiveness, allowing for a test of how Olson’s logic may operate differently when mediated via online communication networks. Second, we examine the effects of three types of appeals: Generic appeals via the organization’s public Twitter account, specific appeals via private direct messages, and tweets from followers of the organization to their own followers. These differ in the extent to

which they depend on social ties, direct contact, and the authority of a trusted organization. Finally, we vary the identity labels used to address our subjects, enabling us to examine whether the salience of specific social identities is associated with the likelihood of contributing to the organization's goals.

To briefly summarize the expectations of the participation and online collective action literature to date, we believe that the work of Bimber, Flanagin and Stohl (2005) would lead to a prediction of larger treatment effects for the (re)tweeting outcome, which involves the relatively low-cost act of making one's support for a cause (more) public. The "slacktivism" hypothesis, by contrast, straightforwardly predicts strong treatment effects for the lowest-cost actions regardless of whether they make information public. This leads to the expectation that subjects across all conditions will retweet LCV's public tweet but that fewer will take the time to sign the petition (and subsequently tweet out the link to it). Finally, the Civic Voluntarism Model (Verba, Schlozman and Brady 1995) predicts that "networks of recruitment" will be most effective at mobilizing followers: The largest treatment effects will be observed when subjects are exposed to tweets from peers.

The Challenges of Experiments on Twitter

In the terminology of network analysis, the social microblogging service Twitter is a directed graph. Users post short, public updates and curate their own networks by "following" others—friends and strangers alike—who may or may not reciprocate. A particular user's Twitter messages, or "tweets," can be read by anyone who visits his or her public feed (also known as a timeline). Since manually reading individual feeds can be cumbersome, users typically take advantage of the Twitter "stream," a real-time aggregation of tweets from users they follow. The result is a never-ending rush of text and photo updates from sources of a user's choosing.

There are several specific ways of communicating on Twitter. Most fundamental

is the tweet, usually restricted to 140 characters (with exceptions for web addresses of reasonable length). “Retweets” (RTs) allow users to quickly resend a tweet from their stream to those of their own followers. (In other words, a retweet copies a tweet from a user’s incoming stream to his or her own feed, with attribution.) In extreme cases, this capability can lead to cascades of retweets of particularly compelling content. Other tweets, while not necessarily retweets, can “mention” another user. This option can lead to extended public conversations, all potentially referring back to an initial Twitter posting. Finally, while Twitter is best known for its public functionality, it also allows users to send private “direct messages” (DMs) to any of their followers. By default, Twitter sends users an email notification when they receive a DM.

Public tweets comprise the bulk of a typical Twitter account’s activity, yet they present challenges for the design and analysis of experiments. Since public tweets and retweets are visible to followers of the sender, a simple experimental design in which some followers are randomized to be shown a public tweet and others in a control group are not is effectively impossible. An alternative design would randomly time a series of public tweets with different messages, but any causal inferences would require strong modeling assumptions concerning the over-time persistence of treatment effects (for example, if the same message is tweeted every other day, followers may become irritated and respond differently from how they otherwise would).

Instead, we take advantage of Twitter’s direct message capability, which allows us to present different messages to different users. Estimation of the relative effectiveness of the messages can proceed in the normal fashion, under an assumption of non-interference between units. The non-interference assumption (sometimes referred to as the Stable Unit Treatment Value Assumption, or SUTVA) requires that subjects’ outcomes not be influenced by the treatment assignments of other

subjects. The interference concern is not trivial: Unmodeled spillovers can lead to biased estimates of treatment effects (Gerber and Green 2012, Chapter 8). For example, if direct messages were highly effective at motivating petition signatures and subsequent tweets, but those tweets exposed subjects in the control group to the same message, a naive difference-in-means estimate would be biased.

A schematic version of an analytic approach to dealing with spillovers of this kind is as follows: First, redefine treatment categories to include “spillover conditions” such as being in the condition of following one subject who received a direct message. Second, calculate the probability that each unit is in each redefined treatment condition. Because Twitter users follow vastly different numbers of other users, these probabilities will vary quite a bit from unit to unit. Third, weight each unit’s outcome by the inverse of the probability of being in its observed condition. Average differences across these redefined treatment categories will reflect unbiased treatment effect estimates.

The trouble with this approach is the prohibitively large number of potential treatment categories: anywhere from following zero treated units to following 601 (the largest out-degree observed in our network). One could instead parameterize the “dosage” of spillovers and estimate a response curve for each extra treated unit (see Bowers, Fredrickson and Panagopoulos 2013 for the method, Coppock 2014 for an application). Such a method, however, requires the researcher to make strong functional form assumptions concerning exposure.

Our solution is to vastly reduce the number of potential spillover conditions by design. As discussed below, we randomly induce a relatively small subset of the network to retweet to their followers, which in turn randomly exposes some users but not others to the petition link.³ Most users only follow one user in this

³Contrast this with the case in which the organization is sending the public tweets—all of their followers are potentially exposed. When a random subset of users tweet the link, however, only a portion of the organization’s followers are exposed, allowing for experimental differences to be

subset, so the number of spillover conditions is quite manageable. Additionally, the variability of exposure probabilities is kept in check. In effect, this approach allows us to randomize “natural” Twitter behavior and directly estimate its consequences.

Because of the difficulties outlined here, it is not surprising that randomized experiments on Twitter have been rare. One study randomly encouraged subjects to follow a Japanese politician on Twitter to test effects on trait evaluations, knowledge, and other post-treatment outcomes (Kobayashi and Ichifuji 2014). A marketing study tested the effectiveness of tweets and retweets on television ratings by collaborating with a media company, as well as “influential” tweeters, to promote a random subset of its TV shows on the Chinese social network Sina Weibo (Gong et al. 2014). The unit of analysis in this design was the shows themselves. Finally, there is ongoing research on the relationship between identity and behavior on social media. A randomized experiment on a web-based social sharing site (but not Twitter) found that cues indicating an account’s identity matter in terms of how users share content associated with that account (Taylor, Muchnik and Aral 2014).

Overview of Network

The League of Conservation Voters is an environmental advocacy organization that “works to turn environmental values into national, state and local priorities,” according to its official website. Its activities include public awareness campaigns, lobbying efforts, and independent expenditures (via political action committees) geared toward electing candidates who support its agenda.

The two field experiments analyzed here were deployed over the network of followers of LCV’s official Twitter account, @LCVoters. The members of this network comprise an “issue public” in the literal sense (Converse 1964; Verba, Schlozman and Brady 1995): highly dedicated to promoting the environment and publicly vis-

observed.

ible in their activism. Studying the impact of mobilization tactics on such a network speaks to the effectiveness of promoting activism among politically engaged individuals, although any generalizations to other populations would have to be qualified.

While the network is highly engaged, it is not a close-knit community. Rather, it is mainly a network of strangers. Members all share a common interest in the environment, but there are relatively few interconnections between them: Out of a possible 44,709,282 connections between nodes, there were 131,474 such edges when the network was scraped before Study 1, yielding a graph density of 0.0029. Put more simply, followers of the organization follow a median of only six users in the organization's network. This structure corresponds most closely to a "Broadcast Network," as described by the Pew Research Center's typology of conversational archetypes on Twitter.⁴ Such "hub and spoke" networks consist of an audience of followers who typically rebroadcast (i.e., retweet) the output of a single source, in this case LCV.

An additional feature of the network, as Figure 2.1 illustrates, is its relatively diffuse nature. Unlike a highly modular social network with various distinct groupings, this one has numerous and overlapping communities that are difficult to distinguish from each other. Network statistics confirm this impression: The Walktrap algorithm (with standard defaults) finds 22 communities in the network, for a modularity of 0.3.

Describing a network with statistics such as the graph density and the modularity is usually insufficient for conveying its structure. Modularity, for example, depends on the number of communities detected; different algorithms come to different conclusions about the number of communities and their membership. In our

⁴See <http://www.pewinternet.org/2014/02/20/mapping-twitter-topic-networks-from-polarized-crowds-to-community-clusters/>.




Figure 2.1: An illustration of the network of 6,687 followers of the advocacy organization's Twitter account scraped before Study 1. Lines illustrate connections between users and are shaded by membership in one of 22 communities as determined by the Walktrap community detection algorithm.

view, the principal utility of these statistics is to provide partial justification for our description of the subject pool as a being a social community only in the loosest sense.

For Study 1, we constructed our universe of subjects by scraping the Twitter ID numbers of LCV's followers, excluding those who had more than 5,000 followers of their own. The reasoning behind this decision was that users with especially large numbers of followers were more likely to be prominent individuals or organizations whose online behavior would not be comparable with the rest of the subject pool. The resulting network contained 6,687 members. For Study 2, conducted five months later, we repeated this procedure and obtained a network with 8,507 members.




Figure 2.2: The number of tweets and retweets sent by LCV per day, from February 2013 to February 2015.

LCV’s Twitter account is fairly active, sending out an average of 6.07 tweets and retweets per weekday from February 2013 to February 2015. However, as Figure 2.2 shows, day-to-day variation in the number of tweets posted is high (s.d. = 10.05). This activity has not stopped the network from continuing to grow, suggesting that its followers, in addition to being dedicated to the cause, are accustomed to frequent Twitter updates from the organization.

This is likely also the case for comparable organizations. Table 2.1 lists the number of followers and average number of tweets per weekday for the top 10 most influential environmental organizations (by 2014 lobbying expenditures as collected by OpenSecrets.org). The mean number of followers among this group, 118,988, is an order of magnitude higher than LCV’s number of followers, and average tweet frequency—about 13 per weekday—is slightly more than twice as high. Finally, the median creation date for these organizations’ Twitter accounts was mid-2008, fairly early in Twitter’s history.

Table 2.1: Descriptive statistics (circa February 2015) of Twitter accounts of the top 10 environmental organizations by 2014 lobbying expenditures, plus LCV. “T/W” = the number of tweets per weekday and counts retweets.

Organization	Account	Followers	T/W	Created
Earthjustice Legal Defense Fund	@Earthjustice	67,004	8.60	2008
Environmental Defense Fund	@EnvDefenseFund	93,619	6.87	2009
Environmental Working Group	@ewg	35,279	7.65	2008
Intl Assn of Fish & Wildlife Agencies	@fishwildlife	1,229	1.30	2010
League of Conservation Voters	@LCVoters	10,880	6.07	2009
Nature Conservancy	@nature_org	387,973	9.60	2008
National Parks Conservation Assn	@NPCA	108,164	11.47	2009
Natural Resources Defense Council	@NRDCFedGov	681	2.10	2012
National Wildlife Federation	@NWF	296,513	14.60	2007
Sierra Club	@sierraclub	143,448	42.67	2009
Wilderness Society	@Wilderness	55,971	28.07	2007

Research Design

We conducted two field experiments with nearly identical designs; lessons learned from the first experiment improved the design of the second. In both studies, LCV first posted a public tweet urging supporters to sign an online petition and retweet the link to their own followers. In a first-stage experiment, subjects were randomly assigned to one of three groups: (1) the baseline or control group, which was exposed to the public tweet only; (2) a condition in which subjects also received a DM with a similar request, referring to them as “followers”; (3) a condition in which subjects also received a DM referring to them as “organizers.” In a second-stage experiment, those who completed the petition were randomly shown a link with an encouragement to tweet the petition to their own followers. We will refer to this treatment as the “tweet encouragement” or the “tweet link.”

Study 1

In Study 1, LCV’s tweet and petition were related to an ongoing campaign to end tax breaks to “Big Oil.” The public tweet, posted on February 5, 2014, was followed

by the DMs,⁵ which had to be sent in 12 daily batches.⁶ (See Appendix B4 for the full text of all messages.) In this version of the design, subjects who completed the petition were required to enter their Twitter usernames in order to connect responses back to treatment assignment.

The initial randomization procedure assigned one-third of the subjects in each of the DM conditions to be shown a tweet encouragement after submitting the petition signature, using complete random assignment.⁷ In total, subjects could be assigned to one of five treatment conditions. We collected outcome data in two concurrent ways: Online petition signatures were collected by survey software, while tweet and retweet behavior was captured by scraping programs that we ran continually for the duration of the experiment. We ran multiple scrapers in parallel to guard against accidental data loss.⁸

Table 2.2 shows the number of subjects within each condition and the proportions signing the petition and tweeting the petition link to their followers. Table 2.2 also reveals an anomalous finding that suggests a potential issue with the randomization: Assignment to be shown the tweet encouragement predicted petition signatures ($p < .01$). Since the encouragement was only displayed to subjects *after* signing the petition, it is possible that randomization failed to eliminate unobserved differences between subjects assigned and not assigned to the tweet encouragement condition. We exhaustively investigated the possible sources of this

⁵“Follower” condition: “You’re one of our most valuable followers! Please RT this petition to your friends to stop tax breaks to Big Oil. [URL to petition]”; “organizer” petition: “You’re one of our most valuable organizers! Please RT this petition to your friends to stop tax breaks to Big Oil. [URL to petition]”

⁶Twitter’s API limits the number of DMs that an application can send to 250 per day. Subjects were randomly selected into batches.

⁷This was done by using two different versions of the petition. See Figure B1 in Appendix B4 for a screen shot of the encouragement, and Figure A2 for the tweet window that popped up if a user clicked on the tweet encouragement.

⁸By “scraping” we mean continually querying the Twitter API in order to capture tweets containing the URLs to either version of the petition. We used this approach because Twitter data is much easier to collect in real time than after the fact.

Table 2.2: Study 1: Design and Outcomes

Treatment Group	N	Signed	Tweeted
Public Tweet	3687	0 (0.0%)	0 (0.0%)
Organizer DM			
Tweet encouragement	500	22 (4.4%)	12 (2.4%)
No encouragement	1000	28 (2.8%)	12 (1.2%)
Follower DM			
Tweet encouragement	500	28 (5.6%)	25 (5.0%)
No encouragement	1000	31 (3.1%)	16 (1.6%)
Total	6687	109 (1.6%)	65 (1.0%)

imbalance, such as day-of-week effects, faulty randomization procedure, and data problems, but we were unable to conclusively pinpoint the cause. The most plausible explanation is that an imbalance occurred simply by chance. In Study 2, we addressed this problem by waiting until subjects clicked through to the online petition to conduct the second-stage random assignment.

In our analysis of the first-stage experiment (the direct message treatments), we will first examine average differences across treatment assignments, with and without covariate adjustment. Using information from users' public Twitter profiles, we were able to gather the following covariates: account type (male, female, organization, or unknown), number of followers, and the number of days the account was open.⁹ We also calculated each subject's eigenvector centrality, a measure of how well-connected the user is within the LCV network. For our analysis of the first-stage experiment, we will rely on a strict assumption of non-interference among units. Table 2.2 contains some indication that this assumption is not wholly unwarranted: Despite 65 tweets of the link throughout the network, a grand total of zero subjects in the public tweet condition signed the petition. At least among this subset, we can be sure of the non-interference assumption. (See also Sinclair, McConnell and Green 2012 for evidence that the non-interference assumption is

⁹See below for the details of this procedure.

well-justified in get-out-the-vote mail experiments.)

The subjects in the second-stage experiment (the encouragement to retweet) were those who met the following criteria: They followed both LCV and users who signed the petition. The petition received 109 signatures; these 109 users were followed by a total of 1,176 other LCV followers. The 1,176 were the pool of subjects randomly assigned to the condition of following someone exposed to the tweet encouragement. Similar to the procedure described earlier, we will weight each observation by the probability of exposure, as those who follow more of the 109 are more likely to be exposed. We will estimate both the intent-to-treat (ITT) effect using ordinary least squares (OLS) and the complier average causal effect (CACE) using instrumental variables (IV). The definition of a complier in the second-stage experiment is a mouthful: a complier is a user who follows one or more petition signers who tweeted the link if and only if shown the tweet encouragement. Though it may seem counterintuitive, the analysis of the second-stage experiment also assumes non-interference—we assume that a unit’s outcomes do not depend on whether or not some other unit follows a petition signer assigned to the tweet encouragement.

One may wonder about users who do *not* follow LCV but may still have been exposed to tweets by virtue of following a petition signer. As it happens, exactly zero petition signatures and subsequent tweets were recorded for non-followers of LCV, providing evidence that among that subsample, our manipulation had no effect on these outcomes. There is one significant exception to this finding, however: In Study 1, 5 users not in the LCV’s network *retweeted* either the public tweet or a tweet from one of the followers. In Study 2 this number was 7. With over 7 million total followers of the LCV’s followers, these magnitudes are minuscule, but they are greater than zero.

Study 2

We implemented a nearly identical research design in Study 2, with some minor improvements to simplify analysis and address the randomization concern described above. To ensure successful randomization in the second stage of the design in Study 2, we used simple random assignment of the tweet encouragement within the survey software (Qualtrics) itself. The three main treatment conditions remained unchanged from Study 1. For those treatments, we used block random assignment by day and number of followers.

We further made two changes to the way the web links (URLs) to the petition worked. First, we were able to incorporate the abbreviated version of the organization's name (LCV) into the URLs themselves in order to boost realism. Second, we passed on the anonymized Twitter IDs of each subject as a query to each URL sent in the DMs so that we could more easily merge individual-level outcomes with treatment assignment.

In Study 2, the subject of the Twitter campaign was more timely: the Environmental Protection Agency's plan to issue regulations mandating reduced carbon emissions from power plants. The public tweet was posted on July 2, 2014, and DMs were sent in 20 batches beginning that day.¹⁰ Between them, the 221 petition signers were followed by 1,990 other users, who constitute the subjects of the second-stage experiment in Study 2. Table 2.3 summarizes the design and basic outcomes of Study 2.

¹⁰"Follower" condition: "You're one of our most valuable followers! Help fight climate change by signing the petition & tweet to your friends! [URL]"; "organizer" condition: "You're one of our most valuable organizers! Help fight climate change by signing the petition & tweet to your friends! [URL]"

Table 2.3: Study 2: Design and Outcomes

Treatment Group	N	Signed	Tweeted
Public Tweet	3495	0 (0.0%)	0 (0.0%)
Organizer DM	2514	107 (4.3%)	28 (1.1%)
Follower DM	2498	114 (4.6%)	36 (1.4%)
Total	8507	221 (2.6%)	64 (0.8%)
Among Subjects Who Signed Petition			
Tweet encouragement	111	111 (100.0%)	50 (45%)
No encouragement	110	110 (100.0%)	11 (10%)
Total	221	221 (100.0%)	61 (27.6%)

Results: Study 1

Study 1's results challenge the conventional wisdom about Twitter's mobilization capabilities on at least two fronts. Perhaps most surprisingly, not a single subject in the public tweet condition either signed the petition or retweeted the petition link. It is important to reiterate that subjects were exposed to a single public tweet, so this result does not rule out the possibility that a more concerted campaign with multiple tweets might have worked.¹¹ Further, it is possible that infrequent Twitter users never saw the tweet at all. The ineffectiveness of the public tweet stands in contrast to the strong showing of the direct messages. Without prior research to guide our expectations, we would not have been surprised at either a null finding or a negative "backlash" effect.¹² One alternative interpretation of these results is that DMs are more effective due to repeat exposure: Subjects may have responded to the DMs because they had already seen a public tweet featuring the same message, but a tweet alone is not enough to drive outcomes.

¹¹Suppose that the true treatment effect is that a public tweet generates a single click per 10,000 followers exposed. With a sample size of 6,687, we would expect to observe zero clicks about 51% of the time.

¹²These also seemed like plausible results *ex ante*; during a pilot study in which I sent automated direct messages to a subset of my followers, several recipients warned about "spam" or a possible virus.

A final finding, but one we interpret with caution given the potential imbalance, is that the “organizer” condition caused subjects to send significantly fewer tweets using the randomly assigned tweet encouragement. We detail these findings and discuss the absence of network effects below.

Main Effects

First, we look at petition signatures as the outcome of interest. As the first two columns of Table 2.4 show, the “follower” and “organizer” treatments both had positive and significant effects at the $p < .01$ level. The follower DM caused an estimated 3.9-percentage-point increase in the proportion of subjects who signed the petition. The organizer DM caused a 3.3-percentage-point increase in participation; these effects are not significantly different from each other ($p = 0.38$). Thus, the best interpretation of the evidence from Study 1 is that receiving any direct message (after potentially seeing a similar public tweet) caused a 3.6-percentage-point increase in petition signing.

Table 2.4: Study 1: Effects of Direct Message Treatments on Participation and Tweeting

	Signed		Tweeted	
	(1)	(2)	(3)	(4)
Treatment: Follower	0.039*** (0.005)	0.040*** (0.005)	0.027*** (0.004)	0.027*** (0.004)
Treatment: Organizer	0.033*** (0.005)	0.033*** (0.005)	0.016*** (0.003)	0.016*** (0.003)
Account Type: Male		-0.004 (0.004)		0.001 (0.003)
Account Type: Organization		-0.018*** (0.004)		-0.003 (0.003)
Account Type: Unknown		-0.012* (0.007)		0.004 (0.007)
Eigenvector Centrality		0.002 (0.002)		-0.000 (0.001)
Number of Followers		-0.001 (0.001)		-0.001 (0.001)
Days on Twitter		-0.001 (0.002)		-0.001 (0.001)
Days on Twitter Missing		0.014 (0.014)		0.010 (0.015)
Constant	0.000 (0.000)	0.006** (0.003)	0.000 (0.000)	-0.000 (0.002)
N	6,687	6,687	6,687	6,687
R ²	0.021	0.024	0.014	0.014

*p < .1; **p < .05; ***p < .01

Robust standard errors in parentheses.

Eigenvector centrality, number of followers, and days on Twitter in standard units and centered at zero.

Next, we turn to the tweet outcome. The last three columns of Table 2.4 show that both DM conditions produced a positive causal effect on tweet activity. The “organizer” message caused fewer tweets than the “follower” message: 1.1 percentage points fewer than the 2.7-percentage-point boost generated by the follower message among subjects assigned to the DM conditions ($p = 0.03$). This evidence is suggestive of a priming effect in which the “organizer” identity reduces the future likelihood of tweeting but not the more immediate task of signing an online petition. We return to this apparent finding below.

The second stage of the experiment was designed to identify the causal effect of the tweet button on subsequent tweet activity by the subject’s own followers. Table 2.5 shows that among those who completed the petition, the effect of being shown a tweet button was large: The treatment caused nearly half of exposed subjects to click and tweet a message about the petition to their followers. (The positive constant may seem counterintuitive, but upon investigation we discovered that it reflects users who independently tweeted a link to the petition without using the supplied functionality—either manually or, for example, using a built-in Twitter app in their web browsers. This also illustrates the advantage of an experimental design, which can distinguish between this baseline activity and tweets caused by the manipulation.)

The interaction in Model 2 reiterates the previous finding that the organizer message appeared to depress tweet activity. Under an additional assumption¹³ that all subjects who signed the petition in one DM condition would have signed the petition in the other DM condition, we can interpret the heterogeneous effect in causal terms: Being primed as an “organizer” *caused* the encouragement to be less effective to the tune of more than 38 percentage points. In other words, the “follower” prime was more than twice as effective at encouraging future tweets as the “or-

¹³See Appendix B1 for a full discussion of this assumption and its plausibility in this application.

ganizer” prime. This is a substantial difference although, again, its interpretation rests crucially on the assumption referenced above.

Table 2.5: Study 1: Effects of Tweet Encouragement

	Tweeted	
	(1)	(2)
Shown Tweet Encouragement	0.454*** (0.086)	0.624*** (0.104)
Treatment: Organizer		0.053 (0.104)
Encouragement X Organizer		-0.384** (0.170)
Constant (Treatment: No Encouragement)	0.186*** (0.051)	0.161** (0.067)
N	109	109
R ²	0.214	0.267

*p < .1; **p < .05; ***p < .01
Robust standard errors in parentheses.

Network Effects

Recall that the subjects of the second-stage experiment were followers of petition signers. In effect, randomly treating some petition signers with the treat encouragement randomly exposed their followers to additional tweets. Column 1 of Table 2.6 shows that this manipulation was quite effective—those who followed exposed subjects were 63 percentage points more likely to have seen a retweeted message. However, despite being potentially exposed to retweets, columns two through five show that treated subjects were not significantly more likely to sign or tweet the message themselves. This finding parallels the ineffectiveness of public tweet sent by our partner organization. At least in this experiment, only direct messages had significant effects on participation.

Table 2.6: Study 1: Effects of Tweet Encouragement on Subjects' Followers

	Shown Tweet		Signed		Tweeted	
	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure: Followed Subject Shown Tweet Encouragement	0.628*** (0.029)	0.005 (0.006)		-0.004 (0.010)		
Exposure: Followed Subject Tweeted			0.008 (0.010)			-0.006 (0.016)
Constant	0.150*** (0.025)	0.010** (0.005)	0.008 (0.006)	0.012 (0.010)	0.013 (0.012)	
N	1,176	1,176	1,176	1,176	1,176	1,176
R ²	0.374	0.0004	-0.001	0.0003	0.001	

*p < .1; **p < .05; ***p < .01

Robust standard errors in parentheses.

All regressions weighted by inverse probability of exposure.

Results: Study 2

The results from Study 2 are broadly in line with those of Study 1.¹⁴ In particular, we replicated the finding that the “organizer” DM condition caused fewer subsequent tweets. The public tweet did not have a significant effect on petition signatures; in this study as well, not a single subject assigned to the public condition completed the petition. As before, additional exposure to direct messages caused a significant number of petition signatures and tweets.

Main Effects

The effect sizes we estimate from Study 2 are somewhat larger than those in Study 1, but they remain substantively comparable.¹⁵ The “follower” and “organizer” messages boosted petition signatures by 4.6 and 4.3 percentage points, respectively (see Table 2.7). The effects of the direct messages on signing were not significantly different from each other ($p = 0.60$). The direct messages also significantly increased tweet behavior. As in Study 1, priming the “follower” identity was more effective than the “organizer” identity, though the difference is no longer statistically significant.

¹⁴We did not find any evidence of balance problems as in Study 1.

¹⁵This could reflect the fact that the campaign was more timely and related to a current political dispute, in addition to the seasonality observed in other types of participation (Rosenstone and Hansen 1993).

Table 2.7: Study 2: Effects of Direct Message Treatments on Participation and Tweeting

	Signed		Tweeted	
	(1)	(2)	(3)	(4)
Treatment: Follower	0.046*** (0.004)	0.046*** (0.004)	0.014*** (0.002)	0.014*** (0.002)
Treatment: Organizer	0.043*** (0.004)	0.043*** (0.004)	0.011*** (0.002)	0.011*** (0.002)
Account Type: Male		0.005 (0.004)		0.005** (0.002)
Account Type: Organization		-0.019*** (0.004)		-0.003 (0.002)
Account Type: Unknown		-0.003 (0.008)		0.002 (0.004)
Eigenvector Centrality		-0.002 (0.001)		-0.001 (0.001)
Number of Followers		0.001 (0.002)		0.003** (0.001)
Days on Twitter		-0.002 (0.002)		-0.002* (0.001)
Days on Twitter Missing		-0.019 (0.014)		0.003 (0.013)
Constant	0.000 (0.000)	0.003 (0.003)	0.000 (0.000)	-0.001 (0.001)
N	8,507	8,507	8,507	8,507
R ²	0.019	0.022	0.005	0.008

*p < .1; **p < .05; ***p < .01

Robust standard errors in parentheses.

Eigenvector centrality, number of followers, and days on Twitter in standard units and centered at zero.

The tweet link caused a 35.0-percentage-point increase in tweeting behavior in the restricted model shown in Table 2.8, but, as in Study 1, there were differential effects by DM condition: The button increased tweets by 47.2 percentage points among subjects sent the “follower” message and by 23.1 percentage points among subjects sent the “organizer” message. Again, we can interpret this difference causally under the assumption that all petition signers in one DM condition would have signed in the other condition.

Table 2.8: Study 2: Effects of Tweet Button on Subsequent Tweets

	Tweeted	
	(1)	(2)
Shown Tweet Encouragement	0.350*** (0.055)	0.472*** (0.077)
Treatment: Organizer		0.037 (0.059)
Encouragement X Organizer		-0.241** (0.110)
Constant (Treatment: No Encouragement)	0.100*** (0.029)	0.083** (0.036)
N	221	221
R ²	0.154	0.181

*p < .1; **p < .05; ***p < .01
Robust standard errors in parentheses.

Network Effects

In contrast to the null network findings in Study 1, Table 2.9 presents evidence that signing the petition was strongly influenced by others’ tweets. Column 1 shows that the manipulation was effective. Column 2 shows an ITT effect of the tweet encouragement of 2.0 percentage points. Column 3 shows the estimated effect among compliers: Subjects who followed others who tweet if and only if they are shown the tweet encouragement were 5.6 percentage points more likely to sign the petition. Considering the relatively low rates of participation generally, these effects

are substantively large. Columns 4 and 5 repeat the analyses for the “tweeted” dependent variable: We observe no significant differences by exposure condition at the $p < 0.05$ level. Still, the fact that network effects on petition signatures are larger in magnitude than the main effect of direct contact from LCV itself is broadly consistent with the notion that “networks of recruitment” are important for promoting participation (Verba, Schlozman and Brady 1995).

Table 2.9: Study 2: Effects of Tweet Encouragement on Subjects' Followers

	Shown Tweet		Signed		Tweeted	
	OLS	(1)	OLS	IV	OLS	IV
Exposure: Followed Subject Shown Tweet Encouragement	0.352*** (0.031)		0.020*** (0.006)		0.004 (0.003)	
Exposure: Followed Subject Tweeted			0.056*** (0.018)			0.013 (0.010)
Constant	0.137*** (0.027)		0.010*** (0.004)	0.002 (0.006)	0.004** (0.002)	0.003 (0.003)
N	1,975		1,975	1,975	1,975	1,975
R ²	0.142		0.005	-0.051	0.001	-0.008

*p < .1, **p < .05, ***p < .01

Robust standard errors in parentheses.

All regressions weighted by inverse probability of exposure.

Treatment Effect Heterogeneity by Account Type

Theories of how online mobilization activities affect political behavior are focused on the individual citizen: Appeals from peers or groups via social networks may induce citizens to make contributions to public goods. The treatments deployed in our two experiments were developed with individual Twitter users in mind; however, the direct messages were sent to the accounts of organizations as well. All else being equal, we would expect individual users to be more likely to both sign the petition and retweet the petition link.

Twitter does not provide account-type information, so we hand-coded the profiles of all experimental subjects. We coded each account as female, male, an organization, or unknown. We relied on users' profile pictures and descriptions to determine account type. Organizations were easy to identify: They typically use language such as "We are a non-profit dedicated to..." in their description fields. Determining gender could be difficult when the profile pictures were not clearly male or female. When possible, we used cues in the description field such as "Activist, educator, and father of two..." When we could not determine gender or organizational status, we coded a profile as being of unknown type. Two of the authors carried out the coding; on a sample of 200 profiles, our inter-coder reliability was extremely high (Cohen's $\kappa = 0.90$).

Figure 2.3 presents the results of our heterogeneous effects analyses. The conditional average treatment effects (CATEs) and 95% confidence intervals are shown for all four account types, broken out by dependent variable and study. In Study 1, we observe some treatment effect heterogeneity on the "signed" dependent variable: Treatment effects are much smaller for organizations compared to individuals. We observe no such heterogeneity for the "tweeted" dependent variable. The second row presents the estimates for Study 2. We see nearly the identical pattern: On the "signed" dependent variable, organizations have much smaller treatment

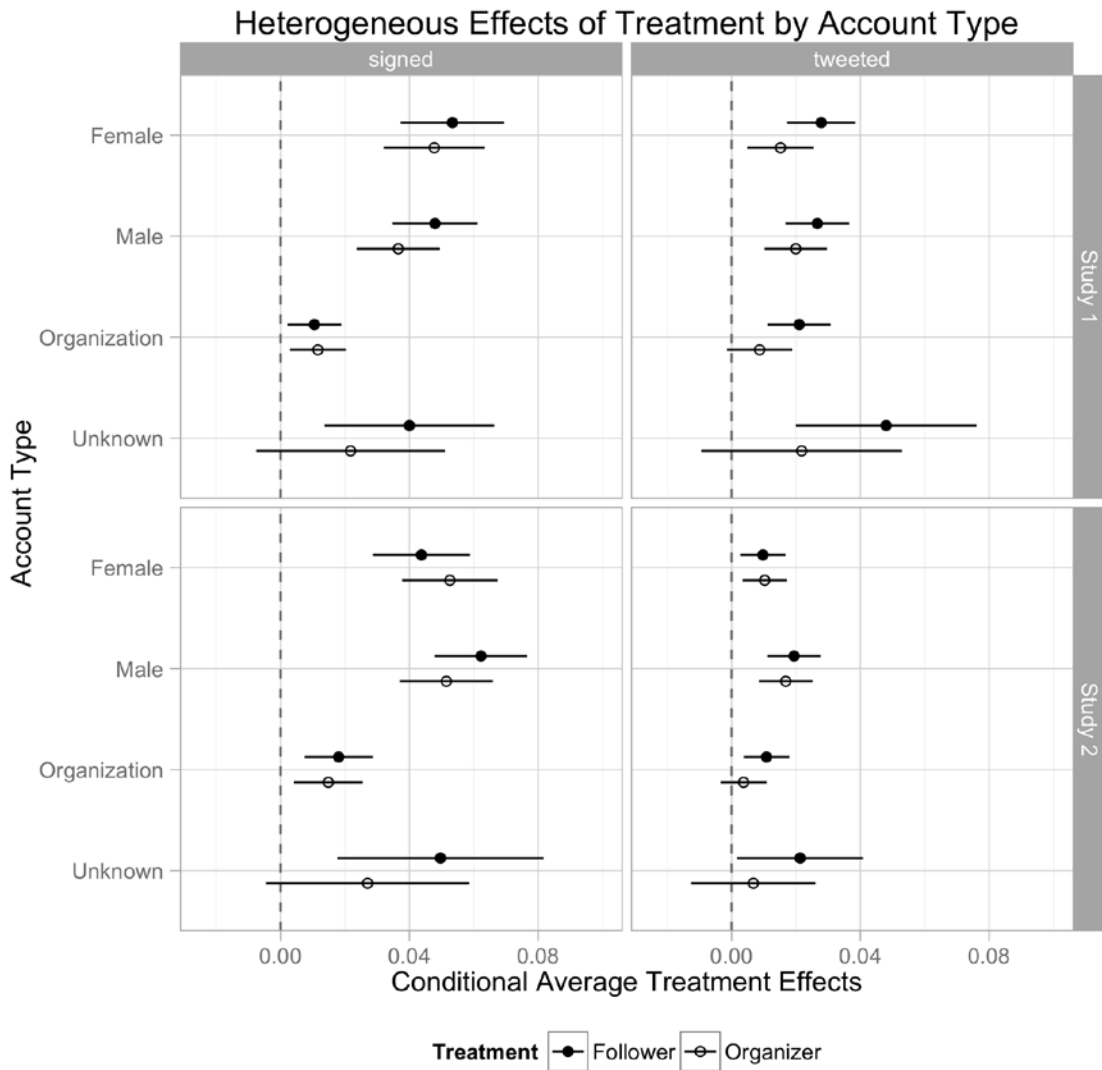


Figure 2.3: Entries are conditional differences-in-means with 95% confidence intervals. In Study 1, the sample was 38.4% male, 30.4% female, 24.4% organizations, and 6.8% unknown; in Study 2, the sample was 39.3% male, 32.9% female, 22.1% organizations, and 5.7% unknown.

effects than individuals, but this difference is not apparent for the “tweeted” dependent variable. Interestingly, there is no consistent pattern for the relative size of treatment effects among men and women; the treatments appear to work equally well for both, regardless of dependent variable.

We present these estimates in a regression format in Appendix B3, along with

heterogeneous effects analyses by subjects' number of followers, number of days their Twitter account was active, and eigenvector centrality.¹⁶ These analyses do not uncover a systematic pattern of treatment effect heterogeneity, though it does appear that treatment effects are marginally smaller for more central users. We interpret this finding cautiously, since organizational accounts tend to have higher centrality scores.

Discussion

This study identifies several robust findings about the effectiveness of different types of mobilization appeals on Twitter. First, direct messages are far superior to public tweets in generating supportive behavior in the form of online petition signatures, tweets, or even retweets. In both of our experiments, not a single subject assigned to be exposed only to the public tweet signed or retweeted the petition. We find that DMs produce approximately a 4-percentage-point increase in clicks. If an organization were to send out 250 direct messages (the maximum) per day for 30 days, they could expect to collect $250 \cdot 30 \cdot 0.04 = 300$ signatures over the course of a month. While this is a modest number, it should be weighed alongside the cost per DM, which is effectively zero.

Our results speak to the literature on social media and collective action. We find some support for the “reconceptualized” collective action theory of Bimber, Flanagin and Stohl (2005) when comparing the magnitudes of our effects. DMs of both types cause an increase in petition signatures greater than the effect on overall tweets to the petition link. However, compared to the effect of randomly assigning subjects who had already completed the petition to see the tweet button (35 to 45 percentage points), we see that the act of sending out a public tweet is arguably much easier to induce than a petition requiring some time or effort to complete.

¹⁶Appendix B2 presents randomization checks using these covariates.

While this is also consistent with some notions of “slacktivism,” we find that overall hypothesis difficult to square with the null effect of the public tweet on retweets in both studies. Overall, the results seem most consistent with the traditional perspective elaborated in the Civic Voluntarism Model, in which network effects are most effective.

Designing experiments on a social network like Twitter is difficult for a number of reasons. Most apparent is the fact that public tweets are potentially visible to anyone, which makes identifying their effects impossible without imposing additional assumptions about over-time persistence and anticipation. Another issue inherent to Twitter is a lack of individual-level exposure measures: Even if there existed a reliable indicator for whether a tweet was potentially visible to a given user (perhaps because a mobile or desktop app was active at the time), it would still greatly overstate whether it was actually seen and retained. This means that practically speaking, the only available estimand will be intent-to-treat (ITT).

One exception is cases in which the treatment is an encouragement to tweet and compliance can easily be measured. Our second-stage experiments, in which subjects were followers of the tweeters, have precisely this design. Its advantage is in distinguishing between the effects of homophily—similar users may already follow each other, be interested in similar issues, and retweet each other’s updates—from the effects of contagion (McPherson, Smith-Lovin and Cook 2001; Fowler et al. 2011). In Study 2 but not Study 1, we find that inducing tweet behavior within the network causes a significant and substantively large number of additional petition signatures. This is a potentially important finding for organizations seeking to launch “viral” campaigns: public tweets may be more effective when sent by followers of an organization than by the organization itself. Alternatively, users may require repeated exposure to public tweets in order for them to be effective. However, since this finding did not replicate across both studies, it should be interpreted

with caution.

Regardless, messages to followers are effective when they take the form of private DMs. While this may seem counterintuitive given Twitter's public network structure, one possible explanation is that individualized contact gives these messages the same essential properties as email (most users in fact receive an email notification when sent a DM). Email can be ineffective for certain purposes, such as mobilizing voter turnout (Nickerson 2007). But if a message is perceived as *solicited* contact from a trusted source, we hypothesize that it can be effective. This is consistent with existing research on emailed recruitment messages for web-based surveys, which emphasizes the torrent of unsolicited email and spam that users face daily. As one study points out (Porter and Whitcomb 2003), despite the relative ease with which spammers can mimic other senders, "it is still difficult to change the credibility of the message itself" (p. 587). In this case, the credibility lies in the fact that the recipient has already chosen to follow updates from the originator of the message (albeit over another medium).

Of course, we cannot rule out the possibility that these effects depend on multiple exposures to the same message (via an initial public tweet and subsequent DM). Future designs might vary exposure to public tweets over time in order to better address this question. It is also possible that public tweets work differently than posts on other social networks, such as Facebook. While we cannot test this possibility here, it seems plausible that the sheer number of tweets posted in real time diminish the effectiveness of any single post, while Facebook's algorithms keep the amount of social content to a manageable level, thus boosting the impact of any individual item. Experimental research has found strong effects of get-out-the-vote posts on Facebook, for example (Teresi and Michelson 2014).

Our final result is the differential effect of the "follower" and "organizer" identity primes. Existing research in social psychology has established that priming in-

dividual traits can have subsequent, unconscious effects on behavior. This can also extend to the ways in which people perceive themselves: When priming specific values (e.g., caring about the environment), subjects will tend to adjust their choices and behavior accordingly—but only if those values are central to their self-concept (Verplanken and Holland 2002). The implications for this study are straightforward. If members of an engaged network of environmental activists view themselves as organizers, priming this identity could bring forward other relevant considerations, such as the commitment it entails.

The differential effects of the “organizer” versus “follower” messages on the probability of tweeting may shed light on two theoretical questions. The first is, How does the authenticity of a message change its effectiveness? Twitter users may find the “organizer” label disingenuous, because in reality, they just subscribe to the advocacy group’s Twitter feed. The second question is, Are messages that prime the costs of collective action less effective? Perversely, encouraging subjects to help overcome free-rider problems with costly actions may primarily serve to reiterate that grassroots organizing is indeed personally taxing.

This is a surprising possibility given the traditional expectation that the Internet has the ability to promote collective action by reducing transaction costs overall (Farrell 2012). Twitter clearly possesses the properties—such as speed, reach, and versatility—necessary for this to be the case. Despite these low structural costs, organizations nevertheless compete for individuals’ limited attention online. Even small changes in perceived costs can reduce the probability of collective action.

Online Media Choice and Moderation

Despite vast changes in the way Americans seek out and receive information about the world around them, political scientists still tend to conceptualize media effects within a largely 20th-century framework. Campaigns target smaller and smaller niches; new media outlets tailor their content to narrower audiences; videos of off-the-cuff remarks spread virally within social networks. But unless the outcome of such processes can be measured in aggregate vote returns or large-scale tracking surveys, they seemingly remain illusory.

Making persuasive inferences about media effects requires surmounting an unusually challenging set of obstacles. Those outlined by Bartels (1993) as particularly difficult in a classic overview remain so today: with aggregate time-series data, media coverage can easily be confounded by the events precipitating it; with individual-level cross-sectional data, self-reported exposure may be correlated with political interest or other unobserved predispositions; and laboratory experiments, while ruling out alternative causal factors, still remain open to criticisms about external validity.

Adding to these challenges, measurement of exposure to media content in the real world tends to be based on unreliable self-reports, irregularly validated proxies, or aggregate traffic data that masks individual-level variation. When innova-

tive solutions to the problem present themselves, they can be leveraged to produce useful observational evidence on media exposure. But to generate convincing evidence on effects, such a measurement strategy needs to be coupled with an experimental approach.

Changes in technology and Americans' media habits both add to the difficulties and point the way toward possible solutions. On the one hand, a handful of network and cable TV channels is now thousands of potential sources of news and information about politics, driven by online-only (and often partisan) new media outlets and feeds on social networks (Mutz and Young 2011). On the other, the proliferation of data on users' online habits means that it is now more possible than ever to track individual behavior (King 2011; Bond et al. 2012).

This study explores how media effects play out online. More and more Americans get their news and information about politics online: According to Pew, in 2013 82% of Americans said they got news on a computer, and over half (54%) did so on a mobile device (Pew 2014). While online media are not yet the primary source of political information for most Americans, they are for the youngest age groups, and the upward trend continues.

In some ways, these changes may make identifying media effects harder than ever. For example, the "forced-choice" paradigm assumed by earlier studies of media effects is not well-suited to an information environment that people can, in part, construct and customize for themselves (Arceneaux and Johnson 2010). Selective exposure affects not only whether individuals see and hear particular content, but how they consciously (and unconsciously) organize their entire media diets (Bennett and Iyengar 2008, p. 724). Theoretically, this means that substantively meaningful self-selection could have occurred well before exposure to any particular content is observed; passive consumption is perfectly compatible with active curation. This insight merely pushes concerns about self-selection of congenial in-

formation back one level: to the strategies and defaults arranged in advance rather than the particular content heard or viewed on a particular occasion.

The most convincing evidence of media effects will come from research conducted in naturalistic settings. This study uses a field-experimental design to identify the causal effect of individuals' informational search strategies online. I combine this approach with direct measures of subjects' Internet behavior to test whether partisan selective exposure mediates the effects of political information online. In doing so, I provide the first test of whether both a "balanced" news diet and selective exposure can coexist within the same framework.

The paper proceeds as follows. First, I provide a brief overview of the most recent evidence on media effects and selective exposure. Second, I outline a new theoretical contribution to the study of media effects in the context of open-ended information search. I then introduce and provide unique observational evidence on online media exposure. Next, I describe the design and measurement strategies of two separate field experiments and report the results of each. A discussion then concludes.

A Motivating Puzzle

Since the original Columbia studies, research within the "minimal effects" paradigm has often relied on the mechanism of selective exposure to explain null findings of media effects (Lazarsfeld, Berelson and Gaudet 1944; Klapper 1960; Bennett and Iyengar 2008; Arceneaux and Johnson 2010). Yet, as a growing number of scholars have documented, the evidence for selective exposure is not as strong as initially claimed. This leads to two related puzzles: First, why does the selective exposure hypothesis find support in laboratory experiments but not observational studies using real-world behavioral measures? And second, if selective exposure is not as common as once believed, why are media effects outside the lab still modest and

fleeting, if they can be identified at all?

These questions cut to the core of longstanding questions about the extent to which democratic citizens are exposed to competing viewpoints, thought to be a prerequisite for informed collective decision-making (Mutz 2006; Jamieson and Cappella 2008; Shapiro 2013). In the worst-case scenario, most forcefully articulated in the context of the 21st-century fragmented media environment by Sunstein (2007), people elect to consume only ideologically congenial information (the “Daily Me”), resulting in an echo-chamber effect and, ultimately, increasing polarization (see also Negroponte 1995). Others have additionally worried that hidden algorithms could speed along this process by replicating and reinforcing people’s preferences, especially on social media (Pariser 2011).

Since the critical review of Sears and Freedman (1967), however, skeptics have questioned whether people are as selective in their choices of how to receive information as previously supposed. More recently, advances in measurement and data collection have backed up these claims. For instance, Gentzkow and Shapiro (2011) use aggregate web traffic data to conclude that ideological segregation online is less severe than for national newspapers or in face-to-face social networks. On Twitter, Barberá (2014) uses panel data to show evidence of cross-cutting follow patterns that lead to moderation rather than polarization.

This body of work sits alongside research uncovering at least suggestive evidence of selective exposure (Garrett 2009, 2013). First, there is experimental evidence from the laboratory: The main finding of the studies conducted by Arceaux and Johnson (2010) is that, consistent with Prior (2007), allowing subjects the choice to “tune out” of political programming moderates the polarizing effect of political news content. But among subjects who select *in* to political news, there is an apparent sorting effect whereby subjects spend more time with pro-attitudinal than counter-attitudinal content. Other research finds a polarizing effect driven by

people who are already fairly extreme in their views (Levendusky 2013). Second, using panel data of over-time dynamics during a presidential campaign, there is evidence of partisan selective exposure (Stroud 2008).

How can these findings be reconciled? It is likely that research using self-reported measures exaggerate the evidence for selective exposure, given that people may be more likely to remember using sources that align with their political identities. Another possibility is that evidence of effects in the laboratory do not generalize to the real world and that given day-to-day distractions and competing demands for time, people's propensity to select sources of congenial information is lower than expected. Finally, it may be the case that under certain conditions, people can be induced to behave in ways consistent with the selective exposure hypothesis, but under others they fall back on more balanced media consumption habits.

I hypothesize that these conditions are determined by the requirements of a given task. Seeking out particular information—for example, in the context of an upcoming election or an issue that could affect an individual personally—means relying on a set of open-ended search strategies that may differ markedly from day-to-day forms of media consumption. Central to the distinction between these two modes is the idea of *passive* versus *active* reception of information. The former occurs within a context of ingrained habits and defaults, while the latter may depend on cues and heuristics that help people navigate less familiar territory (e.g., Popkin 1991).

More concretely, the picture given by the best observational evidence on media choice is that of a vast middle: Most people turn to large, relatively centrist sources of news for political information. This pattern could be driven by a number of mechanisms. Perhaps it is explained by a pattern of preference for ideologically centrist content (Fiorina, Abrams and Pope 2011). Perhaps it reflects a general lack of interest in political information or is even a byproduct of people's preferences

for *non-political* content. Or perhaps it is a result of something else: the defaults and browsing habits of the online public.

I focus on this last possibility because it has the potential to explain how selective exposure as a mechanism could operate even against the backdrop of a relatively balanced aggregate news diet. According to this view, media consumption habits are the result of an interaction between preferences for content (whether for entertainment versus news, or for a certain ideological slant) and the information environment—the number of available sources, the cost of switching, etc. But while contemporary accounts of Internet media tend to assume costs that approach zero, they fail to take into account the hidden obstacles and defaults that structure people’s habits online.

To take the simplest example, modern web browsers come pre-loaded with bookmarks for large news and entertainment sites (such as AOL and Yahoo). Many people still use portals for email and other services which link to headlines, weather and other information. Sometimes, such sites automatically load on startup. It is not hard to customize one’s settings, but doing so already requires preferences over sources—a perceived cost that may be too high for many people. For individuals with passing or intermittent interest in politics, for example, such built-in choices may be sufficient for most day-to-day needs.

Conversely, people may need to rely on different strategies when actively seeking out information (e.g., Iyengar et al. 2008). Rather than passively relying on semi-hidden defaults, it becomes necessary to use search engines, query one’s friends or social networks (crowdsourcing), or consciously think about particular sources of information that would be useful. If the passive mode encourages a tendency toward centrism and homogeneity in media choices, active search creates opportunities for personal predispositions to creep in at every stage. Search engines and social networks can customize results based on past preferences and actions, rein-

forcing past biases; shortcuts based on partisan or ideological affinity could drive decisions about media sources.

One interesting implication of this distinction between active and passive search processes is that the former could eventually become incorporated into one's media diet—novel strategies transforming into defaults. While speculative, this would predict a long-term over-time trend toward more ideologically segregated media diets, even from a relatively centrist starting point. This would explain the findings of Stroud (2008), who found evidence of partisan selective exposure in panel survey data (including respondents who reported online media use).

If this picture is accurate, then the observed lack of evidence for selective exposure in observational data could be a result of choice architecture as much as people's preferences for diversity in media sources (e.g., Sunstein and Thaler 2008). Likewise, when the structure of the information environment favors ideological segregation—as it arguably does on Twitter, where it is easy to find co-partisans and retweet only their content to like-minded followers—aggregate data follows the same pattern (Conover et al. 2011). One question this paper poses is whether the selective exposure mechanism can be separated from the media context, and whether the persuasive effects of media content vary as a result.

This paper is structured in two parts. First, I take advantage of a large and unique survey panel whose members record real-time tracking data on their online browsing habits. Merged with individual-level survey responses, this data offers an unprecedented look at the interaction between political predispositions and media diet. In this first section, I provide an overview of the data and offer an observational look at the political browsing habits of the survey's respondents. Second, I supplement this analysis with two online field experiments combining similar direct measures of respondents' online media exposure with individual-level covariates.

Balance or Bias in Online Media Exposure?

The first step of this investigation is to make use of a unique data set: individual-level survey data merged with continuous tracking of panelists' Internet browsing behavior. This data, which is completely anonymized, was collected by the online polling firm YouGov as part of an ongoing attempt to gauge whether survey respondents are willing to install tracking software (called Wakoopa). This data provides direct evidence of respondents' media habits for political (and non-political) information. It is unique in that it combines data on site visits with individual-level survey responses. For this particular panel, there are no limits to the types of websites that can be included in the data. Moreover, the software tracks web traffic (minus passwords and financial transactions) for all browsers installed on a user's computer and cannot be blocked.

The online tracking panel is currently branded as YouGov Pulse (see Figure 3.1). Panelists are recruited from YouGov's traditional participant pool via incentives. At least initially, these incentives have been very strong: 4,000 "points" for signing up and downloading the Wakoopa software—roughly 8 times the number offered for a typical survey—and 1,000 additional points every month. Participants in online surveys can redeem these points for clothing, prepaid gift cards, and other merchandise.

The data set contains more than 6.3 million observations at the respondent-site level, covering panelists who installed the tracking software on their desktop computers (excluding mobile phones). This sample includes site visits from 1,392 individuals over a three-week period in 2015, from February 27 to March 19. Since respondents were not recruited using random sampling, YouGov typically employs sample-matching weights to make its results representative of the general population (Rivers 2006). For now, I refrain from making overgeneralizations but note that the results below closely match similar analyses from the Mechanical Turk sample

We would like to invite you to take part in **YouGov Pulse** - an exciting new project to find out more about how people use the internet.

Pulse tracks your internet usage and anonymizes it to give a picture of how real people use the internet. We look at search terms, what ads you see (and what ads you close) and the websites you visit.

Pulse does not:

- Collect any usernames or passwords
- Track any online transactions
- Have any impact on the speed of your internet connection or data limits

Pulse does:

- **For a limited time offer** if you download Pulse you will receive **4000 points** within 5 weeks. (Usually 2000 points).
- For every month you have Pulse running on your device you'll receive an additional **1000 points**
- To earn extra points, download to your mobile and tablet too (**it's 1000 points per device, per month!**)

*Note: the Pulse app can be active only on one desktop at a time. Monthly points will be paid at the beginning of the following month if the software is still running at that time.

Signing up for YouGov Pulse is easy: just click the link to download and install the software. Please only do this when you are on the device you want to install it on to.

We will never share your personal information with our clients:



Figure 3.1: Screen shot of an email sent to YouGov panelists on April 8, 2015, inviting participation in YouGov Pulse.

in the second study.

Table 3.1 summarizes the demographic characteristics of the sample. Perhaps most notably, it skews younger and more educated than the general population, although the gender, racial and party breakdown are fairly representative and capture a more diverse cross-section than the MTurk samples discussed below.

Most directly, this data allows me to test whether respondents silo themselves into informational cocoons according to partisanship or ideology. In order to measure the general ideological orientation of individual political websites, I employ data from the Internet analytics firm comScore, which maintains a 12,000-person survey panel of the general Internet audience called Plan Metrix. Employing both direct responses and imputation, comScore provides estimates of the overall de-

Table 3.1: YouGov Pulse Sample: Demographics

Category	Proportion	Category	Proportion
18-30	0.233	Male	0.440
31-40	0.253	Female	0.560
41-50	0.156	Dem	0.367
51-60	0.230	Rep	0.224
61-70	0.113	Indep	0.305
Over 70	0.015	Other	0.038
No HS	0.033	White	0.679
High school	0.218	Black	0.103
Some college	0.377	Hispanic	0.069
College	0.250	Asian	0.059
Postgrad	0.123		

mographic composition of individual sites' audiences. Using these estimates from March 2015, I create two separate audience-based measures of website slant. The first employs the ideological self-placement of site visitors in the Plan Metrix panel. I take the share of respondents who classify themselves as "very conservative" or "somewhat conservative" as a fraction of those who place themselves anywhere on the 5-point ideological scale (also including "middle of the road," "liberal," and "very liberal"). This creates an index of conservative readership that can take values from 0 to 1, although practically speaking the scores do not go above 0.85. The second measure of slant uses the partisanship of Plan Metrix panelists: In a similar way, I compute the Republican share of those who identify with either party.

The measures correlate with each other fairly well ($r = 0.50$). Since this is not perfect, I present results using both measures for completeness. While the estimates have high face validity, one peculiarity is evident from the figures below: most popular sites tend to cluster somewhat left of center (0.5). This is an artifact of the measures' construction: Since no site achieves 100% conservative readership, the distribution is pushed to the left. Relative ideological placements are not affected, however. Table 3.2 displays a sample of the most-visited websites in the YouGov Pulse panel, along with the number of visitors logged over the three-week

period and both measures of slant. Immediately apparent is that MSN News—a mainstream, centrist news and information portal—is by far the most popular source, by almost an order of magnitude. This somewhat matches the available comScore figures, which show that among political and news sites logged in the YouGov data, MSN is the second most visited overall (to CNN.com). Looking at the partisan and ideological slant measures, it is clear that conservative sites such as the Drudge Report and Townhall have higher scores than left-leaning sources such as *Slate*. (One possible anomaly is Daily Kos, a left-liberal site whose ideological score appears more accurate than the partisan one.)

Table 3.2: The partisan and ideological slant scores of some of the most-visited sites in the sample, arranged in reverse order of popularity. All scores are listed in Appendix C2.

Site	Visits	Part	Ideo	Site	Visits	Part	Ideo
msn news	36263	0.447	0.339	bbc	1976	0.445	0.326
yahoo news	5000	0.451	0.341	theblaze.com	1640	0.496	0.348
foxnews.com	4777	0.521	0.372	cnn.com	1565	0.436	0.319
townhall.com	4372	0.571	0.450	breitbart.com	1317	0.507	0.404
buzzfeed.com	4077	0.412	0.285	nbcnews.com	1103	0.461	0.333
huffingtonpost.com	3898	0.453	0.337	wsj.com	994	0.46	0.362
nytimes.com	2532	0.449	0.332	telegraph.co.uk	952	0.432	0.324
news.google.com	2253	0.424	0.331	freep.com	914	0.503	0.350
drudgereport.com	2201	0.624	0.454	slate.com	852	0.414	0.285
daily kos	2148	0.483	0.331	nypost.com	764	0.456	0.362
washingtonpost.com	2025	0.471	0.351				

Before combining these two sources of data together, I categorized the YouGov Pulse panel’s site visits so that I could separate those relating to news and politics from the rest of the web traffic.¹ I also removed visits to local news websites to focus on national sources. Confirming similar findings elsewhere (Flaxman, Goel and Rao 2013), the resulting share devoted to news and information about national politics is strikingly low: 1.6 percent of all visits, or just over 102,000 out of the

¹I used a combination of Wakoopa’s proprietary categorization scheme and keyword searches to categorize roughly 73% of website visits in the sample. The remaining sites comprise a “long tail” with very few visits each.

6,319,441 observations in the sample. I then separately aggregate the number of visits per site for all respondents as well as those who identify as Democrats and those who identify as Republicans. Overall, I could match 208 politics and news websites in the sample to a measure of slant. For Democrats only, this number is somewhat lower, at 175, and for Republicans there were 133 (Table 3.1 shows that there were fewer Republicans in the sample, possibly driving this disparity).

Does the Internet facilitate the Daily Me? Figure 3.2, which plots the density of site visits against the measure of ideological slant, shows that this is generally not the case. Most site visits in the YouGov Pulse sample cluster around a handful of relatively centrist sources such as CNN, MSN, and Yahoo! News. The density curves for Democrats and Republicans are similar to each other and also to the curve for the sample as a whole. The distributions are not bimodal, as extreme ideological segregation would predict. However, there are two smaller bumps at the extremes, corresponding on the right to popular conservative sites Townhall and the Drudge Report. The bump on the left is actually BuzzFeed which, although its overall audience might lean liberal, is generally considered a mainstream news and entertainment site. The corresponding graph using the partisan slant measure rather than the ideological measure is given in Appendix C3. The general pattern is the same, although the scores are somewhat more dispersed.

The evidence, then, is consistent with a view that—among the small fraction of respondents who actively visit news and politics websites—the preponderance of the content encountered is relatively centrist and balanced ideologically. At least in this sample, there also appears to be suggestive evidence of a smaller, intense subgroup of Republicans who (possibly in addition to mainstream sources) consume conservative, but not liberal, news and information about politics. A similar bump on the left, corresponding to the popular viral site BuzzFeed, seems to be an artifact of the ideological measure and not an indication of symmetric echo chambers in

The Online Political Media Diet: YouGov Pulse Data




Figure 3.2: Density plot of aggregate site visits from the YouGov Pulse sample. Site ideological slant on the x-axis is measured using comScore data on audience composition. N = 102,128 visits.

online media diets.

MTurk Experiment Design

It is possible that this observational picture masks substantial variation. In order to gain causal leverage on Internet search behavior in a real-world environment, this study uses an online field experimental design. I take advantage of several innova-

tions to improve measurement and inference. First, I capture outcomes on actual browsing behavior via a small piece of software installed beforehand (with permission) on subjects' computers—the same approach as Study 2 in Chapter 1. Second, I randomly assign panelists to receive an email treatment designed to induce a purposeful, open-ended search for information about a particular, low-salience political issue: the regulation of for-profit colleges and universities. This enables a comparison between two sets of potential outcomes: those of subjects relying on default strategies for processing political information, and those induced to expend additional effort to seek out novel information.

Measurement Strategy

As in Chapter 1, I use a browser plug-in to capture a subset of that history (for a prespecified list of sites and over a fixed length of time) and record it alongside matching survey responses. I adapt this method for the current study's measurement strategy as follows. I created a small piece of software for Google Chrome browsers that saves trace data from a user's web history: whether or not any site from a predetermined list has been visited in the past five days. When manually activated, the browser plug-in takes this snapshot and transmits it to a Qualtrics database, linked with the (anonymized) responses from the same user's survey input. In order to set up a panel of subjects with this software installed on their primary computers, I posted a survey on MTurk collecting pretreatment covariates and requesting that respondents optionally install the software for a 50-cent bonus (on top of the small payment offered for completing the short survey regardless).²

²MTurk respondents could click for a list of the URLs scanned by the browser plug-in. I also provided the source code for the software, written in JavaScript using Google's Chrome API. These steps were intended to reassure potential subjects that the software does not collect identifying information or perform open-ended searches of users' browser history. At the time the study was conducted, Amazon's MTurk terms of service prohibited requiring workers to install software as a condition for being paid.

One limitation of the direct measurement approach is that it assumes that the list of URLs the software can search for in subjects' histories is exhaustive. For generic applications, this seems like a defensible assumption. However, for studies of information search behavior, some users might seek out sources that are hard to predict in advance. To generate the list of sources that the software can capture, I combined two approaches. First, I collected data on average monthly unique visitors from comScore, collecting sites listed in the "General News" category with at least 300 per month, those listed under "Politics" with at least 50, and those in the larger "News / Information" category with at least 1,500. I removed sites that were primarily local, non-U.S., or non-news-related to maintain the focus of the list. Finally, I supplemented this list with the sites of the top 10 newspapers by total average circulation and entries from a number of partisan blog directories. The total number of URLs generated using this first approach is 156.

I augmented this list of general news and politics sources using a second approach more directly tailored to the particular design of this study. Thinking through the possible steps of how one might look up information about a novel political topic, I used as wide a net as possible for collecting URLs to check for: search engine results from various queries related to the issue ("for-profit education," "for-profit colleges," etc.); Wikipedia pages, social media, explainer sites, higher education news sources, partisan sites, and even political non-news sites (such as Senator Tom Harkin's page dedicated to the issue). This generated an additional 72 links, for a total of 228.³

Subject Recruitment

I recruited subjects via Mechanical Turk (MTurk), an online marketplace for requesting and completing self-contained, relatively straightforward tasks Berinsky,

³See Appendix C1 for a full list of sites included.

Huber and Lenz (2012). Of 1,500 initial respondents, N=467 both agreed to install the browser widget and successfully did so. From this experimental sample, I then randomized subjects into either treatment and control, as described in the next subsection. All subjects in the sample were in the U.S. and had at least a 95% approval rate on previous MTurk tasks (sample characteristics: 65.7% male; 81.8% white, 6.2% black, 6.9% Hispanic; 50.5% college-educated; 43.3% Democrat, 14.3% Republican; median age, 29).

Timeline

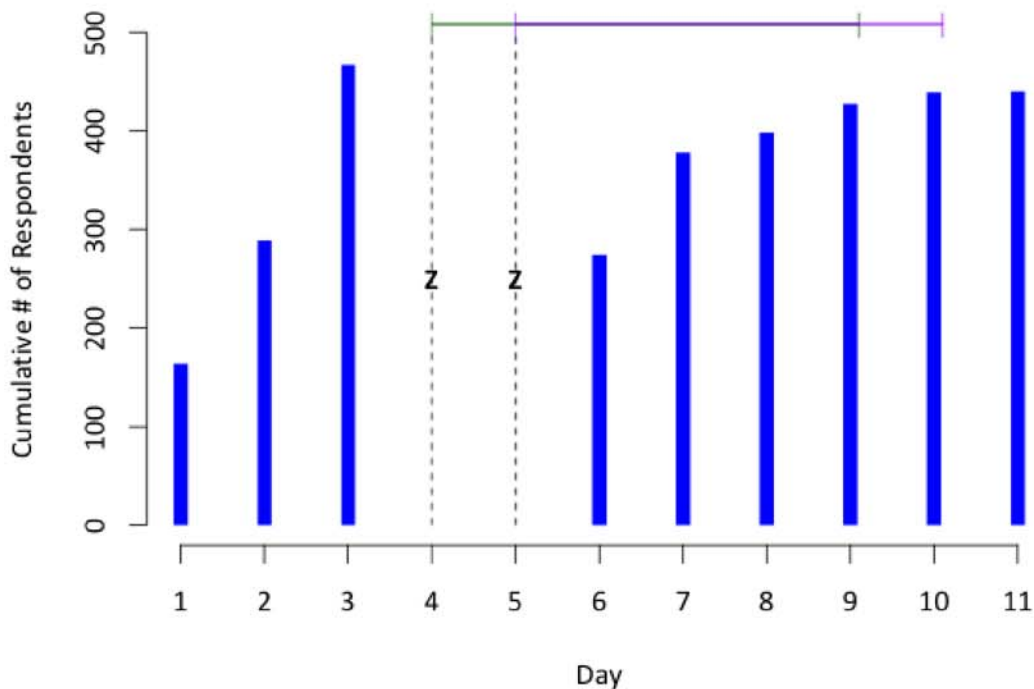


Figure 3.3: Study timeline. The green and purple lines indicate the “memory” of the browser plug-in software, which could only capture site visit data for the previous 5 days.

Figure 3.3 visualizes the timeline of the study by showing the cumulative num-

ber of respondents in each wave in blue. After the initial collection of pretreatment covariates and installation of tracking software, the 467 subjects were assigned to treatment and control via complete random assignment (days 4 and 5 of the study).

The pretreatment survey asked questions about attitudes on a number of different issues, in addition to demographic and political covariates. It was designed in such a way as to make it difficult for a respondent to connect the content of the survey with the subsequent treatment. The treatment itself was an emailed encouragement to seek out information about one of the political topics asked about in the previous survey (sent twice, 1-3 days after the initial survey). The MTurk system allows requesters to send messages to workers who have completed previous tasks (in this case, the pretreatment survey). Here, an otherwise irksome feature of this particular subject pool—the fact that many workers complete dozens of different tasks a day, including many social science surveys—becomes an advantage: Anecdotal evidence from communications with workers suggests that it is unlikely that an emailed request to complete a subsequent task would be tied back to a given earlier survey. This is crucial because while pretreatment covariates are useful for producing efficient estimates, they should be collected in a way that minimizes the likelihood of demand effects. Another advantage of this type of treatment is that it is ecologically valid: MTurk workers often receive electronic requests to complete additional tasks, and the present treatment is delivered in the context of subjects' real-world, day-to-day information environments.

This overall approach is similar in spirit to that of Albertson and Lawrence (2009), who analyzed an encouragement design in which telephone survey respondents were randomly assigned to be asked to watch educational television programs. Follow-up surveys then asked respondents whether they complied and collected post-treatment measures of knowledge, attitudes, and salience. I collected the first two outcomes, although I only measured knowledge post-treatment in or-

der to avoid anticipation or other demand effects. After giving subjects at least two days to comply with the encouragement, I then sent another emailed request (to those in both treatment and control) for responses to a follow-up survey (days 6-11). The recontact rate was high, at 93.4%.

In both the pretreatment and post-treatment surveys, I measured subjects' media exposure via the tracking software. As a result, I can construct measures of subjects' initial (baseline) recent media exposure and the additional sites visited after that point for those who did and did not receive the treatment.

Treatment

In designing the treatment, I selected an encouragement intended to activate the kind of search for novel political information that would uncover the influence of selective exposure, if it is an identifiable mechanism. In particular, I asked subjects to spend some time learning about a political topic (regulation of for-profit colleges and universities) that affects many Americans and could conceivably be mapped onto the partisan divide, but is not salient in current political debates.

To more specifically invoke the mechanisms I was interested in, I worded the encouragement in the following way: "We'd like to follow up soon with a short survey to gauge your responses to a few more questions. In the survey, we will ask about the controversy over for-profit education. To prepare for that questionnaire, we'd like you to familiarize yourself with the issue. Feel free to look up information online the way you usually do, using whatever methods or sources you are comfortable with to learn about for-profit colleges."

MTurk Experiment Results

Here, I analyze the web browser data to look for patterns of selective exposure in the treatment and control groups. First, I show below that the treatment was effec-

tive: It caused respondents in the treatment group to seek out information about an issue that, by and large, they were not highly knowledgeable about. I created an indicator for respondents who visited at least one site on the list of sources specifically about for-profit education captured by the browser plug-in (72 possible), and regressed it on treatment assignment. Given the possibility that some sources were not included in the list of websites searched by the software, this is a lower bound. Table 3.3 shows that being assigned to treatment caused a roughly 10-percentage-point boost in the share of subjects who searched for information about the topic online—a substantively large magnitude, but still a relatively low level of compliance.

Table 3.3

<i>Dependent variable:</i>	
Visited Pages About For-profit Colleges	
Assigned Treatment	0.10*** (0.04)
Constant	0.04 (0.03)
Observations	467
Adjusted R ²	0.02

Note: *p<0.1; **p<0.05; ***p<0.01
OLS, standard errors in parentheses.

Selective Exposure?

In looking for evidence of “selective exposure” in this section, I test whether subjects’ partisan predispositions predict patterns of media choice. Strictly speaking, I cannot make a causal claim: partisan and ideological commitments are not randomly assigned. However, we can observe whether subjects’ behavior in the treatment and control conditions seem to differ systematically in ways that are instruc-

tive. To simplify presentation, I display the results—heterogeneous effects of treatment on Democrats and Republicans’ partisan media diets—graphically below.

First, however, I need an approximate measure of media outlets’ partisan leanings. For the purposes of this analysis, as described above, I again use the survey data from comScore. To construct an individual-level measure of the partisan lean of subjects’ media diets, I simply compute the average of the comScore partisanship index (Republican share divided by the Democratic and Republican share) for each site visited at least once by that person.

I begin with a plot of the correlation between subjects’ party identification and the partisan lean of their (pre-treatment) media diets. Figure 3.4 shows that there is such a correlation but that it is weak—an increase from just under 0.45 to approximately 0.47 going from strong Democrats to strong Republicans. In general most people consume political media near the “center,” where the center in this case is, again, somewhat shifted to the left. Furthermore, one can see that the correlation may partially be driven by extreme outliers: a handful of individuals with homogeneous, very liberal or very conservative media diets that correspond to their stated political leanings.

What happens when we try to induce individuals to seek out new information? Do they rely upon defaults and habits? If so, do these defaults nudge people in the direction of centrist mass media or partisan sources? To shed light on these questions, I run some simple linear models investigating the heterogeneous effects of the treatment on subjects’ media diets. Table 3.4 shows the effects of treatment on post-treatment media diet slant. Model 1 shows that the treatment alone has a small but statistically significant effect in the leftward direction. I make no explicit prediction about this main effect, but it may represent the overall critical nature of commentary about for-profit education online, which would tend to support robust federal action and express skepticism of the industry’s motives. Model 2 adds the

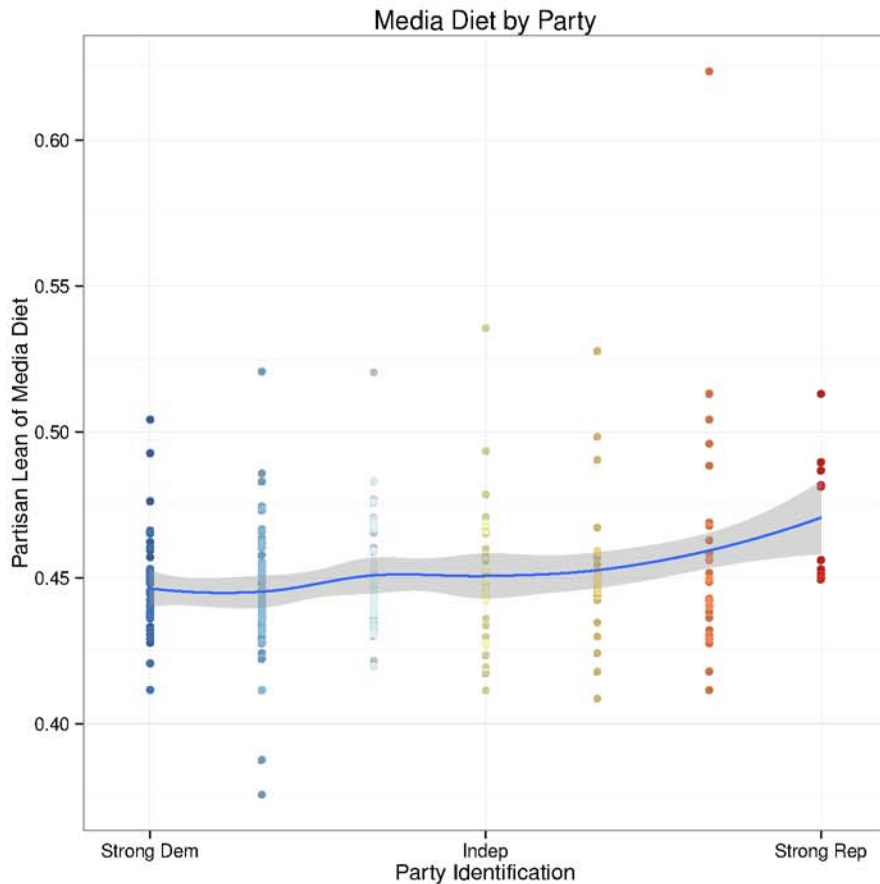


Figure 3.4: This plot shows the correlation between individuals’ political predispositions (as measured by the standard 7-point party identification scale) and the average partisanship of their media diet (before treatments were administered, in the first survey wave).

pre-treatment measure of media diet, which was captured in the first survey wave using the browser extension software. Not surprisingly, people’s over-time media habits are strongly correlated.

Since it is highly prognostic, I include the pre-treatment measure in the subsequent models to improve the precision of the estimates. In Models 3 and 4, we see a consistent pattern regardless of whether we investigate heterogeneity by party or ideological leaning: the negative (but small) statistically significant interaction coefficient implies that the treatment pushes subjects’ media diets in a more leftward direction the more *conservative* they are. This is the opposite of what we would

expect from selective exposure. Rather than seeking out reinforcing information, partisans in the sample sought out heterogeneous sources, resulting in moderation in their diets overall.

Table 3.4

	<i>Dependent variable:</i>			
	Media Diet Partisanship			
	(1)	(2)	(3)	(4)
Assigned Treatment	-0.01*** (0.004)	-0.01* (0.003)	0.01* (0.01)	0.01* (0.01)
Pre-treatment Media Diet		0.49*** (0.05)	0.49*** (0.06)	0.46*** (0.05)
Party ID			0.01*** (0.002)	
Ideology				0.01*** (0.002)
Treated x Party			-0.01*** (0.002)	
Treated x Ideology				-0.01*** (0.002)
Constant	0.46*** (0.003)	0.23*** (0.02)	0.21*** (0.03)	0.23*** (0.02)
Observations	286	229	217	229
Adjusted R ²	0.03	0.29	0.33	0.32

Note: *p<0.1; **p<0.05; ***p<0.01
 OLS, standard errors in parentheses.
 Party and ideology coded using 7-point scales.

Since interactions can be difficult to interpret, I further investigate these effects graphically. Figure 3.5 plots the distribution of individuals' average media diet partisanship for four groups: Democrats assigned to receive treatment, Republicans assigned to receive treatment, Democrats in the control group, and Republicans in the control group. Strikingly, as with the YouGov Pulse data, the same overall pattern of centrist media exposure can be seen here, with one significant exception. Republicans in the control group had a more conservative overall media diet on average, but after treatment was administered, the distribution of Republican

subjects' media diet partisanship moved sharply to the left. Thus it is the relative moderation of Republicans in the sample driving the overall effect heterogeneity. After that shift, the distributions of the four groups is essentially identical, again casting doubt on the selective exposure hypothesis, at least for this particular issue.




Figure 3.5: Density plot of individuals' media diet partisanship, as measured by the mean comScore partisan slant of all sites visited by each respondent post-treatment. Lines denote party and treatment subgroups. Leaners are coded as partisans.

More Comparisons

Finally, below I make further observational comparisons between groups in the MTurk sample. In Figure 3.6, I look at the distribution of overall visits to political news sources by different subgroups. Both plots show essentially the same pattern: more or less identical distributions centered around relatively moderate, popular websites. On the left, we see that treatment (gray lines) appears to boost the overall volume of site visits but does not measurably alter the ideological flavor of people's media choices. And in the right panel, there appear to be few differences in the site traffic of Democrats and Republicans within the treatment group. Again there is a slight bump on the right, corresponding to the Drudge Report, but the picture is not one of overwhelming ideological segregation. (The right panel is comparable to Figure 3.2, which shows similar site visit patterns in the Pulse data.)




Figure 3.6: Left, comparison of treatment (gray) and control (black) site exposure densities. Right, comparison of Democrats and Republicans within the treatment group (including leaners).

YouGov Experiment Design

To supplement the MTurk results, I designed an additional experiment to be run on a subset of the YouGov Pulse panel discussed earlier. There are several advantages to the additional study. First, I hope to alleviate concerns that results are being driven by the unpredictable or idiosyncratic nature of the MTurk respondent pool. Second, by taking advantage of the Wakoopa tracking software already installed on subjects' computers, I was not restricted by the need to compile a comprehensive list of potential sources to scan for; the software records all visits to facilitate coding the results for relevant web traffic on the back end. Third, relevant demographic and political covariates have already been collected on all panelists, allowing my experiment to be embedded in respondents' normal day-to-day survey usage with only a single additional question. Finally, since implementation did not require any steps to install new software, measurement occurred completely unobtrusively and likely without subjects' conscious awareness (but, of course, with their previous consent). Like the MTurk experiment, this study was designed to maximize ecological validity and minimize the potential for demand effects.

For this experiment I chose a different issue, one that was not completely politicized at the time of the survey but that clearly had the potential to polarize along party lines. The issue was whether to allow Syrian refugees to settle in the United States, and the survey was fielded soon after Secretary of State John Kerry had publicly committed to admitting a substantial number of refugees per year (but before the terrorist attacks in Paris on November 13, 2015). Subjects in the Pulse panel were randomly assigned to be given the following question in the course of filling out a daily YouGov survey (those in the control condition were not shown the question at all): "After several years of war and unrest, a large wave of migrants and refugees from the Middle East, mainly Syria, has been fleeing to Europe. Responding to increasing pressure domestically and from European allies, Secretary

of State John Kerry said last Sunday that the United States would accept 100,000 Syrian refugees per year by 2017. What is your view on the number of refugees the United States should accept from that region per year?"⁴

In fact, the answers to the question are not of interest for the current study and were intended merely to encourage thought about the subject. After answering the question, those in the treatment group saw the following text on a new screen: "Thank you for your response! We may be interested in following up with you in a future survey as developments on the refugee crisis continue. Please keep informing yourself about this issue – just as you normally would for a political topic like this one. We are not looking for any 'right' answer, and feel free to use sources you typically turn to (including on the Internet) for news and information." Notably, subjects were not told that subsequent web visit data may be used to analyze the slant of their media diet. This design is similar to the MTurk experiment except that the survey question itself, rather than an email message, was the encouragement to seek out information on a political topic.

The experiment was placed in the daily survey made available to YouGov's survey panelists over the course of four days from September 25-28, 2015. Typically a maximum of 1,200 panelists will respond to the daily survey on a given weekday, and a fraction of those have the tracking software installed. In total, this procedure returned 451 valid responses from Pulse panelists. However, only 120 of those generated Wakoopa data that could be used to measure outcomes (treatment: 61, control: 59).⁵ This is not a threat to inference since subjects were allocated to treatment using simple random assignment (i.e., as they entered the survey). I used Wakoopa

⁴Possible responses were "An unlimited number," "Many more than 100,000," "Somewhat more than 100,000," "100,000 is about right," "Somewhat less than 100,000," "Much fewer than 100,000," and "None at all."

⁵YouGov investigated the possible sources of this issue, but the likeliest reason is that some respondents were part of the Pulse panel but were either using different devices to access the survey or had uninstalled the software from their computers.

data collected for these respondents over a full week, from September 26-October 2.

In order to measure the partisan lean of subjects' media diets, I used a combination of the procedures from the MTurk experiment and the observational portrait of Wakoopa panelists above. Following the latter investigation, I subsetted subjects' traffic to sites about politics and news and merged the visits to partisan slant scores using comScore data. Following the MTurk study, I averaged across all the slant scores visited by a single respondent to create an individual-level mean media partisanship score.

YouGov Experiment Results

Across two different experiments conducted on separate subject pools, with different issues and distinct measurement approaches, the results are strikingly similar. Figure 3.7 replicates Figure 3.5 almost exactly: We see a heterogeneous effect of the encouragement to seek out information on a political topic, with the distribution of Republican leaners' average media slant moving toward that of Democratic leaners (and Democrats' media diets staying much the same in terms of partisan slant). These results clearly rule out a polarizing effect, since any preexisting bias in media diets—as seen in the distinct peaks of the distributions with dotted lines—is reduced to essentially zero. Against much conventional wisdom and existing theorizing, it appears that this type of intervention is more likely to moderate subjects' media diets than to polarize them according to partisan predispositions.

Table 3.5 reports these results in the form of OLS regressions in which the parameter of interest is the coefficient on the interaction between the treatment indicator and partisanship (most notably here Lean Republican).⁶ The first result is

⁶In these results and in Figure 3.7 I fold "leaners" in with party identification. YouGov provides a three-point party ID scale (Democrat, Republican, Independent), which I augment with the five-point ideological self-placement scale. If a respondent is an "Independent" and either "conser-

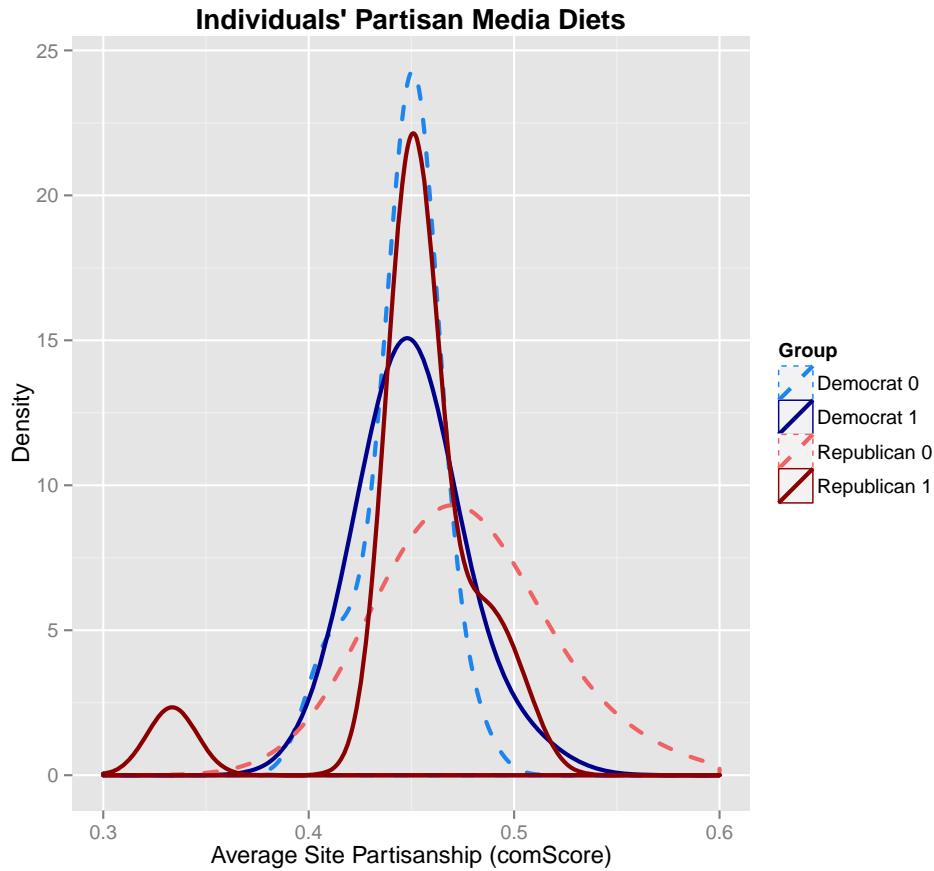


Figure 3.7: Density plot of individuals’ media diet partisanship in the YouGov Pulse experiment, as measured by the mean comScore partisan slant of all sites visited by each respondent post-treatment. Lines denote party and treatment sub-groups. Leaners are coded as partisans.

the lack of a main effect: There is no hypothetical reason why we would expect an overall partisan shift in media diet as opposed to heterogeneous effects (and here, unlike with the MTurk study, the effect is essentially zero). The interaction, however, is estimated to be between $-.03$ and $-.04$, which translates to a sizable leftward shift in the average slant of Republican leaners’ media diets caused by assignment to the treatment encouragement. The magnitude of this shift amounts to roughly 10% of the possible range of media slant scores, and it is statistically significant across all models ($p < 0.05$). Astonishingly, the effect of the treatment is enough

vative” or “very conservative,” I code him or her as a Republican leaner, and likewise for Democrats.

to completely undo the baseline effect of being a Republican who is more likely to visit conservative-leaning news websites in the first place. The precision of the estimate also improves with additional covariate adjustment, as shown in Models 4 and 5 ($p < 0.01$).

It is possible that with other topics, we would see most of the movement among Democrats rather than Republicans. Perhaps most of the information one could encounter on both for-profit education and refugees tended to be on moderate to left-leaning news sources. Theoretically we expect differential effect of treatment by partisan attachment, not necessarily that it is concentrated among one particular group.

Does Exposure Change Attitudes or Knowledge?

Until now, I have examined the effects of the treatment on participants' media habits. In the previous sections, I described how participants can be induced to alter their media diet—in particular, Republicans and conservatives select somewhat more *liberal* sources on average, reducing the conservative lean of their information environment and producing moderation overall. Beyond the immediate effect on exposure to specific content, however, an important question is whether any downstream shifts in attitudes, opinions, or knowledge can be identified.

Table 3.6 summarizes the results, combining responses from the second survey wave in the MTurk study with the original treatment assignment and pre-treatment covariates. It is immediately evident that there are no identifiable effects on views about for-profit colleges⁷ or how to regulate them⁸; the estimated coefficients on the treatment vector are small and highly variable. By contrast, pre-treatment views are

⁷This is a 7-point scale from “For-profit colleges are extremely bad for American education” to “For-profit colleges are extremely good for American education.”

⁸This was a 5-point scale from “The government should regulate for-profit colleges much more strongly” to “The government should regulate for-profit colleges much less strongly,” with more regulation coded as positive.

Table 3.5: Experimental data from YouGov Pulse. Dependent variable is the average partisan slant of each respondent's post-treatment media diet, and the treatment is encouragement to learn about the Syrian refugee issue. There is no main effect, but there is a significant heterogeneous treatment effect among Republican leaners in a more liberal direction.

	<i>Dependent variable:</i>				
	Media Diet Partisanship				
	(1)	(2)	(3)	(4)	(5)
Assigned Treatment	-0.01 (0.01)	-0.01 (0.01)	0.005 (0.01)	0.01 (0.01)	0.01 (0.01)
Independent		0.01 (0.01)	0.01 (0.01)	0.02* (0.01)	0.01 (0.01)
Lean Republican		0.01** (0.01)	0.03*** (0.01)	0.03*** (0.01)	0.03*** (0.01)
Age				0.0004** (0.0002)	0.0002 (0.0002)
Male				-0.004 (0.01)	-0.01 (0.01)
College					-0.01 (0.01)
H.S.					-0.0001 (0.01)
No H.S.					-0.01 (0.02)
Post-grad					-0.02 (0.01)
Some college					-0.01 (0.01)
Treatment x Independent			-0.01 (0.01)	-0.02 (0.01)	-0.01 (0.02)
Treatment x Lean Republican			-0.03** (0.01)	-0.04*** (0.01)	-0.04*** (0.01)
Constant	0.46*** (0.004)	0.45*** (0.004)	0.45*** (0.01)	0.42*** (0.01)	0.40*** (0.03)
<i>Other covariates included?</i>					
Family Income	no	no	no	yes	yes
Race	no	no	no	no	yes
Region	no	no	no	no	yes
Observations	97	96	96	96	96
Adjusted R ²	0.01	0.05	0.09	0.10	0.08

Note:

*p<0.1; **p<0.05; ***p<0.01
OLS, standard errors in parentheses.

highly correlated with post-treatment responses. In the third and fourth columns, however, a more complicated picture emerges for the effects of treatment on knowledge. First, from Column 3 we see that a main effect is not evident. Column 4 shows evidence of heterogeneity by party affiliation, however: Being assigned to receive the encouragement message appears to have had no effect on Democrats ($-0.37 - 0.34 + 0.70 = -0.01$) but actually *reduced* knowledge about for-profit colleges among Republicans ($-0.37 - 0.79 + 1.01 = -0.15$).

These findings are clearly against expectations. Given the limited power of the experiment and the possibility that conditional average treatment effects of this kind may be due to chance, I hesitate to speculate about the significance of these results. Notably, these are intent-to-treat estimates focusing on the effect of being assigned to treatment rather than the effect of *complying* with the treatment—that is, seeking out information about for-profit colleges if and only if encouraged to do so.

Attitude Change: Puzzling Additional Results

Last year, I implemented a pilot version of the same study on MTurk, with minor differences (N=348).⁹ I include the results here to highlight the fact that the findings are not consistent: While I found null effects on opinions above, I found strong effects previously (but none for knowledge). Clearly, more research is needed.

Table 3.7 shows the results of regressions of post-treatment attitudes on treatment assignment, pre-treatment views, partisanship, and interactions. Despite the fact that the treatment did not specify the valence of content to look up, it seemed to have an overall negative effect on subjects' views toward for-profits (possibly as a result of negative press in recent years). Column 1 shows the main effect, a modest

⁹Out of 1,000 initial respondents, 348 both agreed to install the browser widget and successfully did so. From this experimental sample, I then block randomized subjects into either treatment and control. Sample characteristics: 61% male; 79% white, 6% black, 7.5% Hispanic; 47% college-educated; 40.5% Democrat, 14% Republican; mean age, 32.

Table 3.6

	<i>Dependent variable:</i>			
	View	Regulate	Knowledge	
	(1)	(2)	(3)	(4)
Pre-treatment View	0.74*** (0.03)			
Pre-treatment Regulate		0.61*** (0.03)		
Assigned Treatment	-0.25 (0.21)	0.05 (0.13)	0.08 (0.17)	-0.37 (0.26)
Incognito	0.06 (0.47)	0.09 (0.29)	-0.08 (0.59)	-0.16 (0.58)
Chrome	0.25 (0.21)	0.13 (0.13)	0.09 (0.27)	0.17 (0.26)
Democrat	-0.06 (0.26)	0.10 (0.16)		-0.34 (0.32)
Republican	0.27 (0.34)	-0.13 (0.21)		-0.79* (0.41)
Income				0.08** (0.03)
Education				0.08 (0.05)
Treatment x Dem	0.24 (0.29)	-0.11 (0.18)		0.70* (0.36)
Treatment x Rep	0.11 (0.39)	-0.03 (0.24)		1.01** (0.47)
Constant	0.61** (0.29)	1.53*** (0.20)	3.77*** (0.29)	3.35*** (0.40)
Observations	436	436	436	436
Adjusted R ²	0.55	0.49	-0.01	0.03

Note:

*p<0.1; **p<0.05; ***p<0.01
OLS, standard errors in parentheses.

0.38-point decrease along the 7-point attitude scale. Column 2 shows no average difference in views between Democrats or Republicans. Column 3 includes interactions between treatment and the party indicators, which suggest that essentially all of the change in views is driven by Democrats becoming less supportive. Republicans are also estimated to move in the same direction, but this is not significant (likely because the sample had far fewer Republican-leaning subjects). The last three columns show that the heterogeneous effects worked for partisan identification but not ideological self-placement.

Table 3.7: Views toward for-profit colleges, earlier study.

	DV: View (T2)				
	(1)	(2)	(3)	(4)	(5)
Pre-treatment View	0.68*** (0.04)	0.66*** (0.04)	0.68*** (0.04)	0.65*** (0.04)	0.65*** (0.04)
Assigned Treatment	-0.38*** (0.12)	-0.37*** (0.12)	-0.01 (0.17)	-0.43*** (0.12)	-0.19 (0.23)
Democrat		-0.21 (0.13)	0.15 (0.19)		
Republican		0.07 (0.18)	0.35 (0.26)		
Treatment x Dem			-0.69*** (0.26)		
Treatment x Rep			-0.59 (0.37)		
Liberal				-0.33** (0.15)	-0.18 (0.22)
Conservative				-0.15 (0.19)	0.14 (0.26)
Treatment x Lib					-0.25 (0.29)
Treatment x Con					-0.60 (0.37)
Constant	1.33*** (0.15)	1.45*** (0.17)	1.22*** (0.18)	1.67*** (0.21)	1.53*** (0.24)
Observations	318	318	318	304	304
Adjusted R ²	0.51	0.51	0.52	0.51	0.51

Note:

*p<0.1; **p<0.05; ***p<0.01

Weighted regressions (to take into account block randomization), robust standard errors in parentheses.

I show essentially the same results, but for the dependent variable measuring subjects' views about the regulation of for-profit colleges and universities, in Appendix C4. In short, the treatment had a positive effect on subjects' belief in regulating the schools.

In this pilot study, I also tested for effects on knowledge, measured in the post-treatment but not pre-treatment survey, and salience. I found mixed to null results for both, although the estimates for knowledge came close to statistical significance. It is possible that improving the efficiency of the estimates with pretreatment covariates would have allowed me to identify significant effects on knowledge, although I specifically sought to avoid priming subjects with knowledge questions before the treatment was administered.

One possible reason for the discrepancy in findings is a change in the way I recorded the primary dependent variables in the second survey wave. In the pilot study, I simply asked subjects who returned for the second wave their opinions on for-profit colleges and how to regulate them, along with a set of knowledge questions. The attitude and policy opinion questions were worded identically to the ones I asked in the pre-treatment survey. It is possible that treated respondents remembered the previous questions and sought to give the "correct" (more critical) answers given the tenor of the coverage they may have exposed themselves to. In the later MTurk experiment, by contrast, I included the same questions in the second wave—but, as in the pre-treatment survey, I grouped them in with questions on unrelated subjects in an attempt to mask the goal of the study. Thus, it is possible that the earlier findings were essentially the result of demand effects (see, e.g., Green, Calfano and Aronow 2014).

Discussion

In this two-part investigation, I show that—at least among respondents in an opt-in sample—there is no strong evidence that individuals cluster into ideologically segregated information cocoons. Aggregated over more than 100,000 visits, the ideological flavor of most content respondents exposed themselves to is generally centrist. Rather than the strong bimodal prediction made by theorists who express concerns about online echo chambers, the most-visited media sources heavily cluster in the middle. One possible caveat to this conclusion is suggestive evidence of a smaller group of conservatives who consume primarily right-leaning content, but this would have to be replicated with either a more representative sample or with properly weighted data.¹⁰

I couple this observational portrait of panelists' balanced media diets with two experimental designs that likewise do not reveal strong evidence of ideological segregation. In particular, Figure 3.6 shows how aggregate site visits cluster around the center (with another small bump on the right corresponding to Drudge). Randomly inducing participants to educate themselves about a novel, non-salient political issue does not induce patterns of partisan selective exposure; if anything, there is evidence of a heterogeneous effect in which Republicans expose themselves to more liberal content, resulting in overall moderation of media diets.

Regardless of the patterns of exposure I document, evidence of media effects remains elusive. In the main experimental study, I find no measurable effects of being encouraged to seek out information about for-profit colleges on attitudes about the issue or opinions about how to regulate them. There is only suggestive evidence of possible differential effects on knowledge, although it is difficult to interpret. In the pilot study, I found positive effects on attitudes and opinions, but the dependent

¹⁰In the version of the figure produced using the partisan, rather than the ideological, measure of media slant (shown in Appendix C3), there is also a bump on the center-left corresponding to Democrats visiting Talking Points Memo.

variable in wave 2 was collected in a fashion that may have encouraged demand effects for subjects in the treatment group.

It is possible that a different choice of issue might have resulted in different findings. For instance, an issue that is more salient, better maps onto the partisan divide, or is more likely to trigger “hot cognition” may have induced participants to seek out reinforcing information (Redlawsk 2002; Lodge and Taber 2005). Issues of this type may generate a different pattern of results than those reported here. I note that while further tests should be undertaken, the current evidence of moderation on relatively non-salient, potentially cross-cutting policy issues directly addresses the greatest concerns of the informational cocoon theorists. For if even novel information about underpoliticized issues reverberate in mutually exclusive echo chambers, then the prospects for deliberation and consensus are especially dire.

As a result of treatment, subjects likely relied on the kinds of low-cost defaults that typically direct readers to mainstream and centrist (rather than niche and extreme) sources: search engines, homepages, and bookmarks. The results suggest that subjects who sought out information about a complex, non-salient political topic were not generally following a directional motivation or seeking out only congenial content (Kruglanski 1999). Or, at least, that motivation was not strong enough to overcome the cost of overcoming ingrained media consumption habits and the defaults built into the typical Internet user’s daily browsing environment. Perhaps, then, rather than facilitating invisible “filter bubbles,” defaults can actually serve as a moderating filter for new information (Pariser 2011).

Aside from the particular mechanism of the exposure findings, there are multiple potential explanations for the lack of observed media effects. One possibility, as mentioned before, is that some sort of demand effect was at play in which respondents in the pilot study’s treatment group sought to give the “preferred” answer to the post-treatment questions about for-profit colleges and how to regulate them.

In essence, subjects who remembered having been asked to research the issue may have connected that task to the survey items. This would have been less of an issue with the main study, which embedded the primary questions of interest in a larger survey with unrelated items.

Still, if the findings from the study represent the “true” effect, then the question of why effects are muted or nonexistent remains. The era of “minimal effects” in which mass-media messages were thought to filter through community intermediaries and social groups may not resemble 21st-century America (Klapper 1960; Putnam 1995). But on the other hand, people are in some ways more networked than ever—at least online. Has the nature of the intermediaries simply changed? Another possible explanation is that some form of “transactive memory” is in play, in which people effectively “outsource” their knowledge to the Internet. Psychologists have argued, in essence, that a form of mental division of labor can take place within couples and among groups (Wegner, Giuliano and Hertel 1985; Hollingshead 1998; Kozlowski and Ilgen 2006), and recent research has applied this insight to the modern-day reality of ubiquitous access to reference sources online (Fisher, Goddu and Keil 2015). Given these findings, it seems plausible that subjects who complied with the encouragement treatment simply did not retain any new information (as measured by the knowledge items)—even if they felt they had learned about the topic of for-profit education. A final possibility is that there were effects, but that they were simply too fleeting to be captured in the follow-up survey, which for some subjects was several days after receiving the last encouragement message (e.g., Gerber et al. 2011).

As scholars continue to design studies of information dynamics and media choice, especially in online environments, these findings will hopefully reinforce the importance of taking into account the difficult-to-measure contextual factors that guide politically relevant information-seeking behavior.

Bibliography

- Aaker, Jennifer and Satoshi Akutsu. 2009. "Why do people give? The role of identity in giving." .
- Albertson, Bethany and Adria Lawrence. 2009. "After the Credits Roll The Long-Term Effects of Educational Television on Public Knowledge and Attitudes." *American Politics Research* 37(2):275–300.
- Althaus, Scott L. and David H. Tewksbury. 2007. "Toward a New Generation of Media Use Measures for the ANES." *Report to the Board of Overseers of the ANES* .
- Arceneaux, K. and M. Johnson. 2010. Does Media Fragmentation Produce Mass Polarization? Selective Exposure and a New Era of Minimal Effects.
- Barberá, Pablo. 2014. "How Social Media Reduces Mass Political Polarization. Evidence from Germany, Spain, and the U.S." *Working Paper* .
URL: j.mp/BarberaPolarization
- Barr, Dermot and John Drury. 2009. "Activist identity as a motivational resource: Dynamics of (dis) empowerment at the G8 direct actions, Gleneagles, 2005." *Social Movement Studies* 8(3):243–260.
- Bartels, Larry M. 1993. "Messages Received: The Political Impact of Media Exposure." *American Political Science Review* 87(2):267–285.
- Belli, Robert F, Michael W Traugott and Matthew N Beckmann. 2001. "What leads to voting overreports? Contrasts of overreporters to validated voters and admitted nonvoters in the American National Election Studies." *Journal of Official Statistics* 17(4):479–498.
- Benkler, Yochai. 2006. *The wealth of networks: How social production transforms markets and freedom*. New Haven, CT: Yale University Press.
- Bennett, W Lance and Alexandra Segerberg. 2012. "The logic of connective action: Digital media and the personalization of contentious politics." *Information, Communication & Society* 15(5):739–768.

- Bennett, W.L. and S. Iyengar. 2008. "A new era of minimal effects? The changing foundations of political communication." *Journal of Communication* 58(4):707–731.
- Berinsky, Adam, Gregory Huber and Gabriel Lenz. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* 20(3).
- Berinsky, Adam J., Michele F. Margolis and Michael W. Sances. 2013. "Separating the Shirkers from the Workers? Making Sure Respondents Pay Attention on Self-Administered Surveys." *American Journal of Political Science* .
- Bimber, Bruce, Andrew J Flanagin and Cynthia Stohl. 2005. "Reconceptualizing Collective Action in the Contemporary Media Environment." *Communication Theory* 15(4):365–388.
- Bond, Robert M, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle and James H Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489(7415):295–298.
- Bowers, Jake, Mark M. Fredrickson and Costas Panagopoulos. 2013. "Reasoning about interference between units: a general framework." *Political Analysis* 21(1):97–124.
- Bryan, Christopher J., Gregory M. Walton, Todd Rogers and Carol S. Dweck. 2011. "Motivating voter turnout by invoking the self." *Proceedings of the National Academy of Sciences* 108(31):12653–12656.
URL: <http://www.pnas.org/content/108/31/12653.abstract>
- Burton, Scot and Edward Blair. 1991. "Task Conditions, Response Formulation Processes, and Response Accuracy for Behavioral Frequency Questions in Surveys." *Public Opinion Quarterly* 55(1):50–79.
- Conover, Michael, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer and Alessandro Flammini. 2011. Political Polarization on Twitter. In ICWSM.
- Converse, Philip E. 1964. The Nature of Belief Systems in Mass Publics. In *Ideology and Discontent*, ed. David E. Apter. Ann Arbor: University of Michigan Press.
- Coppock, Alexander. 2014. "Information Spillovers: Another Look at Experimental Estimates of Legislator Responsiveness." *Journal of Experimental Political Science* .
- Dilliplane, Susanna, Seth K. Goldman and Diana C. Mutz. 2013. "Televised Exposure to Politics: New Measures for a Fragmented Media Environment." *American Journal of Political Science* 57(1):236–248.
- Drury, John, Christopher Cocking, Joseph Beale, Charlotte Hanson and Faye Rapley. 2005. "The phenomenology of empowerment in collective action." *British Journal of Social Psychology* 44(3):309–328.

- Farrell, Henry. 2012. "The Consequences of the Internet for Politics." *Annual Review of Political Science* 15:35–52.
- Fiorina, M.P., S.J. Abrams and J. Pope. 2011. *Culture War?: The Myth of a Polarized America*. Great questions in politics series Longman.
URL: <http://books.google.com/books?id=s5YZQQAACAAJ>
- Fisher, Matthew, Mariel K Goddu and Frank C Keil. 2015. "Searching for Explanations: How the Internet Inflates Estimates of Internal Knowledge."
- Fisher, R A. 1925. *Statistical Methods for Research Workers*. London: Oliver and Boyd.
- Flaxman, Seth, Sharad Goel and Justin M Rao. 2013. "Ideological segregation and the effects of social media on news consumption." *Available at SSRN 2363701* .
- Fowler, James H., Michael T. Heaney, David W. Nickerson, John F. Padgett and Betsy Sinclair. 2011. "Causality in political networks." *American Politics Research* 39(2):437–480.
- Gaby, Sarah and Neal Caren. 2012. "Occupy Online: How Cute Old Men and Malcolm X Recruited 400,000 US Users to OWS on Facebook." *Social Movement Studies* .
- Garrett, R Kelly. 2009. "Echo chambers online?: Politically motivated selective exposure among Internet news users1." *Journal of Computer-Mediated Communication* 14(2):265–285.
- Garrett, R Kelly. 2013. "Selective exposure: New methods and new directions." *Communication Methods and Measures* 7(3-4):247–256.
- Gentzkow, Matthew and Jesse M. Shapiro. 2011. "Ideological Segregation Online and Offline." *The Quarterly Journal of Economics* (126):1799–1839.
- Gerber, Alan S. and Donald P. Green. 2000. "The Effects of Personal Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94.
- Gerber, Alan S. and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation*. W. W. Norton.
- Gerber, Alan S., James G. Gimpel, Donald P. Green and Daron 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):135–150.
- Gladwell, Malcolm. 2010. "Small Change: Why the revolution will not be tweeted." *The New Yorker* .
URL: <http://www.newyorker.com/magazine/2010/10/04/small-change-3>

- Gong, Shiyang, Juanjuan Zhang, Ping Zhao and Xuping Jiang. 2014. "Tweets and Sales." *SSRN Working Paper* .
- Graber, Doris A. 1997. *Mass Media and American Politics*. Washington, DC: CQ Press.
- Green, Donald P, Brian R Calfano and Peter M Aronow. 2014. "Field experimental designs for the study of media effects." *Political Communication* 31(1):168–180.
- Guess, Andrew M. 2015. "Measure for Measure: An Experimental Test of Online Political Media Exposure." *Political Analysis* .
URL: <http://pan.oxfordjournals.org/content/early/2014/06/20/pan.mpu010.abstract>
- Hindman, Matthew. 2008. *The Myth of Digital Democracy*. Princeton University Press.
- Hollingshead, Andrea B. 1998. "Communication, learning, and retrieval in trans-active memory systems." *Journal of Experimental Social Psychology* 34(5):423–442.
- Humphreys, Macartan, Raul Sanchez de la Sierra and Peter van der Windt. 2013. "Fishing, Commitment, and Communication: A Proposal for Comprehensive Nonbinding Research Registration." *Political Analysis* 21(1):1–20.
- Iyengar, Shanto. 2010. Experimental Designs for Political Communication Research: Using New Technology and Online Participant Pools to Overcome the Problem of Generalizability. In *Sourcebook for Political Communication Research: Methods, Measures, and Analytical Techniques*, ed. Erik P. Bucy and R. Lance Holbert. Taylor and Francis.
- Iyengar, Shanto and Donald R. Kinder. 1987. *News that Matters: Television and American Opinion*. Chicago: University of Chicago Press.
- 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(01):186–200.
- Jamieson, Kathleen Hall and Joseph N Cappella. 2008. *Echo chamber: Rush Limbaugh and the conservative media establishment*. Oxford University Press.
- Karpf, David. 2010. "Online Political Mobilization from the Advocacy Group's Perspective: Looking Beyond Clicktivism." *Policy & Internet* 2(4):7–41.
- King, Gary. 2011. "Ensuring the Data-Rich Future of the Social Sciences." *Science* 331(6018):719–721.
URL: <http://www.sciencemag.org/content/331/6018/719.abstract>
- Klapper, Joseph T. 1960. "The Effects of Mass Communication."
- Kobayashi, Tetsuro and Yu Ichifuji. 2014. "Tweets that matter: Evidence from a randomized field experiment in Japan." *Annual Meeting of the Southern Political Science Association* .

- Kozlowski, Steve WJ and Daniel R Ilgen. 2006. "Enhancing the effectiveness of work groups and teams." *Psychological science in the public interest* 7(3):77–124.
- Krosnick, Jon A. 1991. "Response strategies for coping with the cognitive demands of attitude measures in surveys." *Applied Cognitive Psychology* 5(3):213–236.
- Krueger, Brian S. 2006. "A Comparison of Conventional and Internet Political Mobilization." *American Politics Research* 34(6):759–776.
- Kruglanski, Arie W. 1999. "Motivation, cognition, and reality: Three memos for the next generation of research." *Psychological Inquiry* 10(1):54–58.
- LaCour, Michael J. 2013. "A Balanced News Diet, Not Selective Exposure: Results from Erie to Arbitron." *American Political Science Association Annual Meeting, Chicago* .
- Lazarsfeld, Paul F., Bernard Berelson and Hazel Gaudet. 1944. *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign*. New York: Columbia U. Press.
- Leeper, Thomas. 2013. "Crowdsourcing with R and the MTurk API." *The Political Methodologist* 20(2):2–7.
- Levendusky, Matthew S. 2013. "Why do partisan media polarize viewers?" *American Journal of Political Science* 57(3):611–623.
- Lodge, Milton and Charles S Taber. 2005. "The automaticity of affect for political leaders, groups, and issues: An experimental test of the hot cognition hypothesis." *Political Psychology* 26(3):455–482.
- Lupia, Arthur and Gisela Sin. 2003. "Which public goods are endangered?: How evolving communication technologies affect the logic of collective action." *Public Choice* 117(3-4):315–331.
- Marwick, Alice E and danah boyd. 2011. "I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience." *New Media & Society* 13(1):114–133.
- McPherson, Miller, Lynn Smith-Lovin and James M. Cook. 2001. "Birds of a feather: Homophily in social networks." *Annual review of sociology* pp. 415–444.
- Morozov, Evgeny. 2009. "Iran: Downside to the 'Twitter Revolution'." *Dissent* .
- Mullainathan, S. and A. Shleifer. 2005. "The market for news." *American Economic Review* 95(4):1031–1053.
- Munson, Sean, Stephanie Y. Lee and Paul Resnick. 2013. "Encouraging Reading of Diverse Political Viewpoints with a Browser Widget." *International Conference on Weblogs and Social Media* .

- Mutz, D.C. and L. Young. 2011. "Communication and Public Opinion." *Public Opinion Quarterly* 75(5):1018–1044.
- Mutz, Diana C. 2006. *Hearing the other side: Deliberative versus participatory democracy*. Cambridge University Press.
- Negroponte, Nicholas. 1995. "Being Digital."
- Nickerson, David W. 2007. "The ineffectiveness of e-vites to democracy field experiments testing the role of e-mail on voter turnout." *Social Science Computer Review* 25(4):494–503.
- Obar, Jonathan, Paul Zube and Clifford Lampe. 2012. "Advocacy 2.0: An Analysis of How Advocacy Groups in the United States Perceive and Use Social Media as Tools for Facilitating Civic Engagement and Collective Action." *Journal of Information Policy* 2(0).
URL: <http://jip.vmhost.psu.edu/ojs/index.php/jip/article/view/80/47>
- Olson, Mancur. 1965. *The logic of collective action*. Vol. 124 Harvard University Press.
- Pariser, Eli. 2011. *The Filter Bubble: What the Internet Is Hiding From You*. Penguin UK.
- Pew. 2012. "In Changing News Landscape, Even Television is Vulnerable."
- Pew. 2014. "5 key findings about digital news audiences."
- Popkin, Samuel L. 1991. *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. Chicago: University of Chicago Press.
- Porter, Stephen R. and Michael E. Whitcomb. 2003. "The impact of contact type on web survey response rates." *Public Opinion Quarterly* pp. 579–588.
- Price, Vincent and John Zaller. 1993. "Who Gets the News? Alternative Measures of News Reception and Their Implications for Research." *Public Opinion Quarterly* 57(2):133–164.
- Prior, M. 2007. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press.
- Prior, M. 2009a. "The immensely inflated news audience: Assessing bias in self-reported news exposure." *Public Opinion Quarterly* 73(1):130–143.
- Prior, M. and A. Lupia. 2008. "Money, time, and political knowledge: Distinguishing quick recall and political learning skills." *American Journal of Political Science* 52(1):169–183.
- Prior, Markus. 2009b. "Improving Media Effects Research through Better Measurement of News Exposure." *Journal of Politics* 71(3):893–908.

- Putnam, Robert P. 1995. "Bowling Alone: America's Declining Social Capital." *Journal of Democracy* 6:65–78.
- Redlawsk, David P. 2002. "Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making." *The Journal of Politics* 64(04):1021–1044.
- Rivers, Douglas. 2006. "Sample matching: Representative sampling from internet panels." *Polimetrix White Paper Series*.
- Romantan, Anca, Robert Hornik, Vincent Price, Joseph Cappella and K Viswanath. 2008. "A comparative analysis of the performance of alternative measures of exposure." *Communication Methods and Measures* 2(1-2):80–99.
- Rosenstone, Steven and John M. Hansen. 1993. *Mobilization, Participation and Democracy in America*. MacMillan Publishing.
- Sears, David O. and Jonathan L. Freedman. 1967. "Selective Exposure to Information: A Critical Review." *The Public Opinion Quarterly* 31(2):194–213.
- Shapiro, Robert Y. 2013. "Hearing the Opposition: It Starts at the Top." *Critical Review* 25(2):226–244.
- Shirky, Clay. 2008. *Here Comes Everybody: The Power of Organizing Without Organizations*. Penguin Group.
- Shulman, Stuart W. 2009. "The case against mass e-mails: Perverse incentives and low quality public participation in US federal rulemaking." *Policy & Internet* 1(1):23–53.
- Sinclair, Betsy, Margaret McConnell and Donald P Green. 2012. "Detecting spillover effects: Design and analysis of multilevel experiments." *American Journal of Political Science* 56(4):1055–1069.
- Smyth, Jolene, Don Dillman, Leah Christian and Michael Stern. 2006. "Comparing Check-All and Forced-Choice Question Formats in Web Surveys." *Public Opinion Quarterly* 70(1):66–77.
- Stroud, N.J. 2008. "Media use and political predispositions: Revisiting the concept of selective exposure." *Political Behavior* 30(3):341–366.
- Sudman, Seymour and Norman M Bradburn. 2012. *Asking Questions: A Practical Guide to Questionnaire Design*. San Francisco: Jossey-Bass, 1982.
- Sudman, Seymour, Norman M. Bradburn and Norbert Schwarz. 1996. *Thinking About Answers: The Application of Cognitive Processes to Survey Methodology*. Jossey-Bass.
- Sunstein, Cass R and Richard Thaler. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press New Haven.

- Sunstein, C.R. 2007. *Republic.com 2.0*. Princeton University Press.
- Taylor, Sean J, Lev Muchnik and Sinan Aral. 2014. "Identity and Opinion: A Randomized Experiment." *SSRN* .
- Teresi, Holly and Melissa R Michelson. 2014. "Wired to mobilize: The effect of social networking messages on voter turnout." *The Social Science Journal* .
- Tourangeau, Roger, Lance J. Rips and Kenneth Rasinski. 2000. *The Psychology of Survey Response*. Cambridge: Cambridge University Press.
- Verba, Sidney, Kay Lehman Schlozman and Henry E. Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics*. Cambridge: Harvard University Press.
- Verplanken, Bas and Rob W Holland. 2002. "Motivated decision making: effects of activation and self-centrality of values on choices and behavior." *Journal of personality and social psychology* 82(3):434.
- Wegner, Daniel M, Toni Giuliano and Paula T Hertel. 1985. Cognitive interdependence in close relationships. In *Compatible and incompatible relationships*. Springer pp. 253–276.

Appendices

Appendix A1: Survey Details

Check-all condition

Which of these websites have you visited or used in the past 30 days for news, if any?

Select ALL answers that apply.

- | | | | |
|--------------------------------------|--|--|---|
| <input type="checkbox"/> ABCNews.com | <input type="checkbox"/> Daily Mail | <input type="checkbox"/> NBCNews.com | <input type="checkbox"/> Slate |
| <input type="checkbox"/> AOL.com | <input type="checkbox"/> Drudge Report | <input type="checkbox"/> NewRepublic.com | <input type="checkbox"/> ThinkProgress |
| <input type="checkbox"/> BBC News | <input type="checkbox"/> FoxNews.com | <input type="checkbox"/> NPR.org | <input type="checkbox"/> USAToday.com |
| <input type="checkbox"/> BuzzFeed | <input type="checkbox"/> Google News | <input type="checkbox"/> NYTimes.com | <input type="checkbox"/> WashingtonPost.com |
| <input type="checkbox"/> CBSNews.com | <input type="checkbox"/> Huffington Post | <input type="checkbox"/> Politico | <input type="checkbox"/> WSJ.com |
| <input type="checkbox"/> CNN.com | <input type="checkbox"/> MSN | <input type="checkbox"/> Reddit | <input type="checkbox"/> Yahoo! News |
| <input type="checkbox"/> Daily Kos | <input type="checkbox"/> NationalJournal.com | <input type="checkbox"/> RedState | <input type="checkbox"/> Other (please specify)
<input type="text"/> |

Open-ended condition

Please list any websites or blogs that you have visited in the past 30 days for news. **Take some time to ensure that you think of all the sites you have visited.**

Forced-choice, yes/no condition

Which of these websites have you visited or used in the past 30 days for news, if any?

	Have you visited this site?	
	Yes	No
ABCNews.com	<input type="radio"/>	<input type="radio"/>
AOL.com	<input type="radio"/>	<input type="radio"/>
BBC News	<input type="radio"/>	<input type="radio"/>
Buzzfeed	<input type="radio"/>	<input type="radio"/>
CBSNews.com	<input type="radio"/>	<input type="radio"/>
CNN.com	<input type="radio"/>	<input type="radio"/>

Appendix A2: List of Sites Included in L.C.T.

1 <http://abcnews.go.com/>
2 <http://america.aljazeera.com/>
3 <http://digbysblog.blogspot.com/>
4 <http://dish.andrewsullivan.com/>
5 <http://krugman.blogs.nytimes.com>
6 <http://latino.foxnews.com/index.html>
7 <http://nbcpolitics.nbcnews.com/>
8 <http://news.google.com/> <https://news.google.com/news/section?pz=1&cf=all&topic=e&>
9 <http://news.msn.com/>
10 <http://news.yahoo.com>
11 <http://newswatch.nationalgeographic.com/>
12 <http://www.about.com/> <http://about.com/> <http://www.about.com/newsissues/>
13 <http://www.aljazeera.com> <http://aljazeera.com>
14 <http://www.americanthinker.com> <http://americanthinker.com>
15 <http://www.ann Coulter.com> <http://ann Coulter.com>
16 <http://www.antiwar.com> <http://antiwar.com>
17 <http://www.aol.com/>
18 <http://www.ap.org/> <http://ap.org/>
19 <http://www.bbc.com/news/> <http://bbc.com/news/>
20 http://www.bbc.com/news/world/us_and_canada/ <http://www.bbc.co.uk/news/>
21 <http://www.beforeitsnews.com> <http://beforeitsnews.com>
22 <http://www.billoreilly.com> <http://billoreilly.com>
23 <http://www.blackamericaweb.com> <http://blackamericaweb.com>
24 <http://www.bloomberg.com> <http://bloomberg.com>
25 <http://www.breitbart.com> <http://breitbart.com>
26 <http://www.businessinsider.com> <http://businessinsider.com>
27 <http://www.buzzfeed.com> <http://www.buzzfeed.com/politics>
28 <http://www.buzzya.com> <http://buzzya.com>
29 <http://www.c-span.org> <http://c-span.org>
30 <http://www.cbsnews.com/> <http://cbsnews.com/>
31 <http://www.cbsnews.com/politics/>
32 <http://www.chacha.com/category/politics>
33 <http://www.chicagosuntimes.com> <http://chicagosuntimes.com>
34 <http://www.chicagotribune.com> <http://chicagotribune.com>
35 <http://www.cnn.com> <http://cnn.com>
36 <http://www.cnn.com/POLITICS/>
37 <http://www.commondreams.org> <http://commondreams.org>
38 <http://www.counterpunch.org> <http://counterpunch.org>
39 <http://www.crookedtimber.org/> <http://crookedtimber.org/>
40 <http://www.crooksandliars.com/> <http://crooksandliars.com/>
41 <http://www.csmonitor.com/> <http://csmonitor.com/>
42 <http://www.dailycaller.com> <http://dailycaller.com>
43 <http://www.dailydot.com> <http://dailydot.com>
44 <http://www.dailykos.com> <http://dailykos.com>
45 <http://www.dailymail.co.uk/> <http://www.dailymail.co.uk/ushome/index.html>
46 <http://www.democratichub.com> <http://democratichub.com>
47 <http://www.democraticunderground.com> <http://democraticunderground.com>
48 <http://www.denverpost.com> <http://denverpost.com>
49 <http://www.dickmorris.com> <http://dickmorris.com>
50 <http://www.drudgereport.com> <http://drudgereport.com>
51 <http://www.economist.com/> <http://economist.com/>
52 <http://www.eschatonblog.com/> <http://eschatonblog.com/>
53 <http://www.examiner.com> <http://examiner.com>
54 <http://www.factcheck.org> <http://factcheck.org>
55 <http://www.firsttoknow.com> <http://firsttoknow.com>
56 <http://www.fivethirtyeight.com/> <http://fivethirtyeight.com/> <http://fivethirtyeight.blogs.nytimes.com/>
57 <http://www.forbes.com> <http://forbes.com>
58 <http://www.foreignaffairs.com> <http://foreignaffairs.com>
59 <http://www.foreignpolicy.com> <http://foreignpolicy.com>
60 <http://www.foxnews.com> <http://foxnews.com>
61 <http://www.foxnews.com/politics/index.html>
62 <http://www.foxnewsinsider.com> <http://foxnewsinsider.com>
63 <http://www.freebeacon.com> <http://freebeacon.com>
64 <http://www.freepatriot.org> <http://freepatriot.org>
65 <http://www.freerepublic.com/> <http://freerepublic.com/>
66 <http://www.frontpagemag.com> <http://frontpagemag.com>
67 <http://www.gopusa.com> <http://gopusa.com>
68 <http://www.govexec.com> <http://govexec.com>
69 <http://www.govtrack.us> <http://govtrack.us>
70 <http://www.hotair.com> <http://hotair.com>
71 <http://www.huffingtonpost.com> <http://huffingtonpost.com>
72 <http://www.huffingtonpost.com/politics/>
73 <http://www.humanevents.com> <http://humanevents.com>
74 <http://www.instapundit.com> <http://instapundit.com>
75 <http://www.inthesetimes.com> <http://inthesetimes.com>
76 <http://www.latimes.com> <http://latimes.com>
77 <http://www.littlegreenfootballs.com/> <http://littlegreenfootballs.com/>
78 <http://www.mediaite.com> <http://mediaite.com>
79 <http://www.michellemalkin.com> <http://michellemalkin.com>
80 <http://www.mikehuckabee.com> <http://mikehuckabee.com>

81 <http://www.motherjones.com> <http://motherjones.com>
82 <http://www.mrc.org> <http://mrc.org>
83 <http://www.mrconservative.com> <http://mrconservative.com>
84 <http://www.mtvu.com> <http://mtvu.com>
85 <http://www.nationalgeographic.com/> <http://nationalgeographic.com/>
86 <http://www.nationaljournal.com> <http://nationaljournal.com>
87 <http://www.nationalmemo.com> <http://nationalmemo.com>
88 <http://www.nationalreview.com> <http://nationalreview.com>
89 <http://www.nbcnews.com/> <http://nbcnews.com/>
90 <http://www.newrepublic.com> <http://newrepublic.com>
91 <http://www.newsbusters.org> <http://newsbusters.org>
92 <http://www.newser.com> <http://newser.com>
93 <http://www.newsmax.com/> <http://newsmax.com/>
94 <http://www.newyorker.com> <http://newyorker.com>
95 <http://www.npr.org> <http://npr.org>
96 <http://www.nydailynews.com> <http://nydailynews.com>
97 <http://www.nypost.com> <http://nypost.com>
98 <http://www.nytimes.com> <http://nytimes.com>
99 <http://www.obamacarefacts.com> <http://obamacarefacts.com>
100 <http://www.outsidethebeltway.com> <http://outsidethebeltway.com>
101 <http://www.pbs.org/newshour/> <http://pbs.org/newshour/>
102 <http://www.pbs.org/wgbh/pages/frontline/>
103 <http://www.politicalwire.com> <http://politicalwire.com>
104 <http://www.politico.com> <http://politico.com>
105 <http://www.politicususa.com> <http://politicususa.com>
106 <http://www.politifact.com> <http://politifact.com>
107 <http://www.powerlineblog.com/> <http://powerlineblog.com/>
108 <http://www.reagancoalition.com> <http://reagancoalition.com>
109 <http://www.realclearpolitics.com> <http://realclearpolitics.com>
110 <http://www.realclearworld.com> <http://realclearworld.com>
111 <http://www.reason.com/blog/> <http://reason.com/blog/>
112 <http://www.reddit.com/> <http://www.reddit.com/r/politics/>
113 <http://www.redstate.com> <http://redstate.com>
114 <http://www.rightwingnews.com> <http://rightwingnews.com>
115 <http://www.rushlimbaugh.com/> <http://rushlimbaugh.com/>
116 <http://www.salon.com> <http://salon.com>
117 <http://www.samuel-warde.com> <http://samuel-warde.com>
118 <http://www.slate.com/> <http://slate.com/>
119 <http://www.slate.com/blogs/moneybox.html> <http://slate.com/blogs/moneybox.html>
120 <http://www.smithsonianmag.com> <http://smithsonianmag.com>
121 <http://www.sourcwatch.org> <http://sourcwatch.org>
122 <http://www.stormfront.org> <http://stormfront.org>
123 <http://www.takepart.com> <http://takepart.com>
124 <http://www.talkingpointsmemo.com> <http://talkingpointsmemo.com>
125 <http://www.the-american-interest.com> <http://the-american-interest.com>
126 <http://www.theatlantic.com/> <http://theatlantic.com/>
127 <http://www.theatlanticwire.com/> <http://theatlanticwire.com/>
128 <http://www.theblaze.com> <http://theblaze.com>
129 <http://www.thedailybeast.com> <http://thedailybeast.com>
130 <http://www.thegrio.com> <http://thegrio.com>
131 <http://www.theguardian.com/> <http://www.theguardian.com/politics> <http://theguardian.com/>
132 <http://www.thehill.com> <http://thehill.com>
133 <http://www.themonkeycage.org> <http://themonkeycage.org> <http://www.washingtonpost.com/blogs/monkey-cage/>
134 <http://www.thenation.com> <http://thenation.com>
135 <http://www.theweek.com> <http://theweek.com>
136 <http://www.thinkprogress.org/> <http://thinkprogress.org/>
137 <http://www.townhall.com> <http://townhall.com>
138 <http://www.trueactivist.com> <http://trueactivist.com>
139 <http://www.truthout.org> <http://truthout.org>
140 <http://www.twitchy.com> <http://twitchy.com>
141 <http://www.upi.com> <http://upi.com>
142 <http://www.usatoday.com> <http://usatoday.com>
143 <http://www.usnews.com/> <http://usnews.com/>
144 <http://www.voanews.com/> <http://voanews.com/>
145 <http://www.volokh.com/> <http://volokh.com/>
146 <http://www.washingtonexaminer.com/> <http://washingtonexaminer.com/>
147 <http://www.washingtonmonthly.com> <http://washingtonmonthly.com>
148 <http://www.washingtonpost.com> <http://washingtonpost.com>
149 <http://www.washingtonpost.com/blogs/wonkblog/> <http://washingtonpost.com/blogs/wonkblog/>
150 <http://www.washingtontimes.com/> <http://washingtontimes.com/>
151 <http://www.weeklystandard.com/> <http://weeklystandard.com/>
152 <http://www.westernjournalism.com> <http://westernjournalism.com>
153 <http://www.wonkette.com/> <http://wonkette.com/>
154 <http://www.worldnetdaily.com/> <http://worldnetdaily.com/>
155 <http://www.wsj.com> <http://wsj.com>

Appendix A3: Supplementary Graphs and Tables




Figure A1: Total number of visits per site across all treatment conditions. Top 50 sites in reverse order, pooling data across Studies 1 and 3.

Difference between self-reported and actual exposure: All 3 question types




Figure A2: Plot, with Loess lines, of the difference between total reported and total actual visits per site for each survey question type.

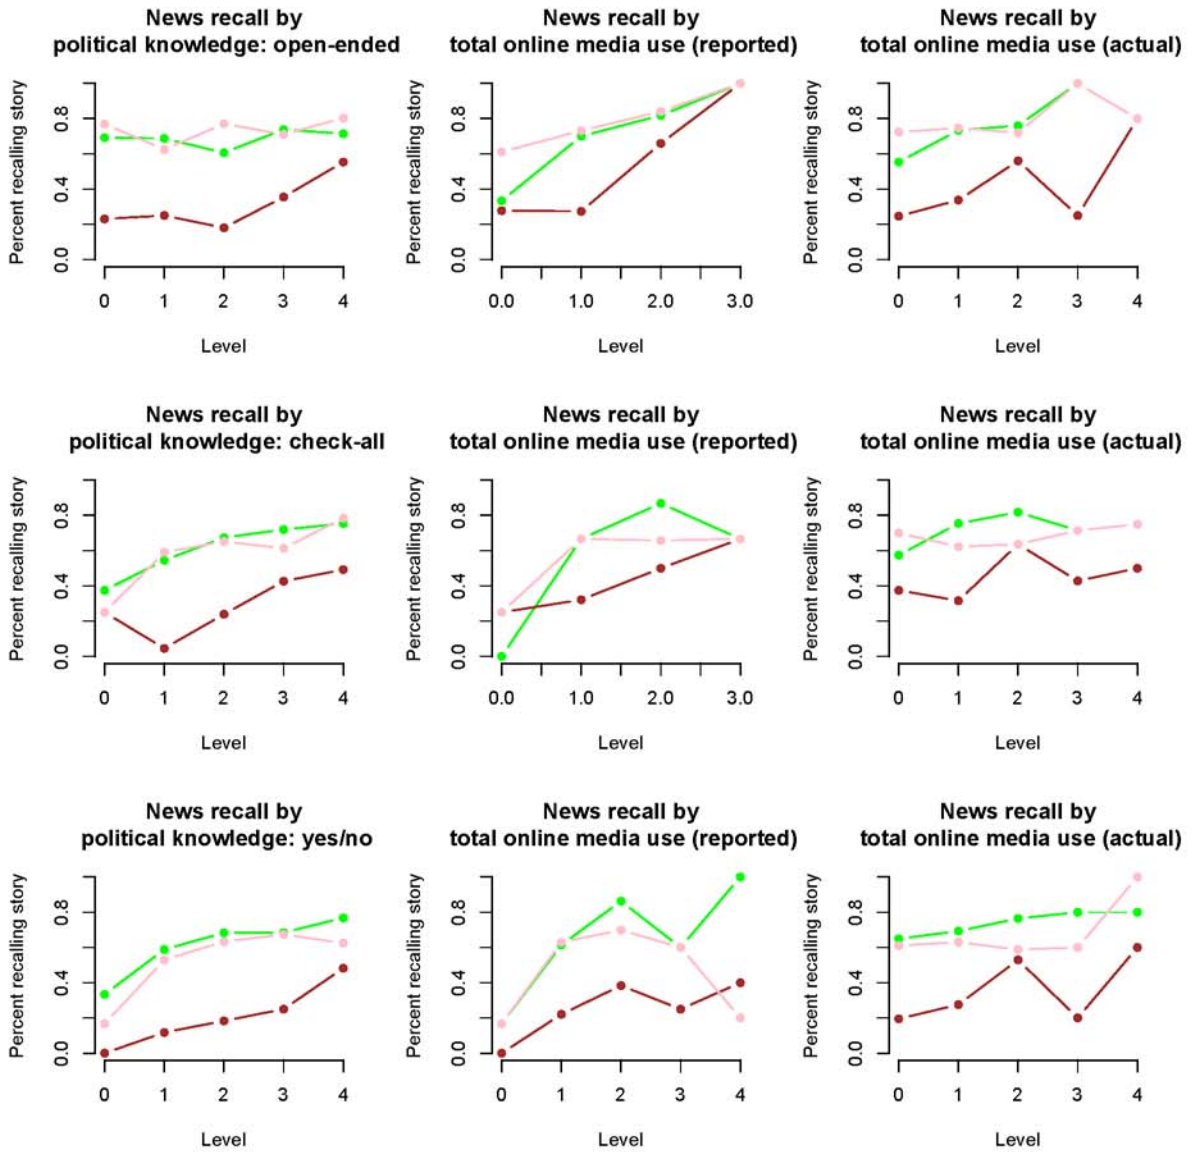


Figure A3: News recall broken down by individual story and predictor. The x-axis plots, from left to right, general knowledge, reported media exposure, and actual media exposure as measured by the *Link Classification Task*. The y-axis plots the average number of stories recalled, out of 3. Each row represents a different treatment condition. The Colorado story is in green, the *Duck Dynasty* story is in pink, and the unemployment benefits story is in brown.

Table A1: Balance checks for randomization, pooled data from Studies 1 and 3 (N=1112).

Note: *p<0.05

	<i>Treatment Group Means</i>			<i>Difference in Means</i>		
	1 (Check-all)	2 (Open-ended)	3 (Yes/no)	1-2	1-3	2-3
Age	31.64 (11.38)	29.58 (9.08)	30.95 (10.51)	2.06 (-0.75)*	0.69 (-0.81)	-1.37 (-0.70)
Race	1.87 (1.75)	1.65 (1.54)	1.74 (1.63)	0.22 (-0.12)	0.13 (-0.12)	-0.09 (-0.11)
Hispanic	0.07 (0.25)	0.09 (0.28)	0.07 (0.26)	-0.02 (-0.02)	0.00 (-0.02)	0.02 (-0.02)
Female	0.38 (0.49)	0.42 (0.49)	0.37 (0.48)	-0.04 (-0.04)	0.01 (-0.04)	0.05 (-0.03)
Income	3.30 (2.45)	3.18 (2.37)	3.21 (2.34)	0.12 (-0.17)	0.09 (-0.18)	-0.03 (-0.17)
Education	4.05 (1.35)	4.05 (1.29)	4.00 (1.34)	0.00 (-0.10)	0.05 (-0.10)	0.05 (-0.09)
Ideology	3.32 (1.50)	3.10 (1.51)	3.05 (5.53)	0.23 (-0.11)*	0.27 (-0.29)	0.05 (-0.29)
Party ID	2.23 (1.32)	2.26 (1.33)	2.38 (1.32)	-0.02 (-0.10)	-0.14 (-0.10)	-0.12 (-0.09)
Knowledge	2.81 (1.04)	2.79 (1.02)	2.76 (1.00)	0.02 (-0.07)	0.05 (-0.08)	0.03 (-0.07)
Attention	2.80 (0.98)	2.84 (1.04)	2.59 (5.39)	-0.03 (-0.07)	0.21 (-0.28)	0.24 (-0.28)
Computer	0.83 (0.38)	0.85 (0.36)	0.80 (0.40)	-0.02 (-0.03)	0.03 (-0.03)	0.05 (-0.03)
Cleared	0.41 (0.49)	0.32 (0.47)	0.34 (0.47)	0.09 (-0.03)*	0.08 (-0.04)*	-0.01 (-0.03)
Purple	0.18 (0.38)	0.20 (0.40)	0.23 (0.42)	-0.02 (-0.03)	-0.05 (-0.03)	-0.03 (-0.03)

Table A2: Quasi-Poisson and negative binomial regressions with robust standard errors.

	<i>Dependent variable:</i>					
	(Quasi-Poisson)			(Negative binomial)		
	Over	Under	Mis	Over	Under	Mis
	(1)	(2)	(3)	(4)	(5)	(6)
Open-Ended	-0.63*** (0.07)	1.41*** (0.37)	-0.03 (0.10)	-0.59*** (0.08)	1.13*** (0.33)	-0.05 (0.09)
Yes/no	0.44*** (0.06)	0.002 (0.15)	0.38*** (0.06)	0.42*** (0.08)	0.12 (0.21)	0.39*** (0.06)
Age	-0.002 (0.003)	0.01 (0.01)	-0.0002 (0.003)	-0.005* (0.003)	0.004 (0.01)	-0.0004 (0.003)
Pacific	0.29 (0.26)	0.84** (0.38)	0.44* (0.23)	0.37 (0.78)	0.26 (0.47)	0.42* (0.23)
Other	0.18 (0.18)	-1.58 (1.40)	-0.23 (0.29)	0.17 (0.20)	-1.86 (1.34)	-0.25 (0.29)
Native	-0.93*** (0.32)	-2.42 (1.98)	-1.48*** (0.45)	-0.86*** (0.16)	-61.94*** (1.62)	-1.49*** (0.44)
Black	0.19** (0.09)	-0.01 (0.33)	0.14 (0.09)	0.19** (0.09)	-0.12 (0.25)	0.12 (0.09)
Asian	-0.08 (0.10)	0.12 (0.27)	-0.03 (0.09)	-0.07 (0.12)	0.09 (0.18)	-0.03 (0.09)
Hispanic	-0.20* (0.12)	2.22 (2.03)	0.33 (0.35)	0.01 (0.14)	1.72 (1.54)	0.36 (0.36)
Female	-0.07 (0.06)	0.08 (0.25)	-0.03 (0.07)	0.03 (0.07)	0.24 (0.22)	-0.01 (0.06)
Income	-0.003 (0.01)	0.07 (0.06)	0.01 (0.01)	-0.02 (0.01)	0.04 (0.05)	0.01 (0.01)
Education	0.03 (0.02)	0.24* (0.12)	0.07** (0.03)	0.02 (0.02)	0.15 (0.11)	0.07*** (0.03)
Ideology	-0.01*** (0.003)	-0.01 (0.02)	-0.01*** (0.004)	-0.01*** (0.004)	-0.005 (0.02)	-0.01** (0.01)
Party ID	0.0002 (0.02)	0.10 (0.15)	0.02 (0.03)	0.004 (0.02)	0.04 (0.15)	0.02 (0.03)
Knowledge	0.13*** (0.03)	0.10 (0.14)	0.12*** (0.04)	0.12*** (0.03)	0.20 (0.14)	0.12*** (0.04)
Attention	-0.01* (0.01)	-0.001 (0.01)	-0.01 (0.01)	-0.22*** (0.03)	0.01 (0.01)	-0.05*** (0.01)
Computer	-0.11* (0.06)	0.39* (0.23)	0.004 (0.07)	-0.16** (0.08)	0.18 (0.30)	-0.002 (0.07)
Cleared	0.28*** (0.06)	-0.37 (0.24)	0.13** (0.06)	0.23*** (0.07)	-0.45* (0.24)	0.13** (0.06)
Purple	0.17*** (0.06)	-0.51 (0.59)	0.02 (0.13)	0.16** (0.07)	-1.91** (0.87)	-0.03 (0.12)
Constant	0.80*** (0.16)	-1.43 (1.17)	0.62** (0.26)	1.55*** (0.22)	-1.95 (1.23)	0.75*** (0.25)

Note:

Appendix A4: Comparison of Samples

While there was substantial attrition in Study 2, it appears that the final sample was not significantly skewed from the full sample (which included respondents who did not install the browser widget). As the table below illustrates, the minor differences in observed characteristics were well within sampling variability. Regardless, “treatment effects” for this study are local to subjects who “complied” by installing the software.

(For reference, category 3 of the income scale corresponds to incomes between \$30,000-39,999, and category 4 of the education scale corresponds to a two-year college degree. Ideology and party identification are the usual 7-point scales.)

	Full Sample	Subsample	Difference
Age	30.34 (10.58)	28.75 (9.56)	1.59 <i>0.16</i>
Income	2.89 (4.98)	3.14 (2.29)	-0.25 <i>0.44</i>
Education	4.0 (1.33)	3.95 (1.29)	0.05 <i>0.75</i>
% Female	38.14 (48.62)	32.94 (47.28)	5.2 <i>0.35</i>
% White	77.86 (41.56)	74.12 (44.06)	3.74 <i>0.46</i>
% Black	6.09 (23.93)	5.88 (23.67)	0.21 <i>0.94</i>
Ideology	2.77 (6.39)	2.59 (1.45)	0.18 <i>0.57</i>
Party ID	2.32 (1.34)	2.27 (1.38)	0.05 <i>0.75</i>
<i>N</i>	568	85	

Table A3: Means and standard deviations (in parentheses) of demographic and political characteristics of respondents in both the full sample and the subsample containing only those who successfully installed and ran the browser widget. The third column lists the differences in means followed by *p*-values (in italics) generated by two-sample *t*-tests.

Appendix A5: Site-Level Determinants of Reported Exposure

The findings from Studies 1-3 illustrate that open-ended questions are substantially better than the other types for the purpose of measuring the overall audience for news and its distribution across different outlets. Here, I explore whether characteristics of the sites themselves might also play a role in respondents' overall tendency to over- or underreport exposure. I focus on two: the popularity of sites as gauged by the number of unique visitors, and the partisan tilt of the sites' readership. If respondents are merely reporting exposure to sources because they are popular—perhaps leading websites to be more accessible in memory—this could lead to overreporting, all else equal. Furthermore, the partisan characteristics associated with particular sites may have an effect on individuals' tendencies to report visiting them. Are people less willing to report exposure to Democratic- or Republican-leaning sites, or perhaps sites with any partisan leaning at all?

To answer these questions, I gathered additional data on total unique monthly visitors per site from comScore, averaged from October to December 2013. For the data on partisan leanings, I relied on comScore's Plan Metrix service, which provides 12-month aggregated data from a running panel survey with approximately 12,000 U.S.-based participants.¹¹ In particular, I used the service's "Composition Index" for self-identified Democrats and Republicans, which captures the degree to which visitors to a given site from either group are over- or underrepresented as compared to the total Internet sample. The Republican Composition Index ranges from 29.3% to 207.3%, and the Democratic index ranges from 43.7% to 174.2%. I simply took the Republican share of the two indexes and divided by 100. To gauge partisan leaning, I computed the standard deviation of each site's Republican and Democratic indexes—a somewhat crude measure intended to capture the variation

¹¹These panelists are themselves a subset of the larger Media Metrix panel. All respondents are age 18 or older.

in how the two groups are over- or underrepresented in the audience for a given site.

I use total reported exposure for each site from Studies 1 and 3 as the dependent variable, pooled across treatment conditions. In all models, I included an indicator for whether a particular site was one of the choices in the check-all and yes/no questions, since, as the above graphs demonstrated, those question formats have a tendency to generate overreporting. In all, there was sufficient data for 116 of the available sites.

Table A4: OLS and quasi-Poisson regressions with Huber-White robust standard errors.

	<i>DV: Total reported visits per site</i>			
		OLS		Quasi-Poisson
	(1)	(2)	(3)	(4)
Total actual exposure		1.30*** (0.18)	1.10*** (0.12)	0.004*** (0.0004)
Unique visitors/1m	2.49** (1.19)	2.60** (1.05)	1.39 (0.92)	0.01*** (0.004)
Partisan leaning			0.06 (0.14)	-0.03*** (0.01)
Share R composition	-17.21 (34.35)	2.27 (14.91)	-1.54 (23.39)	-0.41 (2.37)
Included in check-all/ forced-choice Q's?	148.18*** (33.09)		84.91*** (16.19)	3.97*** (0.25)
Constant	1.46 (17.05)	-11.34 (7.61)	-16.74 (14.47)	0.99 (1.24)
<i>N</i>	116	116	116	116
<i>R</i> ²	0.57	0.80	0.87	
Adjusted <i>R</i> ²	0.56	0.79	0.86	

Note:

p*<0.1; *p*<0.05; ****p*<0.01

I report several different OLS models with robust standard errors in addition to a full quasi-Poisson model that takes into account the overdispersion in the dependent variable. As Table A4 shows, inclusion in one of the closed-ended survey questions is strongly associated with reporting exposure to a particular website. There are two possible reasons for this finding. First, the indicator is likely capturing some of the effects of two of the treatment conditions on overreporting in general. And second, by construction, those question types include selections of well-known sites that respondents are likely to make. This suggests potential multicollinearity in two of the independent variables, total unique visitors and the question indicator, and indeed they are strongly correlated ($r = .527$). This is one possible reason why overall site traffic drops out of significance in the third model. Still, whether as a result of inclusion in survey questions or more directly, the association between a site's total audience and reported exposure is clear in Model 2.

The Republican share of the composition index has no apparent effect on reporting exposure. A partisan tilt in either direction, on the other hand, predicts less reporting overall in the quasi-Poisson model, although this finding does not hold in the OLS models. And finally, even holding the other factors constant, total *actual* exposure as measured by the *Link Classification Task* predicts total reported exposure, demonstrating the validity of the measure.

Appendix B1

The treatment-by-treatment analysis of the heterogeneous effects of the tweet encouragement relied on an assumption that all subjects who signed the petition in the organizer treatment would have done so in the followers treatment as well. This assumption is not guaranteed to hold if there are some subjects who would only respond to one treatment or the other. Table B1 describes the eight theoretically possible types of subjects. The first type, for example, would sign the petition regardless of treatment condition. The second type, however, would only sign the petition if assigned to the public tweet or the organizer direct message treatments—but not if assigned to the follower condition.

Table B1: Possible Subject Types

Type	Public Tweet Only	Organizer	Follower	Population Proportion
1	1	1	1	$\pi_1 = 0$
2	1	1	0	$\pi_2 = 0$
3	1	0	1	$\pi_3 = 0$
4	1	0	0	$\pi_4 = 0$
5	0	1	1	$\pi_5 = ?$
6	0	1	0	$\pi_6 = ?$
7	0	0	1	$\pi_7 = ?$
8	0	0	0	$\pi_8 = [0.955, 0.965]$

We know that the proportions of types 1 through 4 in the population are all equal to zero: no subjects in the public tweet conditions signed the petitions. Together, types 5 through 7 account for approximately 3.6% of the population in Study 1 and approximately 4.5% of the population in Study 2; type 8 accounts for the remainder.

The crucial question for us is the proportion of types 6 and 7, π_6 and π_7 . If they are both equal to zero, then we induce no bias when we condition on DM type in the second-stage experiment. If, however, there are 6's or 7's that sign the petition, then conditioning would in fact induce bias. What evidence do we have that the

proportion of 6's and 7's are both equal to zero?

First, we know that equal proportions of subjects signed the petitions in the organizer and follower DM treatments. In Chapter 2, we describe well-estimated average differences between the two conditions to be very close to zero (and certainly not statistically significantly different from zero). We can therefore infer that $\pi_6 = \pi_7$:

$$E[Y|Z = \text{Organizer}] = \pi_5 + \pi_6$$

$$E[Y|Z = \text{Follower}] = \pi_5 + \pi_7$$

$$E[Y|Z = \text{Organizer}] = E[Y|Z = \text{Follower}]$$

$$\pi_5 + \pi_6 = \pi_5 + \pi_7$$

$$\pi_6 = \pi_7$$

If π_6 and π_7 did not equal zero, then they would have to exactly counterbalance one another, which is possible, but unlikely. It would be especially unlikely for $\pi_6 = \pi_7 = c > 0$ across a wide variety of subjects. A heterogeneous effects analysis of the “organizer” versus “follower” manipulation by network centrality suggested no difference in treatment effects at any level of centrality. This does not constitute conclusive proof that the only types in the population are 5's and 8's, but it is suggestive. The analyses in Chapter 2 rely on this assumption and should be weighed with the plausibility of this assumption in mind.

Appendix B2: Randomization Checks

In this section, we present randomization checks for Studies 1 and 2. In particular, under random assignment of the treatment, we would expect the pre-treatment covariates to be balanced across the three treatment conditions. Equivalently, we would expect that the covariates would not predict treatment status. For each experiment, we will present three randomization checks:

1. Balance tables, presented in Tables B2 and B3. The tables present means and standard errors for four pre-treatment covariates: Account Type (male, female, organization, unknown), Number of Followers, Days on Twitter, and Eigenvector Centrality.
2. Tests of independence for each covariate, shown in the last columns of Tables B2 and B3.
 - Study 1 was carried out using complete random assignment, so we can directly apply the chi-square test to the categorical variable (Account Type) and the f -test of joint independence to the continuous variables (Number of Followers, Days on Twitter, and Eigenvector Centrality).
 - Study 2 was carried out using block random assignment, so we condition the test on the experimental block, and aggregate the tests to form a single p -value using Fisher's method (Fisher 1925, Section 21.1). Additionally, we use Fisher's exact test in lieu of the chi-square test because of the low cell count within a single stratum. The required assumption that the margins are fixed is met by design (a fixed number of treatments are allocated to a fixed distribution of account types).
3. Omnibus test of joint independence of all the covariates from the treatment assignment, presented in the last rows of Tables B2 and B3. This is conducted

using a randomization inference procedure:

- We obtain the likelihood ratio statistic from a multinomial logistic regression of treatment assignment on the covariates.
- We permute the random assignment 1,000 times according to the original random assignment protocol.
- We obtain the likelihood ratio statistics from regressions of these 1,000 simulated treatment assignments on the covariates.
- We construct a p -value by observing the frequency with which the simulated statistics exceed the observed statistic.

Table B2: Experiment 1 Balance

	Treatment Assignment			p -value
	Public Tweet	Follower	Organizer	
Account Type: Female	0.309 (0.008)	0.287 (0.012)	0.307 (0.012)	
Account Type: Male	0.381 (0.008)	0.375 (0.013)	0.401 (0.013)	
Account Type: Organization	0.245 (0.007)	0.254 (0.011)	0.230 (0.011)	
Account Type: Unknown	0.065 (0.004)	0.083 (0.007)	0.061 (0.006)	0.072
Number of Followers	596.240 (14.146)	616.603 (22.099)	635.733 (23.357)	0.312
Days on Twitter	1631.438 (8.983)	1637.362 (14.189)	1637.179 (14.153)	0.910
Eigenvector Centrality	0.039 (0.001)	0.038 (0.002)	0.038 (0.002)	0.960
N	3687	1500	1500	
Omnibus p -value: 0.607				

Table B3: Experiment 2 Balance

	Treatment Assignment			<i>p</i> -value
	Public Tweet	Follower	Organizer	
Account Type: Female	0.315 (0.008)	0.329 (0.009)	0.348 (0.010)	
Account Type: Male	0.405 (0.008)	0.392 (0.010)	0.379 (0.010)	
Account Type: Organization	0.224 (0.007)	0.222 (0.008)	0.214 (0.008)	
Account Type: Unknown	0.056 (0.004)	0.056 (0.005)	0.059 (0.005)	0.425
Number of Followers	585.599 (14.469)	580.738 (16.999)	581.281 (16.996)	0.535
Number of Tweets	1559.503 (10.067)	1554.428 (12.131)	1552.586 (11.988)	0.601
Eigenvector Centrality	0.032 (0.001)	0.031 (0.001)	0.030 (0.001)	0.537
N	3495	2498	2514	
Omnibus <i>p</i> -value: 0.113				

Appendix B3: Heterogeneous Effects of Treatment

Notes for both tables: Eigenvector centrality, Number of Followers, and Days on Twitter in standard units and centered at zero. Robust standard errors in parentheses.

* $p < .1$; ** $p < .05$; *** $p < .01$

Study 1

	signed		tweeted	
Treatment: Follower	0.039*** (0.005)	0.039*** (0.005)	0.027*** (0.004)	0.027*** (0.004)
Treatment: Organizer	0.033*** (0.005)	0.033*** (0.005)	0.016*** (0.003)	0.016*** (0.003)
Eigenvector Centrality	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Follower X Centrality	-0.006* (0.003)	-0.003 (0.002)	-0.003 (0.002)	-0.003 (0.002)
Organizer X Centrality	0.006 (0.009)	-0.002 (0.001)	-0.002* (0.001)	-0.002 (0.001)
Number of Followers	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Follower X Followers	-0.005* (0.003)	-0.005* (0.003)	-0.004 (0.003)	-0.004 (0.003)
Organizer X Followers	-0.002 (0.004)	-0.002 (0.004)	-0.003 (0.002)	-0.003 (0.002)
Days on Twitter	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Follower X Days on Twitter	-0.002 (0.006)	-0.002 (0.006)	-0.002 (0.006)	-0.004 (0.005)
Organizer X Days on Twitter	0.001 (0.005)	0.001 (0.005)	0.000 (0.005)	0.000 (0.003)
Account Type: Male	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Account Type: Organization	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Account Type: Unknown	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Follower X Male	-0.005 (0.014)	-0.005 (0.014)	-0.005 (0.014)	-0.001 (0.010)
Organizer X Male	-0.011 (0.013)	-0.011 (0.013)	-0.011 (0.013)	0.005 (0.008)
Follower X Organization	-0.043*** (0.012)	-0.043*** (0.012)	-0.043*** (0.012)	-0.007 (0.011)
Organizer X Organization	-0.036*** (0.011)	-0.036*** (0.011)	-0.036*** (0.011)	-0.006 (0.008)
Follower X Unknown	-0.013 (0.021)	-0.013 (0.021)	-0.013 (0.021)	0.020 (0.021)
Organizer X Unknown	-0.026 (0.018)	-0.026 (0.018)	-0.026 (0.018)	0.007 (0.016)
Constant	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
N	6,687	6,687	6,687	6,687

Study 2

	signed			tweeted		
Treatment: Follower	0.046*** (0.004)	0.046*** (0.004)	0.044*** (0.007)	0.014*** (0.002)	0.014*** (0.002)	0.010*** (0.003)
Treatment: Organizer	0.042*** (0.004)	0.043*** (0.004)	0.053*** (0.008)	0.011*** (0.002)	0.011*** (0.002)	0.010*** (0.003)
Eigenvector Centrality	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)		
Follower X Centrality	-0.002 (0.003)	0.001 (0.002)				
Organizer X Centrality	-0.006** (0.003)	-0.001* (0.001)				
Number of Followers	0.000 (0.000)	0.000 (0.000)				
Follower X Followers	-0.001 (0.004)	0.003 (0.003)				
Organizer X Followers	-0.004 (0.003)	0.003 (0.003)				
Days on Twitter		0.000 (0.000)			0.000 (0.000)	
Follower X Days on Twitter		-0.002 (0.005)			-0.001 (0.003)	
Organizer X Days on Twitter		0.000 (0.004)			-0.002 (0.002)	
Account Type: Male			0.000 (0.000)			0.000 (0.000)
Account Type: Organization			0.000 (0.000)			0.000 (0.000)
Account Type: Unknown			0.000 (0.000)			0.000 (0.000)
Follower X Male			0.019* (0.011)			0.010* (0.006)
Organizer X Male			-0.001 (0.010)			0.007 (0.005)
Follower X Organization			-0.026*** (0.009)			0.001 (0.006)
Organizer X Organization			-0.038*** (0.009)			-0.007 (0.004)
Follower X Unknown			0.006 (0.020)			0.012 (0.013)
Organizer X Unknown			-0.026* (0.015)			-0.004 (0.008)
Constant	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
N	8,507	8,507	8,507	8,507	8,507	8,507

Appendix B4: Experimental Materials

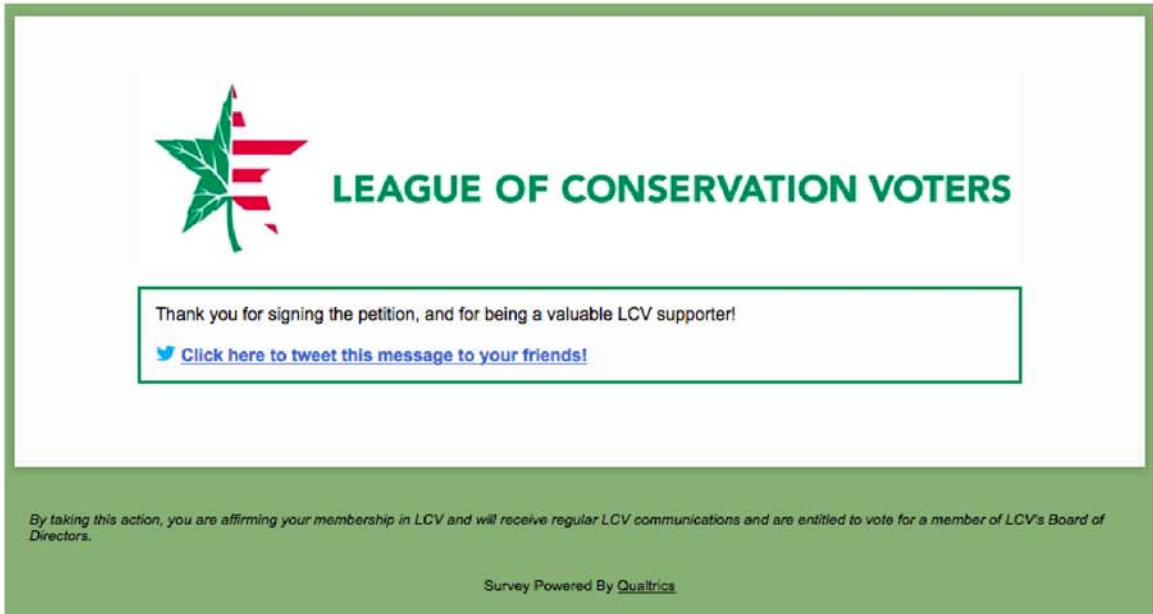


Figure B1: A screenshot of the tweet encouragement randomly shown to respondents in either of the DM conditions who completed the online petition.



Figure B2: A screenshot of the pop-up window shown to respondents who clicked the tweet button.

TELL CONGRESS: END BIG OIL HANDOUTS

It's ridiculous that while Americans struggle to pay rising gas prices and support their families, oil companies are making obscene profits and still getting billions of our taxpayer dollars in subsidies every year.

If you're tired of paying twice for your gas — first at the pump and then again on tax day — tell Congress to end Big Oil handouts. *Fill out your information below to send a message now.*

PETITION: End Big Oil Handouts

Dear Member of Congress,

Over the past decade, the big five oil companies — BP, Chevron, ConocoPhillips, ExxonMobil, and Shell — have enjoyed more than \$1 trillion in profits. Last year alone, these oil giants made \$118 billion in profits.

At the same time, Big Oil continues to benefit from billions in taxpayer-funded government handouts.

Enough is enough.

Sincerely,

Personalize your message:

I urge Congress to end Big Oil handouts.

Figure B3: The top half of the online petition whose link was sent to subjects in the DM conditions in Study 1.

SUPPORT THE EPA'S BIGGEST ACTION EVER ON CLIMATE CHANGE!

The Obama administration has proposed a critical plan to limit carbon pollution from new and existing power plants that fuels global warming. Currently, we limit how much mercury, arsenic, soot and other pollutants power plants dump into our air. But if the Koch brothers, the coal industry, and congressional climate change deniers have their way, polluters will continue to have free rein to dump unlimited amounts of harmful carbon pollution into our air.

The Environmental Protection Agency (EPA) has been working on commonsense carbon limits for power plants, but lobbyists, corporate polluters and Washington insiders are working tirelessly to stop this progress in its tracks. By signing a petition in support of the EPA's efforts, however, you can give the EPA the support it needs to stand up to the polluters.

Sign this petition to the EPA today, and your voice will be heard. Don't miss the chance to help make a huge impact on the fight for our planet's future -- *fill out your information here to support this critical climate change action now.*

PETITION: Cut Carbon Pollution From Power Plants

Dear EPA Administrator McCarthy,

I strongly support the EPA's effort to limit carbon pollution from existing power plants.

Personalize your message:

Figure B4: The top half of the online petition whose link was sent to subjects in the DM conditions in Study 2.



Figure B5: The public tweet from Study 1.

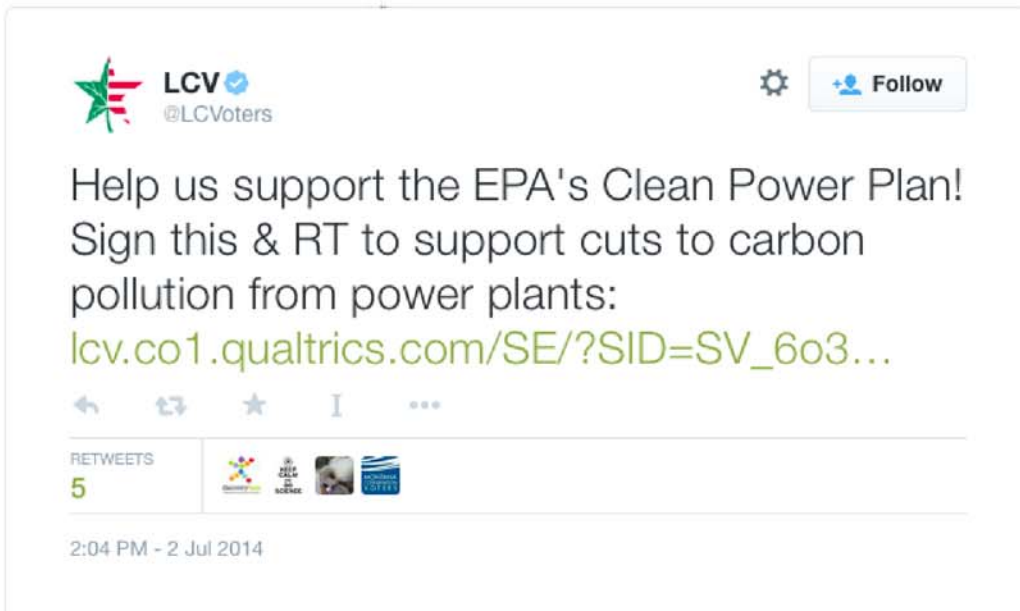


Figure B6: The public tweet from Study 2.

Appendix B5: Privacy and Ethical Considerations

Our research design presents ethical challenges common to field experiments implemented in online environments. In particular, like much unobtrusive field research, we could not obtain informed consent from subjects without compromising our inferential strategy. In proceeding with these studies, we relied on our own judgment that the benefits of the study outweighed any risks to subjects.

Furthermore, since most Twitter activity is public by design, we took a series of steps to protect subjects' anonymity. To ensure that our approach toward consent and privacy met common standards for minimizing any potential harm, we obtained IRB approval from one of the authors' home institutions [details withheld]. Below, we detail several considerations that we believe are crucial to evaluating the ethics of the experiments reported here (as well as others with similar designs).

Twitter's Policies

Twitter's privacy policy, available at <https://twitter.com/privacy>, explicitly informs users that their public profile information and tweets are made immediately available to third parties, including research institutions:

For instance, your public user profile information and public Tweets are immediately delivered via SMS and our APIs to our partners and other third parties, including search engines, developers, and publishers that integrate Twitter content into their services, and institutions such as universities and public health agencies that analyze the information for trends and insights. When you share information or content like photos, videos, and links via the Services, you should think carefully about what you are making public.

This policy, part of the terms of service for all users, ensures that collecting public tweet data is firmly within the bounds of reasonable use.

Data Privacy

LCV's experience mobilizing its members while protecting their privacy generally assuaged our concerns. We acknowledge, however, that in collecting data for this study we make public information somewhat more accessible. No individual tweets are revealed in the study, and personally identifiable information such as user names, descriptions, network connections, and location have been removed from all replication files.

Organization's Goals

A final concern regards the nature of the manipulation. The messages used in both studies were approved by LCV as part of ongoing social media campaigns directly related to its core goals. As shown in Figure 2.2, LCV posts approximately 6 tweets or retweets per weekday on average; the public tweet component of the experiments' design was designed to fit in with the organization's existing day-to-day engagement strategy.

Private direct messages (DM) are less commonly used by organizations, but practically speaking these are no more intrusive than a mass email message. In this case, by signing up for Twitter and voluntarily following LCV's account, subjects assigned to receive a DM in effect opted to receive communications from the latter via the former.

However, we do not take these concerns lightly. Despite the fact that both studies' messages were part of a preexisting social media campaign, we acknowledge that the DM treatments comprised an unorthodox communications strategy. The petitions may also have taken several minutes of subjects' time. In response, we note that LCV's follower count has continued to rise and that we find no evidence of a backlash effect of any kind.

Appendix C1: List of Sites

abcnews.go.com
america.aljazeera.com
digbysblog.blogspot.com
dish.andrewsullivan.com
krugman.blogs.nytimes.com
latino.foxnews.com
nbcpolitics.nbcnews.com
news.google.com
news.msn.com
news.yahoo.com
newswatch.nationalgeographic.com
about.com
aljazeera.com
americanthinker.com
anncoulter.com
antiwar.com
aol.com
ap.org
bbc.com
beforeitsnews.com
billoreilly.com
blackamericaweb.com
bloomberg.com
breitbart.com
businessinsider.com
buzzfeed.com
buzzya.com
c-span.org
cbsnews.com
chacha.com
chicagosuntimes.com
chicagotribune.com
cnn.com
comondreams.org
counterpunch.org
crookedtimber.org
crooksandliars.com
csmonitor.com

dailycaller.com
dailydot.com
dailykos.com
dailymail.co.uk
democratichub.com
democraticunderground.com
denverpost.com
dickmorris.com
drudgereport.com
economist.com
eschatonblog.com
examiner.com
factcheck.org
firsttoknow.com
fivethirtyeight.com
forbes.com
foreignaffairs.com
foreignpolicy.com
foxnews.com
foxnews.com
foxnewsinsider.com
freebeacon.com
freepatriot.org
freerepublic.com
frontpagemag.com
google.com
gopusa.com
govexec.com
govtrack.us
hotair.com
huffingtonpost.com
huffingtonpost.com
humanevents.com
instapundit.com
inthesetimes.com
latimes.com
littlegreenfootballs.com
mediaite.com
michellemalkin.com

mikehuckabee.com
motherjones.com
mrc.org
mrconservative.com
mtvu.com
nationalgeographic.com
nationaljournal.com
nationalmemo.com
nationalreview.com
nbcnews.com
newrepublic.com
newsbusters.org
newser.com
newsmax.com
newyorker.com
npr.org
nydailynews.com
nypost.com
nytimes.com
obamacarefacts.com
outsidethebeltway.com
pbs.org
politicalwire.com
politico.com
politicususa.com
politifact.com
powerlineblog.com
reagancoalition.com
realclearpolitics.com
realclearworld.com
reason.com
reddit.com
redstate.com
rightwingnews.com
rushlimbaugh.com
salon.com
samuel-warde.com
slate.com
smithsonianmag.com

sourcewatch.org
stormfront.org
takepart.com
talkingpointsmemo.com
the-american-interest.com
theatlantic.com
theatlanticwire.com
theblaze.com
thedailybeast.com
thegrio.com
theguardian.com
thehill.com
themonkeycage.org
thenation.com
theweek.com
thinkprogress.org
townhall.com
trueactivist.com
truthout.org
twitchy.com
upi.com
usatoday.com
usnews.com
voanews.com
volokh.com
washingtonexaminer.com
washingtonmonthly.com
washingtonpost.com
washingtontimes.com
weeklystandard.com
westernjournalism.com
wonkette.com
worldnetdaily.com
wsj.com
marketwatch.com
vox.com

en.wikipedia.org/wiki/For-profit_education
harkin.senate.gov/help/forprofitcolleges.cfm
marketwatch.com/story/which-for-profit-education-companies-are-good-bets-2014-09-17
huffingtonpost.com/2014/09/16/cfpb-corinthian-colleges_n_5832540.html
huffingtonpost.com/2014/09/17/education-management-corinthian_n_5832648.html
money.cnn.com/2014/09/07/investing/for-profit-education-stocks/
cnn.com/2014/08/28/opinion/barkley-corinthian-for-profit-colleges/
thestreet.com/story/12777157/1/why-for-profit-education-institutions-should-be-sent-to-detention.html
thinkprogress.org/education/2014/08/19/3472835/for-profit-college-job-interview/
theblaze.com/contributions/obama-wants-to-wreck-for-profit-education-and-the-mainstream-media-want-to-help/
mrc.org/bias-numbers/major-papers-15-1-against-profit-education
ncsl.org/research/education/for-profit-colleges-and-universities.aspx
blogs.browardpalmbeach.com/pulp/2014/09/09/anthem_education.php
thehill.com/blogs/congress-blog/education/218374-real-cost-of-obamas-war-against-for-profit-colleges
forbes.com/sites/robertfarrington/2014/09/10/be-selective-in-choosing-a-for-profit-college/
huffingtonpost.com/julia-mezaros/for-profit-colleges-maint_b_5788466.html
therepublic.com/view/story/674ec780a3d84be0b374f7904737f6a0/PA--Shareholders-Education-Management
usnews.com/news/blogs/data-mine/2014/09/22/is-the-college-admissions-bubble-about-to-burst
wcvb.com/money/forprofit-education-stocks-on-fire/27922708
news.google.com/news?nc1=d6IIF0QUBIG_IvMfB8b-MsTEei9nM&q=for-profit+education&lr=English&hl=en&sa=X&ei=iRgiV0bxEOGqyASin4D4Cg&ved=0CCEQqgIwAA
americanbanker.com/issues/179_179/cfpb-hits-for-profit-education-chain-with-lawsuit-1069960-1.html
topics.bloomberg.com/bloomberg-investigation:-for--profit-colleges/
huffingtonpost.com/2014/09/08/john-oliver-student-debt_n_5784266.html
nacacnet.org/issues-action/LegislativeNews/Pages/For-Profit-Colleges.aspx

web.stanford.edu/group/ncpi/documents/pdfs/forprofitandcc.pdf
moneycrashers.com/for-profit-online-colleges/
oregonlive.com/education/index.ssf/2014/09/anthem_shutters_campuses_in_be.html
esquire.com/blogs/news/the-bleak-picture-of-american-universities-091214
online.wsj.com/news/articles/SB10001424052970203937004578076942611172654
huffingtonpost.com/news/for-profit-colleges/
washingtonpost.com/news/storyline/wp/2014/09/17/how-one-student-got-burned-
by-a-for-profit-college-and-bailed-out-by-occupy-wall-street/
en.wikipedia.org/wiki/List_of_for-profit_universities_and_colleges
huffingtonpost.com/william-g-tierney/understanding-the-controv_b_1764275.html
bloombergview.com/articles/2012-10-25/the-long-and-controversial-history-of-for-profit-colleges
bloomberg.com/news/2014-01-06/clinton-pitches-kkr-backed-college-chain-amid-controversy.html
washingtontimes.com/news/2012/feb/2/obama-romney-divided-on-for-profit-colleges/
as.wvu.edu/asreview/for-profit-colleges-raise-controversy/
nytimes.com/2010/10/25/us/25iht-eduSide25.html?_r=0
insidehighered.com/news/2013/10/10/laureates-growing-global-network-institutions
becker-posner-blog.com/2010/06/the-controversy-over-forprofit-collegesposner.html
motherjones.com/politics/2014/09/for-profit-university-subprime-student-poor-minority
huffingtonpost.com/tag/for-profit-universities/
finance.yahoo.com/news/government-investigates-Corinthian-college-financial-aid-students-react-170452053.html
collegexpress.com/lists/list/popular-for-profit-universities/264/
evollution.com/distance_online_learning/dispelling-some-myths-about-proprietary-education/
nea.org/assets/docs/HE/vol10no4.pdf
usconservatives.about.com/od/education/a/The-Liberal-Attack-On-For-Profit-Education.htm
franklin.edu/blog/non-profit-vs-for-profit-colleges-what-you-need-to-know/
onlinecollegereport.com/for-profit-colleges-vs-not-for-profit-colleges/

degreesinder.com/information-online-degree-programs/faq/what-is-a-for-profit-college/
huffingtonpost.com/2012/06/04/for-profit-colleges-student-debt-dropout_n_1567607.html
huffingtonpost.com/2013/12/16/corinthian-colleges-job-placement_n_4433800.html
bloomberg.com/news/articles/2015-04-23/laureate-said-planning-1-billion-ipo-of-for-
profit-universities
forbes.com/sites/debtwire/2015/04/22/failed-promises-pricey-loans-keep-for-profit-
education-in-regulatory-crosshairs/
businessinsider.com/laureate-education-planning-billion-dollar-ipo-2015-4
theatlantic.com/education/archive/2015/02/the-downfall-of-for-profit-colleges/385810/
help.senate.gov/imo/media/for_profit_report/Contents.pdf
opensecrets.org/industries/indus.php?ind=H5300
npr.org/blogs/ed/2014/12/24/370979991/an-update-on-for-profit-colleges
forbes.com/sites/jamesmarshallcrotty/2014/10/31/is-rapid-decline-of-for-profit-colleges-
good-for-education/
insidehighered.com/news/focus/profit-higher-ed
huffingtonpost.com/news/for-profit-education/
seekingalpha.com/article/3032236-apollo-warns-get-out-of-for-profit-education-now
thestreet.com/story/13037791/1/strayer-for-profit-education-companies
-suffer-through-supply-and-demand-lessons.html
thestreet.com/story/13005580/1/how-obamas-free-community-college-program-may-kill-for-profit-
education-companies.html
topics.nytimes.com/top/reference/timestopics/subjects/f/forprofit_schools/index.html
usnews.com/education/best-colleges/paying-for-college/articles/2014/10/01/3-facts-for-students-to-
know-about-for-profit-colleges-and-student-debt

huffingtonpost.com/davidhalperin/abuses-at-corinthian-are_b_7118546.html
consumerist.com/2015/04/23/consumer-groups-ask-congress-to-ensure-that-for-profit-schools-are-held-accountable/
usnews.com/news/articles/2015/02/24/students-are-returning-to-for-profit-colleges
npr.org/2015/03/31/396636885/students-from-troubled-for-profit-colleges-refuse-to-pay-back-loans
time.com/money/3573216/veterans-college-for-profit/
citizensforethics.org/page/-/PDFs/Legal/10_07_10%20Complaint%20U.S.%20Dept%20Education.pdf
sloperama.com/advice/lesson44.html
as.wiki.studyroom.us/For-profit+education
aetherczar.com/?p=924
vox.com/2014/9/30/6862939/online-for-profit-college-degree-job-market-interview
vox.com/2014/10/30/7130151/gainful-employment-for-profit-colleges-obama-administration
younginvincibles.org/tag/for-profit-colleges/
degreelibrary.org/for-profit-vs-public-college/
jobunlocker.com/blog/the-pros-and-cons-of-for-profit-colleges/
washingtonpost.com/posteverything/wp/2014/07/23/as-a-teacher-i-fought-predatory-for-profit-schools-from-the-inside/
adulthoodeducationadvocates.com/the_pros_and_cons_of_for_profit_colleges
uscollegesearch.org/blog/college-search-2/comparing-types-of-colleges-the-pros-and-cons-of-for-profit-schools
collegeconfidential.com/for-profit-colleges-pros-and-cons/
niyc-alb.org/50th_Anniversary/BewareofFor-ProfitColleges.pdf
insidehighered.com/news/2011/11/21/report-higher-education-research-groups-annual-meeting
healthpronet.org/docs/2011-10-HPN-Jones.pdf
ibtimes.com/john-oliver-slams-profit-colleges-student-debt-cycle-it-will-follow-you-forever-1685864
huffingtonpost.com/davidhalperin/bill-maher-john-oliver-sl_b_5818842.html
youtube.com/watch?v=P8pjd1QEA0c

slate.com/blogs/browbeat/2014/09/08/john_oliver_student_debt_segments_goes_after_for_profit_colleges_video.html
qz.com/261496/watch-john-oliver-make-the-case-that-for-profit-colleges-are-a-waste-of-time-and-money/
consumerist.com/2014/09/08/john-oliver-on-for-profit-colleges-you-might-as-well-go-to-hogwarts/
theweek.com/speedreads/index/267687/speedreads-john-oliver-unpacks-the-special-horrors-
of-student-debt-and-for-profit-universities
laughspin.com/2014/09/08/john-oliver-says-student-debt-is-like-std-
last-week-tonight-video-slams-for-profit-colleges/
gawker.com/john-oliver-calls-bullshit-on-debt-fueled-for-profit-co-1631900831
studentdebtcrisis.org/student-debt-crisis%EF%BB%BF-wins-real-time-with-bill-maher
%EF%BB%BFs-flipadistrict-competition/
ny.thedailydigest.org/2014/09/15/bill-maher-john-oliver-slam-predatory-for-profit-colleges/
hbo.com/real-time-with-bill-maher/episodes/12/329-episode/synopsis/quotes.html#/
salon.com/2014/09/22/higher_eds_obscene_profit_scheme_schools_banks_and_the_government_win_you_drown_in_debt/
twitter.com/RalphNader/status/511533577381875712
c-span.org/video/?307608-5/forprofit-colleges
bloomberg.com/news/2012-09-28/for-profit-colleges-beat-peers-on-loan-breaks-corporate-brazil.html
washingtonpost.com/opinions/for-profit-colleges-shouldnt-blame-the-education-department
-for-their-poor-marks/2014/08/26/7e441fea-2c99-11e4-be9e-60cc44c01e7f_story.html
online.wsj.com/news/articles/SB100014240527023035532045793472307730666584
blogs.reuters.com/felix-salmon/2012/01/16/for-profits-vs-not-for-profits/
aacu.org/liberaleducation/2009/summer/nussbaum
wsj.com/articles/steve-gunderson-making-profit-a-dirty-word-in-higher-education-1415837070
townhall.com/tipsheet/mattvespa/2015/03/12/education-is-not-for-profit-says-
teachers-union-president-making-500k-a-year-n1969789
consumerist.com/2015/04/23/senators-want-to-close-federal-funding-loophole-
exploited-by-for-profit-colleges/

Appendix C2: Measuring Sites' Political Slant

Table C1: List of online news sources and the corresponding partisan slant index, computed by dividing the percentage of the site's audience identifying as Republican in comScore's Plan Metrix panel by the percentage identifying with either party.

Site	% Rep (out of R+D)	Site	% Rep (out of R+D)
shortlist.com	0.230	nytimes.com	0.449
scotsman.com	0.235	yahoo news	0.451
smithsonianmag.com	0.257	dailymail.co.uk	0.453
lifeneews.com	0.307	huffingtonpost.com	0.453
abcactionnews.com	0.337	dallasnews.com	0.454
lifesitenews.com	0.344	therightscoop.com	0.455
dw.de	0.362	nydailynews.com	0.455
mediaite.com	0.373	upi.com	0.456
talkingpointsmemo.com	0.374	nypost.com	0.456
baynews9.com	0.375	cbsnews.com	0.456
rawstory.com	0.378	abovetopsecret.com	0.459
philly.com	0.379	wsj.com	0.460
knoxnews.com	0.379	newsnet5.com	0.461
metro.co.uk	0.395	nbcnews.com	0.461
bostonglobe.com	0.395	ifyouonlynews.com	0.465
chron.com	0.400	mirror.co.uk	0.466
altnet.org	0.402	salon.com	0.467
cnsnews.com	0.403	express.co.uk	0.468
thesun.co.uk	0.403	wn.com	0.468
mentalfloss.com	0.403	aljazeera.com	0.470
usnews.com	0.407	washingtonpost.com	0.471
worldtruth.tv	0.407	huffingtonpost.ca	0.474
newsok.com	0.408	latimes.com	0.476
theatlantic.com	0.409	vox.com	0.476
addictinginfo.org	0.411	usatoday.com	0.479
buzzfeed.com	0.412	beforeitsnews.com	0.480
motherjones.com	0.413	daily kos	0.483
slate.com	0.414	bloomberg.com - politics	0.483
sfgate.com	0.414	newsnow.co.uk	0.484
9news.com	0.419	detroitnews.com	0.491
theguardian.com	0.423	theblaze.com	0.496
takepart.com	0.424	abc7news.com	0.499
news.google.com	0.424	freep.com	0.503
theweek.com	0.425	medium.com	0.503
chicagotribune.com	0.426	breitbart.com	0.507
npr.org	0.428	ynetnews.com	0.510
startribune.com	0.429	mercurynews.com	0.511
examiner.com	0.430	ijreview.com	0.516
cbc.ca	0.430	newsmax.com	0.519
time.com	0.432	foxnews.com	0.521
telegraph.co.uk	0.432	nymag.com	0.523
newsobserver.com	0.433	foreignpolicy.com	0.525
csmonitor.com	0.436	nzherald.co.nz	0.526
cnn.com	0.436	newsbusters.org	0.529
buffalonews.com	0.437	dailystar.co.uk	0.541
hawaiinewsnow.com	0.437	nationalreview.com	0.545
thedailybeast.com	0.442	mynews13.com	0.556
nj.com	0.443	politico.com	0.556
abc news	0.444	newsiosity.com	0.566
iflscience.com	0.444	townhall.com	0.571
vice.com	0.444	redflagnews.com	0.574
bbc	0.445	rightwingnews.com	0.577
msn news	0.447	daytondailynews.com	0.599
newyorker.com	0.447	news-leader.com	0.609
today.com	0.449	drudgereport.com	0.624
newsday.com	0.449	kingworldnews.com	0.718

Table C2: The ideological slant index is computed by dividing the % of the site's audience identifying as "very" or "somewhat" conservative by the % placing themselves anywhere.

Site	% Con	Site	%Con
scotsman.com	0.095	alternet.org	0.291
duluthnewstribune.com	0.103	universetoday.com	0.293
montgomerynews.com	0.119	mentalfloss.com	0.293
msnewsnow.com	0.162	bostonglobe.com	0.293
national geographic	0.162	heraldsun.com.au	0.295
dailyrecord.co.uk	0.168	chron.com	0.299
buzzle.com	0.181	iflscience.com	0.3
starnewsonline.com	0.189	time.com	0.303
polar.com	0.193	dailydot.com	0.303
mansfieldnewsjournal.com	0.193	jamaica-gleaner.com	0.305
guardianlv.com	0.195	newsobserver.com	0.305
smithsonianmag.com	0.2	topix.com	0.305
thenational.ae	0.203	news-journalonline.com	0.306
newstimes.com	0.208	mediaite.com	0.307
newstatesman.com	0.212	whydonyoutrythis.com	0.307
wsws.org	0.213	news-leader.com	0.308
newsnow.co.uk	0.214	takepart.com	0.309
japantimes.com.jp	0.217	onenewsnow.com	0.309
pasadenastarnews.com	0.219	yournewswire.com	0.311
manchestereveningnews.co.uk	0.221	theatlantic.com	0.311
bringmethenews.com	0.221	usnews.com	0.313
irishtimes.com	0.223	huffingtonpost.ca	0.315
shortlist.com	0.227	miamiherald.com	0.316
expressnews.com	0.229	wvmetronews.com	0.316
theaustralian.com.au	0.23	mirror.co.uk	0.316
stuff.co.nz	0.235	euronews.com	0.317
huffingtonpost.co.uk	0.235	daytondailynews.com	0.317
newsone.com	0.237	today.com	0.317
thesun.co.uk	0.237	hawaiinewsnow.com	0.318
presstv.ir	0.242	startribune.com	0.318
lifeneews.com	0.243	cnn.com	0.319
minutemennews.com	0.247	toprightnews.com	0.319
vancouversun.com	0.247	commondreams.org	0.32
enterpriseneews.com	0.247	theweek.com	0.322
collective-evolution.com	0.254	sfgate.com	0.323
sputniknews.com	0.254	vice.com	0.323
addictinginfo.org	0.258	telegraph.co.uk	0.324
catholicnewsagency.com	0.26	about.com	0.324
news.com.au	0.262	npr.org	0.325
motherjones.com	0.27	bbc	0.326
newsok.com	0.273	talkingpointsmemo.com	0.326
globalpost.com	0.276	newspapers.com	0.326
newser.com	0.277	cnsnews.com	0.326
nymag.com	0.277	nationalinterest.org	0.327
cbc.ca	0.28	aljazeera.com	0.328
deseretnews.com	0.281	examiner.com	0.328
rt.com	0.283	countercurrentnews.com	0.328
slate.com	0.285	alarabiya.net	0.329
buzzfeed.com	0.285	dailymail.co.uk	0.33
newrepublic.com	0.285	voanews.com	0.331
smh.com.au	0.285	salon.com	0.331
rawstory.com	0.286	news.google.com	0.331
abc.net.au	0.287	daily kos	0.331
abovetopsecret.com	0.288	nytimes.com	0.332
abs-cbnnews.com	0.288	nydailynews.com	0.332
dw.de	0.291	abcnews.com	0.333
ynetnews.com	0.291	theguardian.com	0.334

Site	% Con	Site	%Con
worldtruth.tv	0.335	americannews.com	0.37
newsherald.com	0.335	dailynews.com	0.371
onenewspage.com	0.335	foxnews.com	0.372
globalnews.ca	0.336	vox.com	0.376
themoscowtimes.com	0.336	theglobeandmail.com	0.379
csmonitor.com	0.337	therightscoop.com	0.384
huffingtonpost.com	0.337	knoxnews.com	0.385
firstcoastnews.com	0.337	news-record.com	0.385
detroitnews.com	0.338	newsday.com	0.392
ifyouonlynews.com	0.338	lifesitenews.com	0.392
msn news	0.339	newsmax.com	0.393
philly.com	0.339	bangordailynews.com	0.396
dallasnews.com	0.34	redflagnews.com	0.399
latimes.com	0.34	breakingnews.com	0.402
chicagotribune.com	0.34	theepochtimes.com	0.404
newsweek.com	0.341	breitbart.com	0.404
upi.com	0.341	newsiosity.com	0.411
yahoo news	0.341	rightwingnews.com	0.416
buffalonews.com	0.342	mercurynews.com	0.419
independent.co.uk	0.345	thedailybeast.com	0.425
madworldnews.com	0.346	nzherald.co.nz	0.427
foreignpolicy.com	0.347	express.co.uk	0.43
theblaze.com	0.348	gulfnews.com	0.431
cbsnews.com	0.348	charismanews.com	0.436
pbs newshour	0.349	news-gazette.com	0.438
nj.com	0.349	beforeitsnews.com	0.45
bloomberg.com - politics	0.35	townhall.com	0.45
freep.com	0.35	standard.co.uk	0.45
independent.ie	0.351	newsdaymarketing.net	0.451
washingtonpost.com	0.351	nationalreview.com	0.451
usatoday.com	0.356	drudgereport.com	0.454
medium.com	0.356	onlinenewspapers.com	0.454
wn.com	0.356	thejournal.ie	0.454
newsbusters.org	0.356	ctvnews.ca	0.457
newyorker.com	0.356	kingworldnews.com	0.465
trueactivist.com	0.357	firstpost.com	0.466
newszoom.com	0.357	spiegel.de	0.47
breakingisraelnews.com	0.36	twcnews.com	0.473
winknews.com	0.36	dailystar.co.uk	0.474
metro.co.uk	0.361	lemonde.fr	0.479
nypost.com	0.362	digitaljournal.com	0.497
wsj.com	0.362	newsminer.com	0.532
jpost.com	0.363	erietvnews.com	0.579
israelnationalnews.com	0.365	marinij.com	0.588
ijreview.com	0.366	news-press.com	0.65
politico.com	0.369		

Appendix C3: Media Diet, Partisan Slant

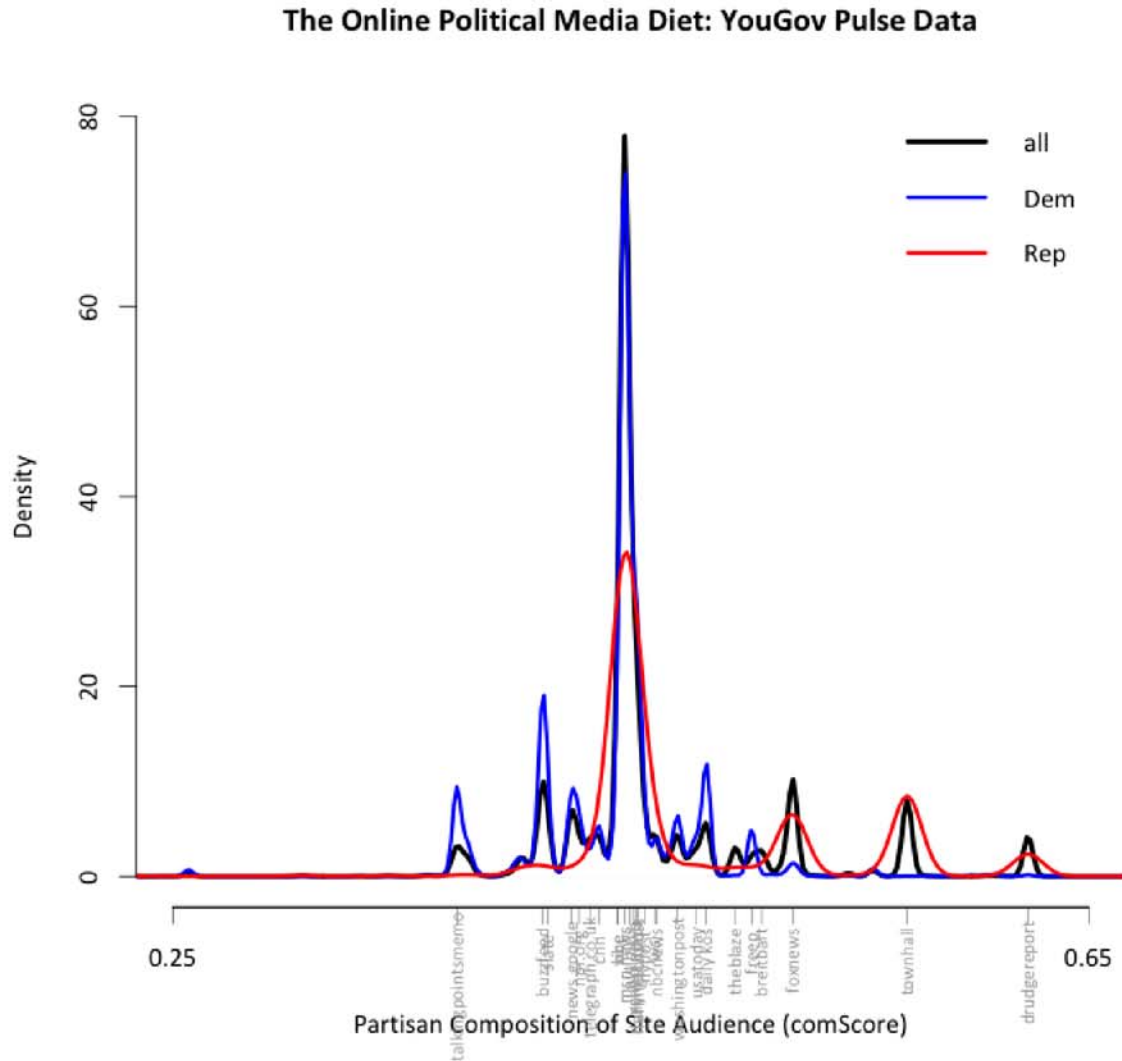


Figure C1: Density plot of aggregate site visits from the YouGov Pulse sample. Site partisan slant on the x-axis is measured using comScore data on audience composition (see Appendix B). N = 102,134 visits.

Appendix C4: Additional Figures and Tables




Figure C2: As this figure shows, most respondents in the YouGov Pulse panel visited no or very few political news sources during the three-week period the data was collected. Barely visible on the right are a tiny proportion of panelists who logged thousands of hits to political sites during that period.

Table C3: Opinions about regulating for-profit colleges, earlier study.

	DV: Regulate (T2)				
	(1)	(2)	(3)	(4)	(5)
Regulate (T1)	0.63*** (0.05)	0.60*** (0.06)	0.60*** (0.06)	0.59*** (0.06)	0.59*** (0.06)
Z	0.17** (0.08)	0.17** (0.08)	0.30*** (0.11)	0.15** (0.08)	-0.07 (0.13)
D		0.11 (0.08)	0.26** (0.12)		
R		-0.15 (0.14)	-0.08 (0.17)		
D:Z			-0.29* (0.15)		
R:Z			-0.12 (0.29)		
Lib				0.02 (0.09)	-0.13 (0.13)
Con				-0.17 (0.14)	-0.40** (0.19)
Lib:Z					0.27* (0.16)
Con:Z					0.43 (0.27)
Constant	1.41*** (0.23)	1.48*** (0.24)	1.41*** (0.25)	1.60*** (0.27)	1.74*** (0.27)
Observations	318	318	318	304	304
Adjusted R ²	0.51	0.51	0.52	0.49	0.50

Note: *p<0.1; **p<0.05; ***p<0.01
Weighted regressions, robust standard errors in parentheses.