MIT Open Access Articles

*How the news media activate public
expression and influence national agendas*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: King, Gary, Benjamin Schneer, and Ariel White. "How the News Media Activate Public Expression and Influence National Agendas." *Science* 358, no. 6364 (November 9, 2017): 776–780.

As Published: <http://dx.doi.org/10.1126/science.aao1100>

Publisher: American Association for the Advancement of Science (AAAS)

Persistent URL: <http://hdl.handle.net/1721.1/119647>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

How The News Media Activates Public Expression and Influences National Agendas

Gary King^{*} Benjamin Schneer[†] Ariel White[‡]

October 6, 2017

Abstract

We demonstrate that the news media causes Americans to take public stands on issues, join national policy conversations, and express themselves publicly more often than they would otherwise — all key components of democratic politics. We recruited 48 mostly small media outlets that allowed us to choose groups of outlets to write and publish articles, on subjects we approved, and dates we randomly assigned. We estimate the causal effect on proximal measures, such as website pageviews and Twitter discussion of the articles' specific subjects, and distal ones, such as Twitter conversation about the general policy area. Our intervention increased discussion in each broad policy area by $\approx 62.7\%$ (relative to a day's volume), accounting for 13,166 additional posts, with similar effects across population subgroups.

Science MS#: aao1100

^{*}Albert J. Weatherhead III University Professor, Institute for Quantitative Social Science, 1737 Cambridge Street, Harvard University, Cambridge MA 02138; GaryKing.org, King@Harvard.edu, (617) 500-7570.

[†]Assistant Professor, Department of Political Science, Florida State University, 113 Collegiate Loop, 531 Bellamy, Tallahassee, FL 32306; bschneer@fsu.edu.

[‡]Assistant Professor, Department of Political Science, MIT, 77 Massachusetts Avenue Building E53-461, Cambridge, MA 02139; arwhi@mit.edu.

The fields of political communications in general and media effects in particular are broad, deep, methodologically sophisticated, and central to much of the social sciences. They have covered persuasion (1), agenda setting (2, 3), attitude formation (4), diffusion, gate keeping (5), priming and agenda setting (6), issue framing (7), and numerous other topics, and are built on an incomparably wide range of intellectual traditions (8, p.174).

We focus here on an aspect of political communications with special relevance to the study of representative democracy — how the news media activates public expression, causing citizens to discuss major issues of policy and politics as part of the ongoing, collective “national conversation”. A well-functioning democracy larger than the sum of individual attitudes and behaviors requires public discussion and engagement among citizens on major issues of the day (9–11). Indeed, “political participation is not merely about trying to influence policy but also about trying to induce others to participate and give voice” (12). Although governments may easily dismiss any individual’s opinion, collective public expression powerfully impacts the behavior of government officials and the public policies they promulgate. The power of collective expression is a central feature of both representative democracy — where “the more the people are aware of each other’s opinions, the stronger the incentive for those who govern to take those opinions into account” (13) — and autocracy (14, 15). Citizens may join this national conversation to deliberate with each other (16), or simply “to give testimony” in the presence of others (17).

We thus study the effects of the media on the classical notion of *expressed* public opinion, a concept predating modern survey research, and with a focus not on changes in individual behavior or attitudes but instead on the content of the national conversation (18, 19). This discussion could once only be measured by collecting “water cooler events” (20), listening to hallway and dinner conversations, reading newspaper editorials and political leaflets, or collecting soapbox speeches from public squares. Today, we can take advantage of the fact that much of the conversation has moved to, and is recorded in, the 750 million social media posts that appear publicly on the web every day.

Unfortunately, estimating the effect of the news media is extremely challenging (21,

p.267). Scholarship dating back more than a century has had to contend with severe endogeneity, since media outlets are businesses competing for readers, catering to their interests. Large scale randomization of news content is normally impossible due to high costs, logistical infeasibility, and even some ongoing miscommunication between the journalistic and scientific communities regarding the norms of the former and goals of the latter. Even if randomization is possible, avoiding spillover effects is difficult because any media intervention can affect all potential research subjects in the nation at the same time. The result is often “profound” biases in estimated effects that can be greater than 600% difference from the truth (22, 23) given common levels of endogeneity, measurement error, and self-selection (see also 24). These biases have been addressed by scholars in some of social science’s most creative observational studies, although these approaches are well suited to answering certain questions (such as those for which instruments are available) but not others (e.g., 25–32). The biases are also addressed via elegant experiments and quasi-experiments, often made possible by studying different quantities of interest, such as individual-level effects or occasionally the effects on aspects of the national conversation (26, 33–42).

We attempt to tackle these methodological issues head on by enlisting a large number of small to mid-sized news media outlets that allowed us to run an unusual set of experiments. We developed and implemented an “incentive-compatible” research design that enables *both* full randomized experimental control in the hands of the researchers, so we could accomplish our scientific goals, *and* full editorial control in the hands of the journalists, fitting into their familiar customs and practices, so they could participate. Forty-eight small to mid-sized news media outlets participated in our research (Section S1.7). Seventeen of these outlets were part of our preliminary trial run experiments, provided information, and were helpful in other ways, and 33 were part of the experimental protocol we now describe (2 participated in both stages on unrelated stories). In addition, 13 others provided information, advice, or proprietary data, but were not part of our randomized interventions.

Our work was aided by the natural interest journalists have in understanding the im-

pact of their work. However, they are also competitors, trying to scoop each other. The difficulty is compounded by the fact that we asked these professionals to take actions few journalists have ever before agreed to, to allow researchers to participate in ways that rarely happen, and to share proprietary information with us that they do not even share with each other. We also needed to secure numerous individual agreements and arrange large scale coordination among competing entities over almost five years. As such, much of our effort involved building relationships, trust, and common understanding. We designed our experimental protocol to ensure that both our scientific goals and the journalists' professional goals were maximized.

In addition to help from our 61 independent news outlets, an industry association (The Media Consortium, representing about half of our outlets) helped us coordinate with the outlets and received funding to offer small financial incentives to some outlets, following their usual funding procedures. Our research team also received some direct funding from the same source. To protect the journalistic integrity of the numerous professionals who participated in our experiments, and the reputation of their publications, we do not reveal the specific articles in our experiment, which outlet published each article for us, or any potentially identifiable individual level aspects of the data we collected. We retained full rights to scholarly publication, without any required review or preapproval. To maintain a high level of realism, we tried to ensure that media outlets followed their standard operating procedures, embedding our treatment within their ordinary routines. The resulting protocol made our design more expensive, logistically complicated, and time consuming, but it should be more generalizable and compatible with the goals and norms of the journalistic and scientific communities.

Our treatment protocol has five parts. First, we choose a broad *policy area* from a set of eleven areas of both major national importance and sufficient interest to our news media outlets: *race, immigration, jobs, abortion, climate, food policy, water, education policy, refugees, domestic energy production, and reproductive rights*. Combining all eleven policy areas together, rather than using only one, greatly expands the representativeness of our study at the potential cost of a larger sample size needed.

Second, we choose a set of news media outlets and induce content correlations across them in ways that mirror common practices. Sometimes referred to as “pack journalism”, these practices include following each other, writing stories on the same subjects, “piling on” immediately after a story is broken by one outlet, occasionally collaborating, and sometimes even coauthoring stories. Although this behavior is sometimes criticized, professional journalists follow these venerable practices to help get stories out and ensure they reach a wide variety of differentiated audiences. We simulate the effects of pack journalism by following a procedure occasionally used by outlets to collaborate before publication, under negotiated ground rules. By employing a *project manager*, a group of outlets agree to collaborate on a specific story for a limited time. Participating outlets share information and publish simultaneously, often with assistance of the outlet hosting the project manager, offering staff, information, visualizations, or promotional materials. These fiercely independent sites even agree to effectively delegate aspects of editorial control to the project manager because, in addition to increasing their collective impact, each site retains the ability to opt out if necessary. This mechanism gives full editorial control over what is part of the collaboration to the project manager but full control over what is published to individual outlets. A prominent recent example is the Pulitzer prize-winning “Panama Papers” investigation (see bit.ly/kppapers and j.mp/ppapers). Playing the role of a project manager, without being based at one of the outlets, had the added advantage of making it easier for the outlets to share information with us that they would not normally share with each other.

We thus intervene for each experiment with what we refer to as a *pack* of 2–5 outlets (with a mean of 3.1 across all our experiments) rather than one. To ensure that outlets had experience in a chosen policy area, and sufficient enthusiasm for the subject matter and their collaborators, we allowed outlets to volunteer to join a pack for each experiment. We then asked them to collaborate as they would normally under this familiar structure. We retained approval rights to the collaboration to satisfy our scientific goals, and journalists and editors retained the right to opt out (prior to randomization) to satisfy journalistic standards; good communication kept either from exercising these rights in practice.

Third, while we control the collaboration as the project manager usually does, we leave the journalists the discretion they normally have. To do this, we have the pack select a specific *subject* for articles they write within our chosen policy area (planning for each outlet in a pack to write one article). For example, if the (broad) area is technology policy, the (specific) subject of the articles might be what Uber drivers think about allowing driverless cars, how a new trade agreement impacts hiring at local technology firms in Philadelphia, or others. The articles can be of any type they normally publish, ranging from large scale investigations, to interview-based journalism, to conceptual or opinion pieces. The journalists and their outlets naturally seek newsworthy articles, and also subjects that will remain of public interest whenever our random assignment mechanism (see Section S1.2) determines they will run. This ruled out some of the most impactful stories based on breaking news. We retained the right to reject a subject if the pack’s choice was outside our policy area or any individual article by an outlet in a pack, and the outlets retained the right to publish whatever they wished outside of our experiment; as above, good communication kept each to a minimum.

Fourth, we implemented a matched pair randomized experimental design. This design generates considerably more statistical power, robustness, and efficiency than classical randomization designs (Section S1.2; (43)). To avoid spillover effects or model dependent inferences, our unit of treatment is the entire nation during an experiment-week, with the treatment being a set of several articles published by a pack of outlets on the publication day (usually Tuesday) of a week we determine. We choose a pair of consecutive weeks matched for similarity of predicted news content (Section S1.2). Then we randomly assign one week to be the “treatment” week, during which the pack runs their stories, and one to be the “control week” where they are asked to behave as usual.

Each news media outlet then distributes its content as it usually would, via its web site, print media, video reports, audio podcasts, etc. As with all modern news media, each outlet also promotes its content with advertising via social media, Google adwords, email lists, and search engine optimization techniques, among others; they also often co-promote with others in the same pack. Near as we can tell, they followed the same

practices for articles in our experiment as those they run ordinarily. We also go to great lengths to ensure that the policy areas, subjects, and articles we choose appear indistinguishable from the normal type and flow of articles written by these news media outlets in the course of their ordinary business practices. And to the best of our knowledge, no outlet received any reader communications about an article or practice that seemed unusual or out of place.

Finally, we avoid intervening on any one news media outlet so often as to get in the way of its normal practices, change the character of its site, or be discovered by readers. This is why we needed to organize a large number of separate outlets, from which we could choose different packs of outlets for each experiment, rather than using only one small pack of 2–5 repeatedly. This procedure adds causal heterogeneity and thus requires a larger n overall, but should generate a more representative causal effect.

Because the cost of collecting each observation in our design corresponds to an entire experiment in most designs, we follow two additional procedures to reduce costs. First, we ensure we collect only as much data as necessary by inverting the usual approach to statistical inference via sequential hypothesis testing, including a novel nonparametric sequential technique we developed (see Sections S1.6, S2.3, S3). And second, we evaluate multiple observable implications of our intervention rather than only one. Thus, Figure 1 portrays points we can measure on the causal pathway from the treatment intervention (far left) to our ultimate outcome variable of interest (far right). The first link is the causal effect of the treatment intervention on the number of articles published. If we found that instructing sites to publish articles in a given week had no effect, we would know to be skeptical of an intent-to-treat effect on social media posts. This is not a deterministic step because unexpected events can cause the media outlets to publish on a chosen subject more than expected in either of the two weeks. Media outlets, as ongoing competitive businesses, may sometimes be forced to respond to unexpected events in ways that violate an experimental protocol. Fortunately, the randomized assignment in our design prevents such “noncompliance” from inducing bias in the intent-to-treat effect, but it could introduce heterogeneity and smaller effects overall, both of which would lead us to need a

larger n given a chosen level of uncertainty.

The next arrow in this causal pathway connects articles published to website pageviews (the number of users visiting a web page). This includes the articles we commissioned and any other in the same policy area in either treatment or control weeks. The second arrow in Figure 1 then refers to a causal effect, which is positive only if more people visit pages with articles in the policy area in treated than control weeks. In our design, the only plausible way for either our treatment or the publication of news articles by media outlets to have an effect on either measure of public expression of opinion is through at least some people reading the articles, usually on the outlets' web pages. We portray this in the figure by the absence of paths, other than through outlet website pageviews, from the randomized treatment or published articles to expression in broad policy areas in social media. However, pageviews can cause social media participants to express themselves publicly on broad national policy issues either directly (see curved arrow in Figure 1) or by reading social media posts written narrowly about the subject of the published articles (see arrows to and then from "Posts on Subject").

We give detailed uncertainty analyses in Sections S2.3 and S3 on the scale of false positive rates; here, we present causal estimates on the scale of our outcome variables and quantities of interest (Section S1.3,S1.4). We present two sets of results, each using model-based and model-free estimation (Section S1.5).

First, Figure 2 reports estimates of the main quantity of interest in our experiment — the average causal effect of a pack of journalists publishing articles, at a time we randomly determine, on the extent to which Americans express themselves publicly on social media in a chosen broad policy area. The left panel gives the effect for each day in terms of a percentage change in social media posts, with corresponding estimates of the absolute numbers of posts in the right panel. Within each panel, we estimate the causal effect for each day following the intervention and the total effect (the horizontal axis). We do this with our model-based estimator (red dots connected with a line, and solid black square for total) and with our model-free estimator (open circles, and black open square for total, respectively).

The figure shows that our experimental treatment causes the number of social media posts appearing on the first day in a broad policy area to increase by 19.4%, according to our model-based estimator (Figure 2, left panel, leftmost red dot, ●). From the red dot in the same position in the right panel, we can see that social media users wrote and published approximately 4,442 additional posts solely as a result of our intervention. Moreover, we estimate that the same articles continue to have effects over the rest of the week. On average, these effects decline as we get farther out from publication day, with approximately zero effect on average by the sixth day (consistent with (44, 45); see also (46)).

The total effect, represented by a black square (■) at the top right of in each panel, indicates that our experimental intervention overall caused a 62.7% increase in social media posts over the week relative to the average day’s volume (or 10.4% relative to the entire week), which on average in a policy area accounts for Americans writing a total of 13,166 additional social media posts solely because of our intervention. The figure also includes estimates from our model-free approach (open circles ○ and open square □). These estimates offer the advantage of avoiding modeling assumptions but have the resulting disadvantage of higher variance. Yet, they clearly convey the same overall pattern in causal effects.

We present detailed uncertainty estimates (see Sections S1.6, S2.3, S3, and S4.1). In addition, given the reasonable hypothesis that the causal effect varies smoothly over days of the week, then the degree to which the model-free estimates (the circles) vary around the model-based results (the line) provides another estimate of the uncertainty of our primary causal effects. As can be seen from this perspective, these estimates have relatively low levels of spread (or uncertainty) around them and are clearly above zero.

Second, in Figure 3, we estimate the effect of our intervention on different subgroups expressing themselves in a broad policy area. The subgroups we were able to define include political party (Democrats, Republicans, and unknown), gender (men, women), region (Northeast, Midwest, West, South), and degree of influence on Twitter (high and low). (The party, gender, and region of social media posts are based on Twitter metadata,

supplemented by analyses of Twitter bios and follower structures; influence is based on numbers of followers and the social graph. See Section S5.) As a reference, we add to each graph a (red) line representing all posts (taken from Figure 2), but we omit the model-free estimates for graphical clarity. The interesting result from this analysis is the lack of a result: the difference between any pair of subgroups within a panel is always small (and never statistically distinguishable from zero). Apparently, the national conversation really is one conversation, at least among those able to participate in social media; even if they do not interact with each other, the evidence indicates that they are being influenced in similar ways by the news media.

Finally, the outcome variable in Figure 2 is based on total number of posts in a policy area, and is designed to measure the national conversation and how it was impacted by our randomized treatment. We now present another observable implication of media effects in Figure 4, counting only the daily number of unique authors of posts rather than the total number of posts. This figure demonstrates that more Americans were engaged by the articles in this policy area (rather than the same people posting more), an important observable implication of our thesis. The causal effect of our intervention, on the first day, was an increase of 23.9% more unique authors (accounting for 3,287 more individuals) participating in the national conversation in a broad policy area. This effect drops to essentially zero by the fifth day. This result also makes bots writing large numbers of posts less likely to account for our results (47) (Section S5.1.2).

Section S2.4 demonstrates that the news media changes the composition of expressed opinion by 2.3 percentage points in the direction of that held by our published articles. And S2.5 shows effects on other observable implications, including effects on website pageviews and on discussion on the specific subject of the articles. Overall, our experiments reveal large news media effects on the content of the national conversation across eleven important areas of public policy, political party, gender, region, and level of social influence.

We put these effect sizes in context, and then discuss their implications. First, the subjects of the articles in our treatments are limited to those which the journalists are willing

to write about, and their outlets are willing to publish, at randomly determined times, days or weeks after they were conceived. Additionally, searching for weeks to constitute good matched pairs, in the service of reducing bias and inefficiency, typically led us to select news periods predicted to be relatively “quiet” (predictions which turned out to be relatively accurate; see Section S6.2). The media effects in other weeks, such as when outlets publish stories to ride a viral social media wave or to satisfy the intense interest of a major breaking story, may of course have different effect sizes than we reported. Our effect sizes are small compared to huge entertainment events (e.g., they are about a hundredth the size of the Twitter frenzy generated by a new episode of the television show *Scandal*; [j.mp/SCandal1](https://www.youtube.com/watch?v=j.mp/SCandal1)). Still, they represent important and substantial increases in policy discussions, and indicate that the media is causing many more people to express themselves publicly, and for each person to express themselves more frequently, than would otherwise be the case.

The intervention we studied was the result of only 2–5 small-to-medium sized outlets. To glean what our effects might have been if we had recruited larger media outlets, we conducted informal observational analyses where randomization or a large n was infeasible. We searched unanticipated *New York Times* stories, on topics where *Times* reporters scooped other outlets or reported on surprise events during periods with few other stories. For example, we found a news story about a previously embargoed scholarly article in the area of fracking affecting drinking water, at a time when little else in the policy area was being discussed ([j.mp/frackH2O](https://www.youtube.com/watch?v=j.mp/frackH2O)). We observed a one day spike in discussion in the broad area of water quality and related issues of over 300% (compared to a 19% effect size in our study). And numerous public policy issues have far higher visibility than fracking, many with far more impactful “interventions”. Although further research is needed to confirm this large outlet effect, it appears that some articles published may have a multiple of this already large effect size. We thus expect that the causal estimates in our experiment are an underestimate of the overall effect of the news media.

Our results should remind us of the importance of the ongoing and interconnected national conversation Americans have around major issues of public policy. This con-

versation is a fundamental characteristic of modern large-scale government, the content of which has major implications for the behavior of office holders and public policies. We also find — among those who participate in social media — that the effects of the media are approximately the same across citizens of different political parties, genders, regions, and influence in social media, further supporting the idea that the conversation is truly national. Given the tremendous power media outlets have to set the agenda for public discussion, the ideological and policy perspectives of those who own media outlets has considerable importance for the nature of American democracy and public policy. The ideological balance across the media ecosystem, among the owners of media outlets, needs considerable attention as well (48). The ability of the media to powerfully influence our national conversation also suggests profound implications for future research on “fake news” potentially having similar effect sizes (49), “filter bubbles” potentially reducing or directing these effects (50), among others.

Social scientists have long been interested in estimating the impact of the media on how Americans participate in the national conversation on important public policy issues, but other important issues, such as media effects on individual attitude formation and persuasion, also need to be subject to randomized experiments. Similarly, further research is needed in areas beyond the 11 policy areas we studied. Further research should also be conducted with outcome variables beyond our social media measures, beyond social media, and with media outlets with different ideological perspectives than in our sample. Finally, although we have been able to estimate the causal effect of some of the news media, we have not measured how often actual media outlets make efforts to move different populations of Americans to express themselves about specific policy areas. We encourage future researchers to take up these measurement challenges, and the numerous other topics that may shed light on the formation, development, and changes in the effect of the media on citizen engagement in the national conversation.

References

1. J. T. Klapper, *The effects of mass communication* (Free Press, 1960).

2. L. Guo, M. E. McCombs, *The Power of Information Networks: New Directions for Agenda Setting* (Routledge, 2015).
3. M. E. McCombs, D. L. Shaw, *Public opinion quarterly* **36**, 176–187 (1972).
4. J. R. Zaller, *The Nature and Origins of Mass Opinion* (Cambridge University Press, New York, NY, 1992).
5. M. A. Baum, T. Groeling, *Political Communication* **25**, 345–365 (2008).
6. S. Iyengar, D. R. Kinder, *News that matters: Television and American opinion* (University of Chicago Press, 2010).
7. T. E. Nelson, in *The Oxford handbook of American public opinion and the media* (Oxford University Press Oxford, 2011), pp. 156–168.
8. W. L. Bennett, S. Iyengar, *Journal of Communication* **58**, 707–731 (2008).
9. W. A. Gamson, *Talking Politics* (Cambridge University Press, New York, NY, 1992).
10. J. G. Gunnell, in *The Oxford handbook of American public opinion and the media* (Oxford University Press Oxford, 2011), pp. 1–18.
11. P. M. Kellstedt, *The Mass Media and the Dynamics of American Racial Attitudes* (Cambridge University Press, New York, NY, 2003).
12. D. Carpenter, C. D. Moore, *American Political Science Review* **108**, 479–498 (2014).
13. B. Manin, *The principles of representative government* (Cambridge University Press, 1997).
14. G. King, J. Pan, M. E. Roberts, *American Political Science Review* **107**, 1–18 (2013).
15. G. King, J. Pan, M. E. Roberts, *Science* **345**, 1–10 (2014).
16. J. Fishkin, *Election Law Journal* **12**, 490–507 (2013).
17. L. M. Sanders, *Political theory* **25**, 347–376 (1997).
18. D. Hopkins, G. King, *American Journal of Political Science* **54**, 229–247 (Jan. 2010).
19. S. Verba, K. L. Schlozman, H. E. Brady, *Voice and Equality: Civic Volunteerism in American Politics* (Harvard University Press, Cambridge, MA, 1995).
20. M. A. Baum, *Soft news goes to war: Public opinion and American foreign policy in the new media age* (Princeton University Press, 2003).
21. L. M. Bartels, *American Political Science Review* **87**, 267–285 (1993).
22. L. Vavreck, S. Iyengar, in *The Oxford handbook of American public opinion and the media* (Oxford University Press Oxford, 2011), pp. 156–168.
23. L. Vavreck, *Quarterly Journal of Political Science* **2**, 325–343 (2007).
24. M. Prior, *The Journal of Politics* **71**, 893–908 (2009).
25. R. R. Bhavnani, G. Nellis, “Can Government-Controlled Media Cause Social Change? Television and Fertility in India”, University of Wisconsin, 2016, <http://j.mp/2yIeZoL>.

26. S. DellaVigna, E. L. Ferrara, “Economic and social impacts of the media”, tech. rep. (National Bureau of Economic Research, 2015), <http://j.mp/2yIkiVd>.
27. S. DellaVigna, E. Kaplan, *The Quarterly Journal of Economics* **122**, 1187–1234 (2007).
28. M. Gentzkow, J. M. Shapiro, M. Sinkinson, *The American Economic Review* **101**, 2980–3018 (2011).
29. M. Gentzkow, *The Quarterly Journal of Economics* **121**, 931–972 (2006).
30. G. A. Huber, K. Arceneaux, *American Journal of Political Science* **51**, 957–977 (2007).
31. A. Reeves, M. McKee, D. Stuckler, *Social science research* **56**, 44–57 (2016).
32. G. Smith, K. Searles, *Public Opinion Quarterly* **78**, 71–99 (1 2014).
33. J. Conroy-Krutz, D. C. Moehler, *The Journal of Politics* **77**, 575–587 (2015).
34. A. S. Gerber, D. Karlan, D. Bergan, *American Economic Journal: Applied Economics* **1**, 35–52 (2009).
35. D. P. Green, B. R. Calfano, P. M. Aronow, *Political Communication* **31**, 168–180 (2014).
36. D. P. Green, L. Vavreck, *Political Analysis* **16**, 138–152 (2008).
37. M. S. Levendusky, *American Journal of Political Science* **57**, 611–623 (2013).
38. D. C. Mutz, J. Soss, *Public Opinion Quarterly* **61**, 431–451 (1997).
39. E. L. Paluck, P. Lagunes, D. P. Green, L. Vavreck, L. Peer, R. Gomila, *PloS one* **10**, e0138610 (2015).
40. C. Panagopoulos, D. P. Green, *American Journal of Political Science* **52**, 156–168 (2008).
41. T. J. Ryan, T. Brader, *Political Science Research and Methods* **1**, 1–21 (2015).
42. M. D. Trujillo, E. L. Paluck, *Analyses of Social Issues and Public Policy* **12**, 113–132 (2012).
43. K. Imai, G. King, C. Nall, *Statistical Science* **24**, 29–53 (2009).
44. A. S. Gerber, J. G. Gimpel, D. P. Green, D. R. Shaw, *American Political Science Review* **105**, 135–150 (2011).
45. S. J. Hill, J. Lo, L. Vavreck, J. R. Zaller, *Political Communication* **30**, 521–547 (2013).
46. L. M. Bartels, *Political Communication* **31**, 532–544 (2014).
47. C. A. Bail, *Proceedings of the National Academy of Sciences* **113**, 11823–11828 (2016).
48. E. Noam, *Media ownership and concentration in America* (Oxford University Press, 2009).
49. H. Allcott, M. Gentzkow, *Journal of Economic Perspectives* **31**, 211–235 (2017).
50. F. J. Zuiderveen Borgesius, D. Trilling, J. Moeller, B. Bodó, C. H. de Vreese, N. Helberger, *Internet Policy Review* **5** (2016).

Acknowledgements

Our thanks to Jo Ellen Kaiser and John Schwartz for suggestions and advice over many years; Matt Baum, Kosuke Imai, Jim Katz, Jonathan Ladd, Walter Mebane, Dan Moskowitz, Diana Mutz, Jamie Robins, Lynn Vavreck, and Xiang Zhou for helpful comments on earlier versions; Julia Ansolabehere, Manolia Charlotin, Elizabeth Garcia, Kevin Hornbeck, Nicole Levin, Gilberto Lopez, Caroline Marshall, and Devi Nair for expert research assistance; Crimson Hexagon for data and computational support; and Voqal, The Media Consortium, and Harvard University's Institute for Quantitative Social Science for research support. We also thank Harvard University's Center for American Political Studies and Graduate School of Arts and Sciences for graduate fellowships to Benjamin Schneer and Ariel White; the Program on Inequality at Harvard's Kennedy School and Radcliffe Institute for graduate fellowships to Ariel White; and the Institute for Quantitative Social Science and the Tobin Project for graduate fellowships to Benjamin Schneer. Replication data is available in the Harvard dataverse at ***.

Figure 1: The causal path from randomized treatment (first point) to public expression on broad policy areas (last point)

Figure 2: Causal effect of news media on public expression, denominated in percent change (left panel) and absolute change (right panel) in numbers of social media posts in a broad policy area. Effects appear as the percent change in social media posts for each day of the week (estimated by our model-based estimator, ●, and our model-free estimator, ○) and the total overall (■ and □, respectively).

Figure 3: Causal effect of the news media on the percent change in social media posts by political party, gender, region, and influence on Twitter, with axes defined as in the left panel of Figure 2.

Figure 4: Causal effect of randomized treatment on the number of unique authors expressing themselves in the same policy area as the intervention, in percent change (left) and absolute numbers of posts (right), for each day (red dot, ●) and total overall (black square, ■).