



Social Media & Democracy:

Assessing the State of the Field & Identifying Unexplored Questions

Report from a conference organized by the **Media & Democracy** program at the **Social Science Research Council**

Kris-Stella Trump

In collaboration with
Jason Rhody, Cole Edick,
and Penelope Weber

April 19–20, 2018
Stanford University



The Frances C. Arrillaga Alumni Center

Introduction

This report summarizes the highlights of a discussion that took place at Stanford University on April 19-20, 2018. The conference, titled “Social Media and Democracy: Assessing the State of the Field and Identifying Unexplored Questions,” convened leading social scientists to discuss the state of the field with regard to research on social media and democracy.

The conference format encouraged participants to discuss what they saw as the most important recent evidence on the relationship between social media and democracy, to highlight gaps in existing knowledge, and to suggest relevant topics for future research. Rather than giving traditional research presentations, panelists were asked to share their reflections on the state of the field and identify outstanding research questions, followed by a discussion among all gathered scholars and the audience.

Reflecting the spirit of the conference, this report focuses on high-level takeaways from the discussion, and we do not summarize individual presentations or contributions here. We chose this report format to encourage all participants to speak freely and exchange ideas, in the knowledge that their specific statements would not be attributed to any one individual but would rather be summarized in a collective manner. Thus, this report gives a bird’s-eye view of points of agreement, items of discussion, and recommended future research questions.

We hope that readers—scholars and practitioners alike—will find this report thought-provoking and stimulating. Above all, we hope that this report will inspire future research, foster interdisciplinary collaboration, and help consolidate the current moment in a fast-changing field of scholarly inquiry.

About the Organizers

The conference was convened as part of the Media & Democracy program (mdn.ssrc.org) at the Social Science Research Council (www.ssrc.org). The Media & Democracy program encourages academic research, practitioner reflection, and public debate on all aspects of the close relationship between media and democracy.

The lead conveners of the conference were Nathaniel Persily (James B. McClatchy Professor of Law at Stanford University) and Diana Mutz (Samuel A. Stouffer Professor of Political Science and Communication at the University of Pennsylvania), with assistance from Media & Democracy program staff, led by program codirectors Kris-Stella Trump and Jason Rhody.

This report was produced by Kris-Stella Trump in collaboration with Jason Rhody, Cole Edick, and Penelope Weber.

The conference was cosponsored by the Center on Philanthropy and Civil Society at Stanford University and the Center for Advanced Study in the Behavioral Sciences at Stanford University.

This event was made possible by the generous funders of the Media & Democracy program: the John S. and James L. Knight Foundation, the Democracy Fund, and the William and Flora Hewlett Foundation.

Participants

Jonathan Albright

Director of Research, Tow Center for Digital Journalism
Columbia University

Matthew Baum

Marvin Kalb Professor of Global Communications
Harvard University

Susan Benesch

Director, Dangerous Speech Project
Berkman Klein Center for Internet and Society
Harvard University

Adam J. Berinsky

Mitsui Professor of Political Science
Massachusetts Institute of Technology

Damon Centola

Associate Professor of Communication
University of Pennsylvania

Renée DiResta

Policy Lead
Data for Democracy

Annie Franco

Data Scientist
Facebook

Kelly Garrett

Associate Professor of Communication
Ohio State University

Matthew Gentzkow

Professor of Economics
Stanford University

Bryan Gervais

Assistant Professor of Political Science and Geography
University of Texas at San Antonio

Shanto Iyengar

Professor of Political Science
Stanford University

Nina Jankowicz

Global Fellow, Kennan Institute
Woodrow Wilson International Center for Scholars

Linda Kinstler

Graduate Student, Rhetoric
University of California, Berkeley

Diana Mutz

Samuel A. Stouffer Professor of Political Science and Communication
University of Pennsylvania

Jennifer Pan

Assistant Professor of Communication
Stanford University

Nathaniel Persily

James B. McClatchy Professor of Law
Stanford University

David Rand

Associate Professor of Psychology
Yale University

Jaime Settle

Assistant Professor of Government
College of William and Mary

Monica Stephens

Assistant Professor of Geography
State University of New York at Buffalo

Emily Thorson

Assistant Professor of Political Science
Syracuse University

Joshua Tucker

Codirector, Social Media and Political Participation Lab
New York University

Theme Statement

The 2016 American elections intensified popular as well as scholarly interest in the relationship between media and democracy. The role of social media has featured particularly prominently in debates over fake news, information bubbles, and algorithmic propaganda. Increased scholarly interest in these themes is manifest in regularly occurring scholarly convenings that explore technological changes in the media, social interactions online, and their relationship to the quality of our democracy.

We convened this conference to take stock of the current moment in a fast-developing field, to discuss recent findings, and to identify key remaining research questions. To frame the issues at hand, the conference opened with remarks from Joshua Tucker (Professor of Politics, New York University). Prof. Tucker set the stage with a summary of findings from the recent report on "[Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature](#)," commissioned by the William and Flora Hewlett Foundation.

The ensuing panels addressed five different themes, all united under the umbrella of social media and democracy. The topics addressed by speakers included the prevalence of hate speech online, the quality of online and offline civic discourse, the distribution of fake news, the correction of false beliefs, ideological "echo chambers," and the use of social media by hostile foreign actors. The following sections in this report summarize highlights from each panel session. Each section also features a short list of "related reading" suggestions. These suggestions are not intended to be exhaustive but rather to reflect a selection of articles or books relevant to the subtopic raised during discussion.

Inflammatory Speech and Incivility Online

Are inflammatory speech and uncivil discourse particularly prevalent on social media? If so, why and with what consequences? What, if anything, is distinctive about the online environment that may encourage incivility?

ECHO CHAMBERS AND INCIVILITY: ONLINE VS. OFFLINE

In order to determine whether inflammatory speech and incivility are more common online, we first need to know how common they are in offline conversations. However, our knowledge of offline political exchanges—both civil and uncivil—is very incomplete. Therefore, the first task for researchers is to expand our understanding of all types of political exchanges, including those that take place offline. Research questions include: How often do political conversations occur? Who participates? How often do conversations occur across partisan lines? How often do they turn negative or inflammatory?

As an illustration of the above point, one possible explanation for why political interactions online might be more inflammatory is that cross-partisan exchanges are more common online, and we are unaccustomed to talking to people across ideological lines in other venues. In other words, if we are unaccustomed to dealing with cross-partisan exchanges offline, then suddenly engaging in these interactions online could spark incivility. This possibility, of course, contrasts with the alternative expectation that ideological echo chambers are more common online and that the insulation of these echo chambers may contribute to incivility. Research on political conversations offline as well as online would help us start disentangling these possibilities.

Future research should also expand existing efforts to better understand ideological echo chambers. The relative prevalence of such echo chambers online and offline is a contested question. We do not know what the consequences of a one-sided social media diet are; it is possible, and consistent with some recent

evidence, that increased cross-partisan dialogue would in fact amplify affective polarization. Difficulties defining, measuring, and comparing online echo chambers have prevented a consensus from emerging, and more research is needed before we can assess the net consequences of echo chambers with any level of certainty.

THE PREVALENCE AND CONSEQUENCES OF INCIVILITY AND HATE SPEECH

Recent research on hate speech suggests that it comprises a very small share of total activity on social media (the case studies under discussion primarily involved Twitter), and that it occurs in “bursts” of activity with a relatively short shelf life. This finding contrasts with more popular perceptions that hate speech usage on social media is endemic and/or rising. Further, it is not clear that hate speech on social media increased either during or after the 2016 US presidential campaign. Additional studies using alternative methods to identify and measure hate speech will help consolidate knowledge in this area. This also suggests an outstanding research question: If hate speech online did not become more frequent during this time period, why is there a common perception that it did?

Additionally, there is an outstanding need for more nuanced studies of the consequences of incivility and hate speech online. For example, there is limited and conflicting evidence on whether, how, and under what conditions incivility and hate speech online can change offline behaviors. More generally, research is needed on the wide range of variables that may influence the consequences of incivility. Potential research questions include: What impact does incivility have on targets and observers? Does the seniority and fame of the source of incivility matter? What are the effects of observing incivility that is congruent with one’s own attitudes? We need more work that disaggregates different types of incivility, as well as the range of consequences that incivility may have (for example, withdrawal from online discussions, fear, and negative affect toward outgroups). Finally, studies could also focus on audio/visual forms of incivility (including memes and GIFs). Cross-platform data collection of memes and GIFs, with accompanying metadata, may be an important step toward expanding our knowledge of how this form of incivility spreads, and its consequences.

FACTORS THAT ENCOURAGE OR ENABLE INCIVILITY

What factors enable online incivility? A common explanation for the spread of incivility is the ability to post anonymously, but anonymity by itself is probably not a sufficient explanation. For example, rates of inflammatory and/or hate speech vary considerably across platforms that enable anonymous posting and across communities on the same platform (see, for example, norm variation across subreddits). Such variation could be the result of socialization, reputation-building, self-selection, and/or accountability. In particular, anonymity is not the same as lack of identity/reputation; the ability to invest in the latter (even in formally anonymous social systems) may provide leverage points to increase accountability and reduce incivility.

Lack of (feelings of) accountability is likely to be a key variable for explaining incivility. Accountability may, for example, be enhanced by thinking about one's interlocutor as a human being, having geographic proximity to them, or sharing membership with them in offline communities. More research is needed on factors that produce accountability in anonymous spaces or that encourage reputation-building and discourage incivility.

RELATED READINGS

Gervais, Bryan T. "Incivility Online: Affective and Behavioral Reactions to Uncivil Political Posts in a Web-based Experiment." *Journal of Information Technology & Politics* 12, no. 2 (2015): 167–185. <https://doi.org/10.1080/19331681.2014.997416>.

Kumar, Srijan, William L. Hamilton, Jure Leskovec, and Dan Jurafsky. "Community Interaction and Conflict on the Web." *Proceedings of the 2018 World Wide Web Conference* (2018): 933–943. <https://doi.org/10.1145/3178876.3186141>.

Munger, Kevin. "Tweetment Effects on the Tweeted: Experimentally Reducing Racist Harassment." *Political Behavior* 39, no. 3 (2017): 629–649. <https://doi.org/10.1007/s11109-016-9373-5>.

Mutz, Diana. *In-Your-Face Politics: The Consequences of Uncivil Media*. Princeton: Princeton University Press, 2015.

Saleem, Haji Mohammad, Kelly P. Dillon, Susan Benesch, and Derek Ruths. "A Web of Hate: Tackling Hateful Speech in Online Social Spaces." *arXiv preprint* (2017). <https://arxiv.org/abs/1709.10159>.

Siegel, Alexandra, Evgenii Nikitin, Pablo Barber, Joanna Sterling, Bethany Pullen, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker. "Trumping Hate on Twitter? Online Hate Speech and White Nationalist Rhetoric in the 2016 US Election Campaign and its Aftermath." In progress.

Stephens, Monica, and Ate Poorthuis. "Follow thy neighbor: Connecting the social and the spatial networks on Twitter." *Computers, Environment and Urban Systems* 53 (2015): 87–95. <https://doi.org/10.1016/j.compenvurbsys.2014.07.002>.

Suhay, Elizabeth, Emily Bello-Pardo, and Brianna Maurer. "The Polarizing Effects of Online Partisan Criticism: Evidence from Two Experiments." *The International Journal of Press/Politics* 23, no. 1 (2018): 95–115.

Thorson, Kjerstin, Emily Vraga, and Brian Ekdale. "Credibility in Context: How Uncivil Online Commentary Affects News Credibility." *Mass Communication and Society* 13, no. 3 (2010): 289–313. <https://doi.org/10.1080/15205430903225571>.



Monica Stephens, Assistant Professor of Geography, State University of New York at Buffalo

Distribution and Effects of Fake News

The question of how to reliably identify false news sources and articles is a key ongoing discussion in this field. While questions of classification were discussed, this panel deliberately focused primarily on research that goes beyond issues of definition. What do we still need to learn about how—and with what consequences—misinformation spreads online?

WHO SHARES AND SEES WHAT

Recent evidence suggests that, on average, American citizens did not see many fake news stories in the runup to the 2016 elections. Against this baseline of low exposure, there is additional evidence that in the 2016 election, Republicans and independents shared more fake news than Democrats, although it is important to note that this may be the result of there being a larger supply of pro-Trump fake news stories. Despite the possibility of low average rates of exposure to fake news, we do not yet know whether these averages conceal small but highly exposed/targeted subgroups. For example, we do not yet know whether some geographic regions were disproportionately exposed to fake news stories.

Another recent finding is that older people are more likely to share fake news than younger people. In general, we know more about how fake news spreads than we do about the consequences it has on recipients. As with other topics in this rapidly moving field, more basic and descriptive research on who shares and reads what and on which platforms is needed.

As part of this important descriptive work, future research should pay attention to and test various definitions of exposure and engagement (e.g., appearance in feed, clicking through, etc.). Scholars should also carefully examine what engagement tells us about the state of mind of the individual (for example, whether people engage more with items whose veracity they question or those whose conclusions they believe).

THE IMPACT OF FALSE INFORMATION

More research is needed to understand the impact of false information. How do the effects of false information vary across

recipients, sources, disseminators, and timeframes? Specific research questions in this area could include: Do the dynamics of false information look different when it is spread during a breaking news event, as opposed to campaigns that take place over several months or years? Does online exposure to falsehoods have different consequences than offline exposure? Is online information more readily accepted when it is shared by a member of one's physical or geographic community? What factors make social media more or less likely to be a source of false information compared to other media forms? How is engagement with a story on social media influenced by conditions such as social endorsement, the presence or lack of editorial oversight, or the inadvertent exposure to false news (which can occur if social media is used primarily for entertainment purposes but includes occasional political posts)?

In addition to studying the consequences of false information on individuals who are exposed to it, we need a better understanding of its consequences on a platform level. For example, some recent evidence suggests that repeated exposure to a piece of information increases acceptance of it but reduces sharing. How does this affect the aggregate dynamics of false information spreading on social media platforms? Related research questions include: If the sharing of fake news on a social media feed goes up, does this affect the sharing of news of all kinds? Does it affect overall engagement with the platform?

Looking even more broadly, we face an even more challenging but crucial new area of research: the consequences of overall media diets. Researchers should develop ways to learn about an individual's whole media diet, including entertainment as well as news consumption across social media platforms, newspapers, streaming services, and other websites. This may require the creation of online panels of individuals who agree to have their news consumption anonymously tracked—such a panel could combine features from the Nielsen television rankings with those from online survey panels. Research with data of this kind could explore the correlates of overall media diets in the short term as well as the medium term.

POTENTIAL INTERVENTIONS

The creation of effective interventions to counter the negative effects of false information should be central within this research agenda. While these interventions could be developed from insights about the spread and consequences of false information, drawn from the research agenda outlined above, it will be crucial that interventions themselves also be tested for effectiveness.

One potential intervention, based on existing experimental evidence, might prompt people to reflect before sharing. Experimental evidence suggests that prompting people to reflect on the veracity of news makes them subsequently less likely to share news stories whose veracity they are unsure of.

Other potential interventions could build on a variety of research questions. Would interventions be more successful in online communities that are based on physical or geographic communities? We know that misperceptions persist more strongly when they are shared among a community of believers; how should interventions take this into account? What interventions to increase media literacy would be successful in places like churches or senior centers (in addition to the more usual location of schools)? Is the correction of false information easier or harder in an environment where exposure to news content is accidental (e.g., otherwise entertainment-oriented social media feeds)?

RELATED READINGS

Allcott, Hunt, and Matthew Gentzkow. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31, no. 2 (May 2017): 211–36. <https://doi.org/10.1257/jep.31.2.211>.

DiResta, Renée. (2016). "Social Network Algorithms Are Distorting Reality by Boosting Conspiracy Theories." *Fast Company*, May 11, 2016. <https://www.fastcompany.com/3059742/social-network-algorithms-are-distorting-reality-by-boosting-conspiracy-theories>.

Garrett, R. Kelly, Brian E. Weeks, and Rachel L. Neo. "Driving a Wedge Between Evidence and Beliefs: How Online Ideological News Exposure Promotes Political Misperceptions." *Journal of Computer-Mediated Communication* 21, no. 5 (2016): 331–48. <https://doi.org/10.1111/jcc4.12164>.

Guess, Andrew, Brendan Nyhan, and Jason Reifler. "Selective Exposure to Misinformation: Evidence from the Consumption of Fake News during the 2016 U.S. Presidential Campaign." Working Paper (2016). <https://www.dartmouth.edu/~nyhan/fake-news-2016.pdf>.

Guess, Andrew, Briony Swire, Adam Berinsky, John Jost, and Joshua Tucker. "Rumors in Retweet: Social Media and the Spread of Political Misinformation." In progress.

Pennycook, Gordon, and David G. Rand. "Crowdsourcing Judgments of News Source Quality." *SSRN Electronic Journal* (2018). <http://dx.doi.org/10.2139/ssrn.3118471>.

Pennycook, Gordon, and David G. Rand. "Susceptibility to Partisan Fake News Is Explained More by a Lack of Deliberation Than by Willful Ignorance." *SSRN Electronic Journal* (2018). <http://dx.doi.org/10.2139/ssrn.3165567>.

Correcting Disinformation

Once a person has been exposed to, and accepted as true, an inaccurate piece of information, how can their belief be corrected? What do we not yet understand about the possibilities and pitfalls of correcting misinformation?

FACT CHECKING

Fact checking may currently be the most popular tool used to correct false beliefs and stop the spread of misinformation. Available evidence regarding the effectiveness of fact checking is mixed. Some studies identify positive effects, but it is clear that fact checking as an activity faces substantial hurdles. Two key factors inhibit fact checking success: 1) the volume of information to be checked is overwhelmingly large, and 2) the audience can be resistant to the fact checker's message for a number of reasons.

First, the volume of information is an important hurdle to fact checking, and research is needed on both automated and human-led methods that can flag and check individual stories. Because the supply of individual stories is so large, some fact-checking solutions categorize entire websites as likely purveyors of misinformation. However, even sites with high rates of false information can have a significant proportion of content drawn straight from trusted sources like AP or Reuters, making such categorizations problematic. One intermediate-level suggestion to solve the problem is to create a Wikipedia-style database of fake news and fact checkers. The individual units with entries in this wiki could be websites and news organizations, and the score of each listed unit could continuously be updated based on recent flows of information and spot checks.

The problem of voluminous material in need of checking is further compounded by the technological ease with which numerous websites can be set up in order to spread large volumes of misinformation. In addition to developing technological solutions that enable more efficient fact checking, research should also

study nefarious uses of publishing tools such as WordPress and Amazon Web Services, in order to better understand the supply side of false information.

The second concern with fact checking is that appeals to independent authority do not always work. In addition to important research agendas on source credibility and the roles of identity and emotion in this process (discussed below), research questions about fact checking effectiveness could include: When is fact checking successful, and how long do the effects last? Is fact checking more successful when it occurs closer in time to the original exposure to false information? What role do civic education and media literacy play in acceptance of fact checking?

SOURCE CREDIBILITY

One key reason that corrections of misinformation fail is a lack of source credibility, which is an important quality not only for fact checkers but for all communicators, including politicians and news organizations. While decreasing and politically polarized trust in news organizations is posing problems for their credibility in the eyes of Americans, recent evidence suggests that the extent of polarization in source credibility is sometimes overstated. While trust in Fox News is heavily politically polarized among the American public, there is more cross-partisan agreement on the trustworthiness of other media outlets; for example, more conservatives trust The New York Times than trust Breitbart News. Even taking this factor into account, however, future research on the determinants and consequences of source credibility would be useful for advancing our knowledge of how to combat false information.

Research to date suggests that one key source of credibility is being seen as nonpartisan; however, this is a difficult task in a hyper-partisan political environment. Another way to gain (topic-specific) source credibility is to speak out against the party line and thus one's own self-interest (e.g., Republicans debunking the "death panels" rumor that surrounded the Affordable Care Act). Both of these are likely to be effective strategies, but they are difficult to promote in a politically polarized situation. Research on how to encourage speaking out against an ideologically congruent position would be welcome. As an example, such research could build on prior evidence that media competition among newspapers on the same side of the political spectrum can encourage critical coverage of copartisans.

THE MOTIVATIONS OF INFORMATION CONSUMERS

Most research on misinformation assumes, implicitly or explicitly, that people consume news in order to find truth. However, this assumption may not be accurate. What would research on misinformation look like if it did not assume that people are motivated by truth? Research that relaxes this assumption could forge new ground, for example, by exploring the impact of different user interfaces on media consumption. What role do emotion-based reaction and sharing on Facebook play in shaping media

diets? If a person is online for purposes outside of finding accurate information, then their ability to discern when information is false may not be relevant to their choice to consume the information. If that is the case, how could user interfaces be designed to minimize the spread of misinformation?

More generally, how strong is the demand for accuracy and truth? Under what conditions do people demand accuracy? One area in which people may be more likely to search for accuracy is in policy-relevant but not highly politicized facts. Topics such as the details of welfare regulations or who owns US foreign debt are important for citizens who seek to make sense of politics, but these topics are not as politicized as some others, like climate change. Research on the dynamics of correcting misinformation in these spaces can help us understand the American political information landscape more fully.

ADDITIONAL AREAS OF POTENTIAL STUDY

Other potential areas of research include: studying the dynamics of extremely local cases of mis- and disinformation (as distinct from national or international cases); exploring the frequency with which news organizations rely on social media content as sources for television, web, or newspaper stories, including the consequences this may have on media “laundering” of disinformation; and studying the costs (social, mental, physical, and financial) that accompany the increasing professional requirement for journalists to operate in social media spheres.

RELATED READINGS

Albright, Jonathan. “Welcome to the Era of Fake News.” *Media and Communication* 5, no. 2 (2017): 87–89. <https://doi.org/10.17645/mac.v5i2.977>.

Berinsky, Adam. “Rumors and Health Care Reform: Experiments in Political Misinformation.” *British Journal of Political Science* 47, no. 2 (2015): 241–262. <https://doi:10.1017/S0007123415000186>.

Gentzkow, Matthew, Jesse M. Shapiro, and Michael Sinkinson. “Competition and Ideological Diversity: Historical Evidence from US Newspapers.” *American Economic Review* 104, no. 10 (2014): 3073–3114. <https://doi.org/10.1257/aer.104.10.3073>.

Lazer, David M. J., Matthew A Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, et al. “The Science of Fake News.” *Science* 359, no. 6380 (2018):1094–1096. <https://doi.org/10.1126/science.aao2998>.

Thorson, Emily. “Contextual Fact-Checking: A New Approach to Correcting Misperceptions and Maintaining Trust.” *Knight Foundation White Paper* (2018). <http://kng.ht/thorson>.

Pennycook, Gordon, and David G. Rand. "The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Stories Increases Perceived Accuracy of Stories Without Warnings." *SSRN Working Paper* (2017). <https://ssrn.com/abstract=3035384>.

Pennycook, Gordon, Tyrone Cannon, and David G. Rand. "Prior Exposure Increases Perceived Accuracy of Fake News." *Journal of Experimental Psychology: General* (in press). <https://ssrn.com/abstract=2958246>.



Jonathan Albright, Director of Research, Tow Center for Digital Journalism, Columbia University

Homophily in the Social Media Sphere

Many are concerned that online political discourse is taking place in self-selected or algorithmically supported “information bubbles,” though evidence diverges on how serious this problem is. Do we know whether social media is exceptional in enabling or limiting exposure to politically heterogeneous information?

HOMOPHILY IN SOCIAL GROUPS

Homophily in social groups comes about as a result of self-selection as well as social influence. We know that people tend to self-segregate into homogenous groups and that they are subsequently also influenced by these groups to become even more similar to other members. As mentioned earlier in this report, the relative prevalence of political homophily in online vs. offline social networks is a question under active research, but it is highly likely that both selection and influence are at work in both online and offline communities.

Further research is needed on online interventions related to homophily. The ultimate aim of interventions should be to reduce polarization and increase (well-placed) trust; whether this will require discouraging the development of homogenous networks is an open empirical question. Promising research agendas on this topic might address: developing finer-grained definitions of what content counts as cross-cutting vs. aligned with partisan preferences; improving our understanding of which topics citizens consider political; improving our ability to track media consumption across platforms; and finding ways to avoid the Hawthorne effect when observing subjects’ media consumption and sharing patterns.

SELECTION EFFECTS

People’s networks tend toward homogenous political affiliations, such as among circles of friends on social media. One recent study finds that selection into communities is a stronger explanatory variable for political homogeneity than the subsequent effect of being exposed to one-sided information. In addition, the powerful

influence of partisan selection appears to be increasing over time: Today, people are a lot more likely to marry a copartisan than they were fifty years ago. This trend toward higher rates of selection into politically homogenous couples is lower for selection on nonpolitical traits, suggesting that increased selectivity is particularly pronounced in the political domain.

Outstanding research questions on the consequences of selection into homogenous groups include: Do people vary in their propensity to sort by political or other traits? By what mechanism does online sorting occur? What are the different roles played by actions like (de)friending, muting, and blocking? How does political sorting affect the broad set of users who are not interested in politics?

INFLUENCE EFFECTS

Our social networks also create influence effects: We become more similar to those we spend time with (online and offline). Recent evidence shows, for example, that interpretations of politically charged factual data regarding climate change are affected by social context. Participants in one study improved their interpretation of such data when they were placed in a group of anonymous others, mutually exchanged interpretations of the data, and were then able to update their own interpretations. When the partisanship of others was known, social learning still took place, but partisan bias was present in final interpretations. Results like these suggest that social media and other platforms can be designed to encourage forms of interaction that enable civil learning and exchange without exacerbating political polarization.

Future research concerning influence effects on social media could address questions such as: Does social media enhance majority illusions (the perception that most people agree with you), and with what consequences? How do influence effects vary across strong and weak ties? What, if any, effect does exposure without engagement have (e.g., simply scrolling past something on one's feed)? What happens when people encounter political information incidentally in an otherwise nonpolitical feed? How can we measure the contextually revealed motivations behind different kinds of social media engagement (for example, distinguishing supportive from sarcastic retweets of partisan news sources)?

RELATED READINGS

Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. "Exposure to Ideologically Diverse News and Opinion on Facebook." *Science* 348, no. 6239 (June 5, 2015): 1130–32. <https://doi.org/10.1126/science.aaa1160>.

Centola, Damon. *How Behavior Spreads: The Science of Complex Contagions*. Princeton: Princeton University Press, 2018.

Iyengar, Shanto, Tobias Konitzer, and Kent Tedin. "The Home as a Political Fortress; Family Agreement in an Era of Polarization." Working Paper (2017). <https://pcl.stanford.edu/research/2017/iyengar-home-political-fortress.pdf>.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. "The Origins and Consequences of Affective Polarization in the United States." Working Paper (2018). <https://www.dartmouth.edu/~seanjwestwood/papers/ARPS.pdf>.

Settle, Jaime. *Frenemies: How Social Media Polarizes America*. New York: Cambridge University Press, 2018.

Wojcieszak, Magdalena E., and Diana C. Mutz. "Online Groups and Political Discourse: Do Online Discussion Spaces Facilitate Exposure to Political Disagreement?" *Journal of Communication* 59, no. 1 (2009): 40–56. <https://doi.org/10.1111/j.1460-2466.2008.01403.x>.



Jaime Settle, Assistant Professor of Government, College of William and Mary

Globalization of the Marketplace of Ideas

The use of social media for political ends is not limited to the United States, nor to traditional state actors. What do we need to learn about the use of social media to push the political interests of foreign, state, or nonstate actors?

POLITICAL USES OF SOCIAL MEDIA

States' attempts to meddle in one another's internal affairs are not new; rather, it is the tool of social media that is novel. This technology has given both state and nonstate actors a new way to achieve their pre-existing political goals, such as sowing domestic discord among their adversaries or preventing uprisings domestically. Countries like China and Russia think of internet security as being synonymous with national security; in the eyes of these governments, the internet is a tool that can be used to exert influence over the behaviors and opinions of people both at home and abroad.

Research has shed light on some of these online activities; for example, China uses online censorship to remove posts that have the potential to go viral and that refer to political mobilization. The Chinese government also fabricates social media posts—known as fifty-cent party posts—as if they were the opinions of ordinary people. In these posts, the Chinese regime's strategy is to avoid arguing with skeptics of the party and the government and discussing controversial issues. The goal of this massive secretive operation is to distract the public and change the subject, as most of these posts involve motivational, positive, cheerleading content. While we know that China engages in this type of activity domestically, we know less about its existing or potential international operations; this would be an excellent area for research if the significant difficulties with access to empirics and data can be overcome.

In the case of Russia, we know more about its internationally oriented online activities. Russia frequently uses disinformation

tactics on social media; these are a continuation of its previous destabilization activities. Previous Russian tools employed for the same ends have included clandestine support for in-person protests, propping up “alternative” publications, disseminating talking points to homegrown media outlets, and providing fake experts. Many of these tactics were used in Eastern Europe before being deployed against the United States more recently.

RESPONSES AND OUTSTANDING QUESTIONS

When it comes to Russian interference in the 2016 United States election, studying whether and how individual-level opinion changed as a result of Russia’s online activities is exceedingly complicated. However, it is possible to trace how our national discourse changed in the aftermath of the election and the revelations about Russia’s involvement. With regard to the latter, Russia has been successful in its aim to increase its own standing in the political imagination. Considering this, how can liberal democracies, including the United States, respond to challenges of this type (originating in Russia or elsewhere)?

One option is for states to enact restrictions on speech; as an example, consider the Network Enforcement Act recently passed in Germany. Enforcing speech controls is difficult, as monitoring and moderating online content involves substantial resource consumption and is inherently complicated, raising the question of who should do the moderation. If platforms are tasked with enforcement (and are at significant threat of penalty from the state if they underenforce), then they may reasonably err on the side of caution. However, this could result in (overly) aggressive deletion of posts and raise important issues regarding freedom of speech.

Outstanding research questions in this area include: How and under what conditions do politically or commercially motivated social media campaigns change offline behavior? When it comes to protections against hostile cross-border campaigns, how do we think about jurisdiction? What does it mean to put legal restrictions on what companies like Facebook can show their users, and what happens when these restrictions vary by country? What are/should be the boundaries of legal speech, particularly in private, corporate spaces? What can we expect to happen when Chinese platforms and content providers expand their presence globally? How do we improve access to data and empirics from nondemocracies? What security measures should be put in place to protect potential targets, including informative institutions (such as Wikipedia)?

RELATED READINGS

Applebaum, Anne, Peter Pomerantsev, Melanie Smith, Chloe Colliver. "Make Germany Great Again: Kremlin, Alt-Right and International Influences in the 2017 German Elections." *Institute for Strategic Dialogue Report* (2017). <http://www.isdglobal.org/wp-content/uploads/2017/12/Make-Germany-Great-Again-ENG-081217.pdf>.

Jankowicz, Nina. "The Disinformation Vaccination." *The Wilson Quarterly* (Winter 2018). <https://www.wilsonquarterly.com/quarterly/the-disinformation-age/the-disinformation-vaccination/>.

King, Gary, Jennifer Pan, and Margaret E. Roberts. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107, no. 2 (2013): 1–18.

King, Gary, Jennifer Pan, and Margaret E. Roberts. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, not Engaged Argument." *American Political Science Review* 111, no. 3 (2017): 484–501.

King, Gary, Jennifer Pan, and Margaret E. Roberts. "Reverse-engineering Censorship in China: Randomized Experimentation and Participant Observation." *Science* 345, no. 6199 (2014): 1–10.

Kinstler, Linda. "Can Germany Fix Facebook?" *The Atlantic*, November 02, 2017. <https://www.theatlantic.com/international/archive/2017/11/germany-facebook/543258/>.

Maréchal, Nathalie. "Networked Authoritarianism and the Geopolitics of Information: Understanding Russian Internet Policy." *Media and Communication* 5, no. 1 (March 22, 2017): 29. <https://doi.org/10.17645/mac.v5i1.808>.

