13

Conclusion: The Challenges and Opportunities for Social Media Research

Nathaniel Persily and Joshua A. Tucker

We began this volume by noting the relationship between basic research on the impact of social media and politics and the open policy questions concerning social media regulation. The preceding chapters have demonstrated that scholars have learned a great deal in a relatively short period of time about social media's impact on political communication and elite and mass political behavior. It is equally clear that many, many important questions remain to be answered in the coming years.[1] Moreover, as we hope this volume has demonstrated, the answers to these questions are desperately needed to inform major decisions regarding public policy around the world. Responding to an environment of panic surrounding social media's impact on democracy that may very well be amplified by the Covid-19 pandemic, regulators and other political actors are rushing to fill the policy void with proposals based on anecdote and folk wisdom emerging from whatever is the most recent scandal. The need for real-time production of rigorous, policy-relevant scientific research on the effects of new technology on political communication has never been more urgent.

In this concluding chapter, we turn to a different aspect of the link between research and social media policy: the need for new policies to guarantee the continued production of high-quality research to ensure society has answers to the many crucial questions concerning the relationship between social media and democracy. We begin by laying out the many serious challenges facing the field in terms of access to the necessary data to conduct this research, including political, legal, and logistical factors. We then provide what we hope will serve as a set of key principles that can underline arguments regarding the importance of data access for research and a framework for thinking about how such access

[1] Indeed, Kevin Munger has argued that, because of the speed at which the underlying architecture of major social media platforms and networks change, we should – in addition to tackling new unanswered questions – be revisiting questions that we think have already been answered in order to ensure the *temporal validity* of the original findings (Munger 2019).

can be provided moving forward. We close with an assessment of where we are now and some recommendations for how to get to where we need to be.

CHALLENGES TO RESEARCH ON SOCIAL MEDIA AND DEMOCRACY

To some extent, it has been the best of times and the worst of times when it comes to social media research. As the first half of this book reveals, we are beginning to gain important insights into the dynamics of the communication revolution underway. However, despite these achievements and the widely recognized importance of this research, unique constraints have hindered the necessary concerted academic effort to answer the most important empirical questions. The key social media datasets to answer these important questions are not as readily available as were politically relevant datasets of years past. Moreover, unique legal barriers prevent analysis of such data, and related ethical and privacy concerns have arisen that have chilled academic inquiry.

First, the difficulties in obtaining access to the relevant data cannot be overstated. Unlike most politically relevant datasets, the data necessary for social media research are largely controlled and "owned" by private companies. Whereas most political science data analysis, until recently, has utilized administrative data produced by the public sector, such as election returns and census data, or data produced by researchers themselves, such as surveys or experiments, a large portion of the data necessary to investigate the Internet's effect on democracy and elections is locked inside firms, such as Facebook and Google. Although different platforms have exerted different levels of effort to make data available for outside research, it remains the case that making data accessible for outside research has not been – and is highly unlikely to be in the future – part of the core mission of these companies. Indeed, it can often get in the way of a platform's profit-making mission, especially (as has often been the case of late) if outside researchers discover problems with the product or identify potential damage it may cause to society.

As a result, the research agenda for studying the effect of social media on democracy – as well as the scientific insights produced from such research – run the risk of being biased by the kind of data platforms make available to researchers. For example, the vast majority of the research studies on which we report in this volume are analyses of Twitter data; clearly, this is not because there is a consensus that Twitter is the most politically consequential social media platform. Although Twitter is certainly important for politics in many countries, this imbalance in research occurs because Twitter data have historically been among the most easily accessible for outside research, especially compared to Facebook data.

Twitter pays a price for this openness, however. Officials at Twitter are quick to recount how journalists and scholars can paint a misleading picture of what happens on the platform by merely reporting the volume of problematic content without giving context as to the share of user Twitter feeds in which the content appears. However, moving beyond counts of phenomena to more valuable measures, such as ratios (i.e., measuring the denominator in addition to the numerator), exposure, and longitudinal trends, requires time- and cost-intensive data-gathering strategies. These are often hampered by the platform's own terms of service, to say nothing of policies that remove the activities of malicious actors (such as foreign influence campaigns) or whole classes of data (such as exposure or recommendation) from outside access.[2] The most influential analyses of Facebook data, by contrast, have largely been written by Facebook researchers themselves or by academics working through special arrangements with Facebook, often requiring prepublication approval from the company.

In the summer of 2018, Facebook initiated a data-sharing initiative with Social Science One, an academic effort seeking to make Facebook and other industry data available to the larger scientific community (King and Persily 2019).[3] Through the Commission created by Social Science One, which was funded not by Facebook but by a set of nonprofit foundations, academics have gained access to some Facebook datasets. Working with Facebook, Social Science One issued a Request for Proposals (RFP) for analysis of a specific dataset, beginning in July 2018. Scholars who seek access to the data must have projects approved by their universities' Institutional Review Board and undergo a separate evaluation for legal and ethical compliance; their universities must then sign a research data agreement that ensures protection against unauthorized disclosure. It took almost twenty months for Facebook to produce a dataset similar to the one promised in the original RFP – a dataset of almost 38 million links comprised of several trillion cell entries delineating the numbers and types of people who saw and engaged with the URLs, but with statistical noise added to the data to protect users' privacy. The dataset also contained information about the URL, such as whether it was fact-checked or labeled hate speech, but was limited only to URLs that were shared publicly approximately 100 times or more (a serious limitation to the data). The fact that it took ten times longer than expected for anything close to the original dataset to become available for analysis illustrates the fundamental challenges facing large internet companies in finding a way to make their data available for outside research. Indeed, the now-released URLs dataset was supposed to be the less interesting, "easy" dataset for Social Science One, as it did not contain data at the individual level. Even this "easy" limited dataset turned out to be incredibly challenging to produce, however, given the privacy-related

[2] Twitter does deserve praise for its recent efforts to make data produced by foreign influence campaigns available for scholarly research after having removed these posts from the platform; see "Information Operations" in the Twitter *Transparency Report*. https://transparency.twitter.com/en/information-operations.html.

[3] Disclosure: Social Science One is cochaired by Gary King and one of us (Persily); the other one of us (Tucker) chairs an advisory committee on disinformation and election integrity, and several of the authors appearing in this volume are also involved with the effort.

challenges and data infrastructure requirements for such a broad and inclusive research effort.

Moreover, Facebook's cooperation with Social Science One stands as an exception to the general rule of tech corporations denying access to user data. In general, these companies view themselves as having much more to lose from sharing data than they have to gain. On top of potential reputational and financial damage caused by findings that put the firm in a bad light, the firms are reluctant to risk the possibility of unauthorized disclosure and breaches of privacy. Today, academic requests for access to these kinds of data are seen in the light of the now-infamous Cambridge Analytica scandal. As is now well known, in 2014 a researcher at Cambridge University, acting in his personal capacity, placed a psychological questionnaire on Facebook's platform. Users who took the survey consented to deliver data about their profile and activity on Facebook and that of their friends (who never consented to the survey). That researcher transferred the data to Cambridge Analytica, a political consulting firm that was working with, among many others, the campaign for then-candidate Donald Trump. As a result, some data of at least 50 million Facebook users were delivered to a political consulting company that said it had developed and employed new methods of psychographic profiling that could be used for political advertising and other forms of campaign targeting.

Cambridge Analytica was a political and corporate scandal, but it was an academic scandal as well. An academic misused the Facebook data of tens of millions of Facebook users. (Whether the researcher actually violated Facebook's terms of service remains an object of debate, but, regardless, the violation of unconsenting friends' privacy highlighted a problem with making the Facebook social graph data accessible under those circumstances.) Even if he was acting in his personal, rather than academic, capacity, his misdeeds have had a chilling effect on academic (and other) access to the critical stores of data social media firms possess on politically relevant questions. Often invoked in policy discussions, the scandal has lived on beyond its particular facts to embody a larger controversy regarding firms' misuse of users' private data, which can be exploited for political, economic, or other purposes. In the wake of Cambridge Analytica and other data privacy scandals, the platforms have reevaluated their data accessibility protocols for researchers and all other users and in the process reduced or shut down altogether preexisting access to APIs.[4]

[4] APIs, or Application Programming Interfaces, are tools by which data produced online can be directly delivered to external (or, for that matter, internal) researchers without, for example, having to actually render a web page. They offer much greater speed than the alternative means of collecting social media data – which is scraping the web page directly – and provide the platforms with a degree of control over what can and cannot be downloaded, as well as how much data can be collected within a given time period. Of course, platforms can also choose to charge for access to APIs, as well as to make scraping information from web pages directly against its terms of service.

In the midst of all this, regulators around the world have, predictably, flexed their muscles to constrain the platforms' ability to make private data accessible to anyone outside the firm and, in some cases, to prevent collection of certain data by the firm itself. Since 2011, Facebook had been under a consent decree with the US Federal Trade Commission (FTC). That decree, which arose out of Facebook's failure to comply with its articulated privacy policies, constrains all kinds of potential data access for academics and others. It also places Facebook under intense and continuous oversight by a federal agency. Based on its perceived breach of the consent decree in the Cambridge Analytica scandal, the FTC entered into a new settlement and decree with Facebook, which involved a $5 billion fine and additional future oversight of Facebook procedures (FTC 2019).[5]

The most influential law governing researchers' (and any outsiders') access to social media data, however, is the European Union's General Data Protection Regulation (GDPR). That law has presented new and somewhat indiscernible constraints on the use and release of social media data for research. GDPR contains exceptions for certain types of research (Article 89), but both the rule itself and, in particular, the member states' divergent implementation of it have been less than clear as to the boundaries of this exception. Confusion revolves around the required procedures to minimize risks to privacy, as well as the degree of anonymization (or pseudonymization) required for social media data. Although individual names and identifying information can be removed from social media datasets, doing so might not be sufficient to pass the legal bar of anonymization given the richness of those datasets and the outside chance that researchers might theoretically be able to reidentify people if they were committed to combining multiple datasets from other sources.

To be sure, arguments about GDPR and other privacy laws sometimes appear as pretexts from one or another platform to justify restricting outside access to data. However, in this environment of legal uncertainty, the platforms have taken (understandably, from their standpoint, but frustrating from the research community's perspective) a very restrictive view in terms of sharing individual-level data, often regardless of whether the data are anonymized or could lead to reidentification. Moreover, Facebook has said that it will apply GDPR worldwide, in part because many other governments have either passed or are considering similar legislation. To address widespread misperceptions about the chilling effect of GDPR on research, the European Data Protection Supervisor issued in January 2020 "A Preliminary Opinion on data protection and scientific research" (European Data Protection Supervisor 2020). The guidance also attempts to reinforce the message of the importance of research and its consistency with the goals of GDPR. Nevertheless, even given this further guidance, lawyers at the platforms remain very conservative in their

[5] One potentially valuable use for a portion of this and similar fines would be to support independent research on the impact of the platforms on society.

interpretation of when, what, and how data can be made accessible to researchers.

Given this impasse, the research community needs more than a mere clarification of GDPR. It needs a clearly defined safe harbor or a research pathway sanctioned by the European Commission. This could involve the designation of secure facilities and computers to analyze data, government vetting of researchers requesting data access, surveillance and recording of researchers while they analyze data, auditing of research results to ensure privacy is protected, pre-review of publications to ensure no privacy leakage, and significant penalties for any researcher who seeks to reidentify individuals in the dataset provided. (Such procedures, while seemingly draconian, are already in place for other sensitive datasets involving census, tax, or health data.) What would make the harbor safe, however, is that, in exchange for delivering data under these conditions, the platforms would receive complete legal immunity for granting data access to researchers. Indeed, given the value of research access to society for a whole host of reasons, not only should the platforms be immunized when they are willing to provide data; they should be legally compelled to do so. Governments need to spell out the legally safe pathway for granting researcher access, and then they need to require that the platforms follow it. Only then will lawyers inside the platforms recalibrate their legal risk assessments so as not to overcorrect on GDPR. Indeed, the Kofi Annan Commission on Elections and Democracy in the Digital Age (2020) recommended this specific proposal of legally compelled researcher access to platforms' privacy-protected data in a recent report on how to address the challenges that new technologies pose for democracy. (Disclosure: Persily was a commissioner on the Kofi Annan Commission that issued the report.)

The privacy-related obstacles to research access are not limited to those legislated by governments, however. In the wake of Cambridge Analytica, other privacy scandals, and governments' regulatory responses, a powerful privacy movement has arisen in civil society. The privacy policies of the platforms themselves, as well as surveillance by governments, are the main targets of this movement. The movement is both necessary and salutary given the real dangers to privacy that the evolving digital environment portends. Academic research, however, has become collateral damage in this battle between privacy advocates and the platforms. At one end of the divide are those who argue that individuals who provide data to social media platforms do not do so with the intent that these data will be used for purposes beyond simply sharing their posts with the intended audiences, and therefore outside academic research should not be permitted. These "other purposes," of course, include the bread and butter of social media platforms' business models – targeting ads – but also potential uses of digital trace data for social good, including but not limited to scholarly research in the public domain. From this perspective, academics should be permitted to analyze only data that the user has expressly made public or has been specifically designated for academic analysis (e.g., through a survey instrument designed to gain consent for research).

At the other end of the spectrum are those who consider social media data as akin to "administrative data." Administrative data are generally defined in the research community as data that are produced for one reason (e.g., giving students grades in a class to illustrate their competency in a subject area) but can be analyzed for another purpose (e.g., to discern the most efficient way to spend tax dollars on education). Political scientists have long analyzed administrative data at both the aggregate level (e.g., election results, protest participation, unemployment rates) and the micro level (e.g., voter registration data, census data) without assuming that the analysis of such data requires individuals to have consented to be considered as subjects in a research study. For administrative data, therefore, it is unnecessary for those individuals to provide explicit consent in order for the data to be analyzed. Requiring explicit consent for research on administrative data would prohibit, for example, any study of election results or employment rates.

Of course, to describe social media data as administrative data has the potential to diminish its sensitive personal nature. Although some social media data seem "administrative" (e.g., number of friends, popularity of URLs, whether one has posted using a mobile or desktop device, time zone, etc.), other data appear qualitatively different in that they do not contain records of behavior so much as the personal thoughts and observations of individuals. However, some forms of administrative data, such as those involved with medical treatment and health statistics, are equally sensitive – perhaps even more so – than social media data. Yet, we may be moving toward a data access regime in which personal medical data may be more accessible to researchers than what users share or read on Facebook.

A privacy paradigm that requires explicit user consent for social media research will prevent scholars from answering some of the most important questions surrounding social media's impact on democracy. The most basic questions, such as how much disinformation the average user sees in a newsfeed or which categories of users see questionable or polarized sources, will not be answerable from data available only from a statistically biased set of users willing to provide it. For example, as recounted earlier in this volume, researchers have begun to find evidence that consumption and forwarding of disinformation, at least in the United States, is concentrated among older people. To be confident in this conclusion, though, researchers must be able to analyze – in the aggregate – exposure to different media sources conditional on the age of the user. However, many people do not publicly reveal their age on Facebook or other platforms. Although preventing researchers from deliberately uncovering a particular individual's age without user consent seems a perfectly reasonable barrier to protect privacy, preventing aggregate analysis of different age cohorts of users on the same basis would necessarily prevent us from understanding how social media exposure varies based on age.

A research paradigm based on consent as the touchstone would prevent both kinds of inquiries.

THE NEED FOR A NEW DATA-SHARING PARADIGM

As we think through possible data-sharing paradigms, it is important to begin with an understanding of the fact that prohibiting the sharing of social media data for analysis by the scholarly community – or any researchers who are committed to sharing findings in the public domain – does *not* mean that social media data are not being mined for insights. Rather, it means that *employees of the platforms will be the only ones* analyzing the data and learning the answers to the most pressing questions as to social media's impact on democracy and other social phenomena.[6] Insights and expertise will therefore flow solely to a small number of very large (and politically influential) corporations, which can then pick and choose on their own which questions to ask and what conclusions to share with the public at large. Recognizing this inconvenient truth, the question as to research access and privacy is not whether user data should be analyzed for insights, but whether the platforms should have a monopoly on such access or inquiry.

Commentators often describe Google and Facebook as information monopolies. Usually, this accusation provides fodder for arguments about antitrust and competition law – such as whether the companies should be broken up into their constituent parts or regulated as public utilities (Stigler Center 2019). However, they are also information monopolies in a more literal sense – the firms control the information necessary to understand basic facts about contemporary society. As dangerous as these information monopolies may be for purposes of economic competition, such dangers are compounded when only those who work for the firms and share in their corporate missions are able to gain social insights from the data they possess. Social media companies control both the information valuable to their competitors and the personal data valuable to their users. More importantly for academics and those hoping to use rigorous scientific research to inform policymaking, though, the platforms control the information that most richly describes politics and society and therefore the data necessary to make sound judgments across virtually all major policy domains.

[6] Moreover, such questions will be studied only if the platforms choose to devote corporate resources to trying to answer these sorts of questions in the first place. In most cases, these data will not be analyzed to answer questions to advance scientific knowledge but rather to bolster efforts to maximize profits. To be clear, we do not and should not expect the platforms to substitute a public mission for their economic interests. Instead, we seek to point out that the use of these data for societal good – such as understanding the impact of social media on democracy, the goal of this volume – by necessity requires those outside the platforms to have access to these data.

A recent decision of the US Court of Appeals for the Ninth Circuit appears to recognize the pervasive impact of platform control of information. In *hiQ v. LinkedIn*, No. 17–16783 (9th Cir. 2019), the Court protected a company's right to scrape user-provided data on LinkedIn. As the Court explained, "giving companies like LinkedIn free rein to decide, on any basis, who can collect and use data – data that the companies do not own, that they otherwise make publicly available to viewers, and that the companies themselves collect and use – risks the possible creation of information monopolies that would disserve the public interest." Admittedly, the Court's decision arose in the special context of the Computer Fraud and Abuse Act's potential prevention of scraping of publicly available data by a private company, but the parties in the case, as well as the opinion itself, noted that academic researchers often must resort to scraping to get the information under platform control. The decision echoes arguments that academic researchers have themselves made about the need, and perhaps even the right, to scrape social media data, when doing so is in the public interest but against the terms of service of a given platform (see Freelon 2020).

The platforms' monopoly control over relevant social scientific data, which admittedly derives from their users' private communication and behavior, requires a reframing of the debate around access to social media data.[7] We need to move beyond the normatively pleasing paradigm of "should the platforms respect the privacy concerns of their users?" – with which, of course, everyone agrees in the abstract – to one that fully embraces the trade-offs inherent in making data accessible to outside researchers. Such a framework might be oriented around several key principles:

1. Social media platforms' business models are entirely dependent on insights gained from analyzing data provided by their users;
2. There are legitimate privacy (and legal) concerns when the platforms grant access to social media data to third parties for research purposes; *but*
3. There are real differences between private actors who analyze these data in order to support for-profit businesses with no obligation to release findings to the public (and indeed may even have obligations to shareholders not to do so) and other actors in society whose goals are to analyze these data in order to advance scientific knowledge and share their findings in the public domain, or to build tools for (nonprofit) social good;
4. There are real gains – economic, political, and social – that can result from the public sharing of insights from analyzing social media data. These benefits run the gamut from medical discoveries to disaster prevention to identifying and preventing foreign interference with elections. There are

[7] Indeed, a recent issue of *The Economist* features an entire special report on the new data-driven economy, with a section of the report devoted to questions around data access (see The Economist 2020).

> also dangers when public policy is made without the advantage of the insights that can be gained through analysis of social media data.

Thus, the question of whether social media data ought to be shared more or less widely than they currently are is not merely a question of how platforms can better respect the privacy concerns of their users. Rather, policymakers and advocates need to consider the *trade-offs* between a world in which *data are shared less frequently but gains from analysis accrue only to large for-profit companies* and a world in which *data are shared more frequently but gains from analysis can accrue to the public at large*. Under the former, privacy can (usually) be better protected but net social gains are likely to be smaller; under the latter, privacy would be more at risk without the appropriate safeguards but the opportunities for social gain are larger as well.

One could go even farther than acknowledging the trade-off between privacy concerns and the benefits accrued by research in the public domain to raise the question of whether it is even appropriate to think of social media platforms "owning" the data provided by users of the platform, with a concomitant right to be the only entity allowed to accrue knowledge from analysis of the user-provided data. Such an argument could start from acknowledging the role that the major digital platforms play in contemporary society: They are not merely places people visit online but rather central nodes of our social, economic, and political lives.[8] Thus, social media – much like jobs reports, election results, or even census data – are a crucial component of our understanding of contemporary social, economic, and political systems. However, unlike traditional sources of administrative data, digital platforms are distinguished by the fact that they are both wholly privately owned *and* highly concentrated. Thus, society has a special claim on these data precisely because (1) they will not otherwise be made available in the public domain (the way, for example, census or election data will be) and (2) these companies provide not a single service (like airline or automotive companies) but rather a collection of services that inextricably link these platforms to society's social, economic, and political life. Taken together, these arguments suggest that, although such data are currently considered a private good owned by the corporations, they ought to be considered a public good.

Therefore, we may need to begin thinking about updating our concept of the public's right to data in the context of these information monopolies. This right should supersede the proprietary right of companies to enjoy exclusive access to the digital trace data created by users of their products at some point when those data are necessary for a complete understanding of the basic social, economic, and political functions of society. Of course, it would be folly to suggest that the builders of the platforms would not have access to the data they are collecting,

[8] We thank Sam Gill of the Knight Foundation for suggesting this line of reasoning in response to having read an earlier draft of this chapter.

but ought they be the only ones allowed to do so? Perhaps these firms are better conceptualized as "data stewards" (or "information fiduciaries" to use the term coined by Yale Law Professor Jack Balkin; see Balkin 2016) entrusted with managing the data that they have acquired for the good of their users *and* society at large, in addition to their shareholders. From this perspective, providing access to data for public-facing research would be considered an obligation for data stewards above a particular size.

We can consider two even more radical arguments in this regard. One would be to shift the policy focus from who "owns" the data – for example, as long as Facebook or Google "owns" their data, they can do what they want with it but cannot be permitted to transfer the data to other actors – toward a "use" approach to privacy rights.[9] For example, companies that collect data could be entitled to use the digital trace data collected from their users' behavior to maximize profits but not to actively support particular political actors in domestic political competitions. Similarly, scholars could be permitted to use these data for the purpose of scientific research but not to reidentify any individual users.

Another idea that involves a more radical rethinking of existing policy is to introduce the concept of a "data tax" on the platforms to be paid back to society. By this we mean not a traditional monetary tax that would be assessed based on how much data a company holds but literally a tax "paid" by the contribution of some proportion of user-provided data into a repository for independent analyses, the results from which would need to be put into the public domain to inform citizens and policymakers alike. Companies can be compelled, of course, to surrender a portion of their earnings to fund the state that provides the infrastructure (e.g., courts, roads, security, etc.) that allow those companies to conduct business. Why not similarly require provision of a portion of their data, as well, to be returned to the public as a way of contributing to the society in which their users reside? Such an argument would potentially be even more compelling if the data were being used in a manner to address potential problems caused by the platforms themselves, such as in the case for research addressing the impact of the platforms on elections and democracy.

We would also suggest that the scholarly community plays an especially important role in leveraging the gains that accrue to the public at large. To be sure, academics and other researchers can engage in malfeasance, conduct unethical research, pursue narrow "academic" questions, or publish erroneous results. However, scholars are incentivized to release the results of their research publicly. Indeed, the currency of academia is *publications*, with an emphasis on the public part of that word: we get paid to publish, and we get rewarded when people reference our research. Although other nonprofit efforts have made great use of social media data in a range of domains including public health, election protection, disaster preparedness, and consumer welfare, their

[9] Again, we thank Sam Gill for suggesting this point.

incentives do not always run in the direction of publication of research results, whatever the conclusions.

Furthermore, there are peculiar features of social media data that make time-consuming and methodologically complex academic research particularly important. Namely, it is exceedingly easy (and often misleading) to find cursory evidence of *anything* on social media because there is so much of it and, by dint of being digital trace data, it is almost by definition optimized for search. Thus, if pundits and journalists want to find evidence of some particular phenomenon, they can probably find a Facebook group, a YouTube channel, or some choice Tweets to illustrate that the phenomenon exists. Even simple counts over time (such as a Google trends timeline) are easy enough to produce. Indeed, a search for any particular type of pathological content can often return millions of results – which appears overwhelming but may be misleading unless one knows the share of total content that is problematic or the likelihood that any given person may have been exposed to it (see, e.g., Siegel, Chapter 4, this volume). The more important questions, such as the relative prevalence of a phenomenon, trends over time, or assessments of causal relationships, however, require much more complex research designs, sustained research efforts, and (often) sophisticated methodological tools. This is where the public-facing research community, and especially academic researchers, have a crucial role to play.

It is worth acknowledging that the social media platforms can, of course, choose to release the results of their own internal research. Indeed, many seminal papers using Facebook data referenced throughout this volume were written or coauthored by internal Facebook researchers. The concern here is that, to the extent that we are reliant on employees of the platforms for research, we have to wrestle with the implications of firms reserving the right to approve papers written by their employees before they are submitted for peer-review. This is a version of what is known in academia as "the file drawer problem," a term used to refer to a concern that positive results are more likely to be published than negative results due to journals' preferences for positive findings, thus presenting a skewed overall view of what we know about a particular topic (Franco, Malhotra, and Simonovits 2014). However, the problem can be even more pernicious when we consider for-profit companies playing the role of gatekeeper, where the assumption would be that research making the company look bad would be more likely to be withheld. To be sure, there are important works that have been published by data scientists working for the platforms that have less than flattering implications for the firms themselves and have contributed significantly to our understanding of how these platforms work. However, if a key goal for the public and policymakers is to learn the impact of a particular new technology/platform/product on relevant outcomes in society (e.g., political polarization), then having access only to published research that has survived vetting from that company is deeply problematic. Imagine letting tobacco companies vet research on the effect of smoking on health, for example. Moreover, we believe that, in the long term, solving this file drawer problem would benefit the platforms as well. Today, if a firm releases studies that show positive societal benefits from usage of its platform, how can policymakers know if there were fifty other studies the company had run that showed the opposite effect? Thus, as long as the company only selectively releases research, the results from any studies released publicly that paint the company in a positive light are likely to be greeted with skepticism. Conversely, if public-facing researchers have access to data to confirm results published by internal researchers, the public is more likely to trust the firm's research, in general.

A closely related question is whether studies *funded* by one of the platforms, but not carried out by researchers who are employees of the platforms, would suffer from the same concerns. Clearly, if funding from a platform came with a right of prepublication approval by the platform (or any funder for that matter), it would raise the same type of file drawer problem. Under such conditions, one might also worry that, even if a funder did not have the right of refusal on research publications, the lure of potential future funding might be sufficient to cause researchers to be selective about what they choose to publish. Further, this problem could be equally serious if the platform was providing access to normally inaccessible data. This is similar to the type of problems that can arise in the context of medical research if researchers tailor their work in an effort to make it more likely that a drug manufacturer will fund further research.

One potential way out of this problem would be to establish a rule or norm prohibiting academic researchers from accepting funding from the platforms. However, there are costs to this approach as well – namely, funding for social media research is quite scarce and the platforms have money. Moreover, it is a legitimate question to ask whether, if research is necessary to understand the impact the platforms are having on the world, the platforms ought to be footing the bill for some of this research. (Indeed, the fact that the platforms now hire away a considerable number of graduate students who otherwise might join the academic ranks also begs the question whether they owe some compensation for the university-paid training provided to their employees.) The trade-off is perhaps even starker when considering the question of data access. When money is the barrier to research, it is always hypothetically possible that someone else could supply funding. If data access is the primary barrier to research, then the cost of not cooperating with a platform in some cases will be that the research simply will not take place or, alternatively, that the research will be carried out by employees of the platform.

However tricky these questions are, some common-sense guidelines seem like they would go a long way in addressing this particular challenge:

1. Academic researchers studying social media or utilizing social media data should not accept funding from the platforms when a condition of the

funding is that the platform has the right to approve papers before they are submitted for publication.[10]

2. Academic researchers who receive funding from social media platforms should be transparent regarding (i.e., disclose) all funding as part of all relevant publications. (This should also apply to any paid consulting relationships.)
3. Academic researchers who are involved in data-sharing partnerships directly with platforms should be transparent regarding (i.e., disclose) these partnerships (and their conditions, if any) as part of all relevant publications.[11]

The question of how scholars navigate the issue of accepting funding from the platforms to conduct social media-oriented research will continue to be an important one because this type of research is expensive. Funding for necessary research comes from a limited range of sources – individuals, foundations, corporations, or governments – each with its own set of risks and limitations. (We should take this occasion to reiterate our thanks to the John S. and James L. Knight Foundation for funding this volume.) While we are extremely encouraged by the provision of nontrivial amounts of funding from foundations in recent years in this space and we remain hopeful that governments around the world will prioritize the study of social media and politics for public funding, the question of the extent to which the platforms themselves fund this type of research is not going to disappear. Ideally, independent institutions dedicated to social media research could form a trust that can serve as a repository for funds from these different actors. Doing so might assist in creating financial distance between any funder (corporate or otherwise) and the researcher to prevent both the reality and the appearance of funder control of the research. Moreover, if a diverse group of corporate, foundation, individual, and government funders participate, no single contributor could be said to be sponsoring the research. One promising way to seed such a trust would be with a portion of the recent $5 billion FTC fine leveled against Facebook.

WHAT THE FUTURE HOLDS

We had hoped to write in this conclusion that the future of social media research is not only certain but bright. Yet we remain quite unsure how the tensions we

[10] To be clear, we are referring to unconstrained rights of refusal to allow publication. We are not suggesting, for example, that platforms that have provided data to researchers under particular conditions (e.g., that publications only release aggregated as opposed to individual-level data) do not have the right to ensure compliance with terms of data-use agreements.

[11] We are grateful to the Knight Foundation and the Social Science Research Council (SSRC) for convening a meeting in the fall of 2018 – in which one of us (Tucker) served as a co-convener and the other (Persily) was a participant – where these questions and potential guidelines were explicitly discussed.

have identified here will be resolved. Indeed, we see cause for optimism and concern.

Optimists might point, for example, to the development of new technologies of differential privacy that might help us out of the privacy versus access trade-off. These new methods, which have met with mixed success as part of the research effort of Social Science One, usually add statistical noise to datasets in such a way that one could prove mathematically that no particular individual in the dataset can be identified. For some, such methods will not resolve the question about consent for research subjects – that is, those who see a personal dignity interest in data might not be assuaged by the analysis of that data, even when users cannot be reidentified. Others will consider these new datasets to be truly "synthetic" – that is, not composed of real data and therefore not actually requiring the consent of anyone for analysis. Yet, if a platform can demonstrate mathematically that there is little risk of reidentification, many of the concerns about user privacy will be alleviated. The US Census Bureau has adopted this approach and has promised to protect user data in the 2020 Census through a system of differential privacy.[12]

At the same time as researchers are developing ways to protect user privacy in social media datasets, the platforms themselves are moving in directions that might make collection of most user data impossible. After Mark Zuckerberg declared in early 2019 that "the future is private," Facebook announced its plan to tie together several of its products (Instagram, Messenger, and WhatsApp) into a suite of encrypted communications. Given concerns about privacy, doing so is understandable and even desirable from the users' standpoint. However, once the platforms move toward widespread encryption, outside analysis will become particularly difficult. As difficult as it is to conduct social media research on platforms already skittish about sharing data, it becomes even more difficult when the platforms do not have access to the communication itself.

At least for now, encryption makes it theoretically impossible for the platforms themselves to see any of the text of communications, let alone for scholars to get access to these data.[13] To be clear, this does not mean that no "data" can be analyzed; it is still possible, for example, to see how often an account is sending out messages and therefore, for example, to make assessments as to whether the account is more likely to be controlled by a human or an algorithm (or even whether a set of accounts appears to be

[12] Many questions remain to be answered as to the appropriateness of existing statistical methods for analyzing differentially private data, as well as how fast new methods for doing can be developed and validated. See, for example, Evans and King (2019), as well as accompanying software available at https://github.com/georgieevans/PrivacyUnbiased.

[13] We say "at least for now" because the rise of quantum computers may challenge conventional understanding of the impenetrability of encrypted data; see Gidney and Ekera (2019).

operated by the same actor). Nevertheless, if one wanted to explore whether actors were spreading misinformation or, potentially more worrisome, calls for violence, the platforms would be unable to do so. Members of an encrypted chat group would still have access to the text of messages, though, raising the temptation for analysts to join groups for research purposes.

Indeed, social media researchers, particularly in the developing world, are already confronting difficult ethical questions as to when and how they can participate and observe communication occurring on encrypted platforms such as WhatsApp. Researchers in Africa, Latin America, and South Asia, where WhatsApp has become a dominant social media platform, are adopting the techniques of anthropologists by embedding themselves in the communities they study. Yet, unlike the anthropologists who became part of the communities they studied, social media researchers can lurk on WhatsApp groups outside of view. Even if they announce themselves when they first join, as ethical guidelines require, groups will change over time and participants are unlikely to be aware that their communications are being surveilled. However, to the extent that invitations to join political WhatsApp (or Telegram, Signal, or other encrypted messaging apps) groups are posted publicly,[14] there is a legitimate question as to whether people who join such groups should have a reasonable expectation of privacy or not. Herein lies the ethical rub: If such groups are having an important political impact – and misinformation spread on WhatsApp groups has been blamed for interethnic violence in many countries (McLaughlin 2018; Arun 2019) – then scholars are going to want to understand that impact and policymakers are going to need the results of such research to inform public policy. Yet, as long as the data remain encrypted, these types of ethically challenging strategies to recover the content of such conversation will be the easiest – and perhaps only – option available.

In addition to the technological challenges and opportunities posed by developments such as differential privacy and encryption, the field will also continue to wrestle with the policy debates surrounding privacy and access. Indeed, we hope that one contribution of this volume is to help us better understand the parameters of the trade-offs between limiting the spread of users' data out of concern for user-privacy versus the potential scientific progress that can be made when digital trace data are made available for scholarly analysis. On the one hand, the preceding chapters have presented a large amount of knowledge that has entered into the public domain due to the fact that scholars have managed – through a variety of suboptimal processes – to secure access to some social media data during the first decade of the Web 2.0 era (as well as to come up with many creative research designs that do not rely

[14] See, e.g., Narayanan and Ananth (2018); PTI (2019).

on access to social media data but speak to the political phenomena related to social media). These lessons, insights, and discoveries are testament to the remarkable potential of social media data to drive the accumulation of knowledge and, in particular, knowledge about the effects of the platforms themselves. At the same time, the volume also highlights the costs of restricting access to data: Time and time again, chapters have referred to what we do not yet know. Of course, there are always remaining questions to be asked in social science research, but it is notable how often the authors in this volume cite limitations in access to social media data as an important cause of these research gaps.

Most "state-of-the-field" edited volumes end with a list of important next steps. These often include research questions and aspirations for new types of data collection. For research on the study of social media and democracy, we find ourselves in a somewhat unusual position. The data we need to conduct our research are plentiful – indeed, the amount of data out there is far beyond our wildest dreams as compared to even a decade ago. Even if one is concerned about the generally observational nature of social media data, the opportunities for experiments abound. Although we cannot prove it – which is sort of the point – we are quite confident that there were more experiments on behavioral outcomes carried out by Facebook and Google this year than the sum total of those carried out by members of the American Political Science Association's Experimental Political Science section.

Yet we also live in a time when a whole host of factors outside of the academy can have huge effects on the degree to which scholars can access these data. These factors include, but are not limited to, policy decisions by government authorities such as the US FTC and the European Data Protection Board and internal business-related choices by platforms to restrict access to APIs. We are truly at the mercy of outside forces that do not often elevate the importance of academic research in their decision-making processes.

Taken together, we would like to suggest, then, that there are essentially three paths to ensuring and expanding the continued production of the type of research featured in this volume. The first is to work *with the platforms*, through efforts like Social Science One, other forms of research partnerships, or by directly lobbying for APIs that are optimized for academic research. The second option is to *work independently of the platforms*, to try to come up with creative ways to collect social media data that are not dependent on platform cooperation; this could include, for example, deprivation studies (where people agree to go off the platforms for a period of time and then agree to be surveyed by researchers), "citizen science" efforts (where citizens are encouraged to download their own data from the platforms and donate them to academic research), or traditional off-platform survey efforts. Finally, researchers can *work with governments* to ensure that data access for outside research is properly valued and considered a crucial component of any attempt to regulate social media platforms.[15] Owing to the many obstacles facing each of these strategies, our best path forward is to pursue all of them at once.

With these goals in mind, we hope that this volume can alert all the relevant private and public players as to the value of research using social media and digital trace data. At the same time, this book represents a clarion call for making social media data available for research, with results concomitantly released in the *public* domain, even while recognizing the importance of users' privacy concerns and the legal and business interests of the firms. If we want the public to know more about hate speech online, the relationship between digital media and political polarization, and the pathways of misinformation through modern communication networks, then we need to ensure that scholars who publish in the public domain have access to the data necessary to carry out these studies. Moreover, if we want policymakers to make informed choices in setting policies regarding digital advertising, regulating new media, and addressing harmful content online, they need to be able to draw on rigorous scientific research conducted with the appropriate data. So much of the research reviewed in this volume concerns topics we might not even have imagined fifteen years ago; it is literally cutting-edge research. Yet it is also research that informs crucial questions of public policy, meaning that the failure to move this research agenda forward will have consequences that reverberate far beyond academia. We hope this concluding chapter, as well as the entire volume, represents a first step on the path toward a future of greater understanding of the challenges social media presents for democracy and, by consequence, a future with better informed policies to address those challenges.

[15] We are encouraged by the fact that some policymakers are beginning to recognize the importance of data access for independent research. Indeed, in Elizabeth Warren's "Fighting Digital Disinformation" plan, she included the following component: "Open up data for research: Research by academics and watchdog organizations has provided the public with important insights into how disinformation spreads online, but these efforts are greatly limited by social media platforms' unwillingness to share data. Platforms like Facebook currently provide only limited and inconsistent access. Research can help evaluate the extent of, and patterns within, disinformation on social media platforms. It can also offer the public an objective evaluation of how the features that platforms offer, including those that allow for rapid dissemination of content, contribute to disinformation. Social media companies must provide an open and consistent application programming interface (API) to researchers" (Warren Democrats 2020).

REFERENCES

Arun, C. (2019). On WhatsApp, rumours, and lynchings. *Economic & Political Weekly*, *54*(6), 30–35.

Balkin, J. (2016). Information fiduciaries and the First Amendment. *U.C. Davis Law Review*, 49, 1183–1284.

The Economist. (2020). Special report: Are data more like oil or sunlight? *The Economist*, February 20. www.economist.com/special-report/2020/02/20/are-data-more-like-oil-or-sunlight

European Data Protection Supervisor. (2020). *A Preliminary Opinion on Data Protection and Scientific Research*. https://edps.europa.eu/sites/edp/files/publication/20-01-06_opinion_research_en.pdf

Evans, G., & King, G. (2019). Statistically valid inferences from differentially private data releases. Working paper. http://j.mp/38NrmRW

FTC (Federal Trade Commission). (2019). FTC imposes $5 billion penalty and sweeping new privacy restrictions on Facebook. Federal Trade Commission press release, July 24. www.ftc.gov/news-events/press-releases/2019/07/ftc-imposes-5-billion-penalty-sweeping-new-privacy-restrictions

Franco, A., Malhotra, N., & Simonovits, G. (2014). Publication bias in the social sciences: Unlocking the file drawer. *Science*, *345*(6203), 1502–1505.

Freelon, D. (2020). Computational research in the post-API age. *Political Communication*.

Gidney, C., & Ekera, M. (2019). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. arXiv.org. https://arxiv.org/abs/1905.09749

King, G., & Persily, N. (2019). A new model for industry–academic partnerships. *PS: Political Science & Politics*, 1–7. https://doi.org/10.1017/S1049096519001021

Kofi Annan Commission on Elections and Democracy in the Digital Age. (2020). *Protecting Electoral Integrity in the Digital Age*. Report. www.kofiannanfoundation.org/app/uploads/2020/01/f035dd8e-kaf_kacedda_report_2019_web.pdf

McLaughlin, T. (2018). How WhatsApp fuels fake news and violence in India. *Wired*, December 12. www.wired.com/story/how-whatsapp-fuels-fake-news-and-violence-in-india/

Munger, K. (2019). Temporal validity. OSF, September 2. osf.io/3mnzu

Narayanan, D., & Ananth, V. (2018). How the mobile phone is shaping to be BJP's most important weapon in elections. *Economic Times*, August 23. https://economictimes.indiatimes.com/news/politics-and-nation/how-the-mobile-phone-is-shaping-to-be-bjps-most-important-weapon-in-elections/articleshow/65508743.cms

PTI. (2019). 2019 polls: BJP to form chain of WhatsApp groups to strengthen communication between party workers. *Economic Times*, December 23. https://economictimes.indiatimes.com/news/politics-and-nation/2019-polls-bjp-to-form-chain-of-whatsapp-groups-to-strengthen-communication-between-party-workers/articleshow/67219816.cms

Stigler Center. (2019). *Digital Platforms and Concentration*. Stigler Center for the Study of the Economy and the State. https://promarket.org/wp-content/uploads/2018/04/Digital-Platforms-and-Concentration.pdf

Warren Democrats. (2020). Fighting Digital Disinformation. Plan. Warren Democrats website. https://elizabethwarren.com/plans/fighting-digital-disinformation