CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

WORKING PAPER

DECEMBER 2020

Research Collaboration on Influence Operations Between Industry and Academia: A Way Forward

Jacob N. Shapiro, Natalie Thompson, and Alicia Wanless

Research Collaboration on Influence Operations Between Industry and Academia: A Way Forward

Jacob N. Shapiro, Natalie Thompson,
and Alicia Wanless

Carnegie's Partnership for Countering Influence Operations (PCIO) is grateful for funding provided by the William and Flora Hewlett Foundation, Craig Newmark Philanthropies, Facebook, Twitter, and WhatsApp. PCIO is wholly and solely responsible for the contents of its products, written or otherwise. We welcome conversations with new donors. All donations are subject to Carnegie's donor policy review. We do not allow donors prior approval of drafts, influence on selection of project participants, or any influence over the findings and recommendations of work they may support.

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.

CONTENTS

Summary

Research on influence operations requires effective collaboration across industry and academia. Social media platforms are on the front lines of combating influence operations and possess a wealth of unique data and insights. Academics have rigorous training in research methods and relevant theories, and their independence lends credibility to their findings. The skills and knowledge of both groups are critical to answering important questions about influence operations and ultimately finding more effective ways to counter them.

Despite shared interest in studying and addressing influence operations, existing institutions do not provide the proper structures and incentives for cross-sector collaboration. Friction between industry and academia has stymied collaboration on a range of important questions such as how influence operations spread, what effects they have, and what impact potential interventions could have. Present arrangements for research collaboration remain ad hoc, small-scale, and nonstandard across platforms and academic institutions.

Federally funded research and development centers (FFRDCs) provide a compelling model for multi-stakeholder collaboration among those working to counter influence operations. Federally funded research institutions—such as the RAND Corporation, the Institute for Defense Analyses, or the MITRE Corporation—have hosted successful cross-sector collaboration between the federal government and academic institutions for more than seventy years. Academic and industry researchers should seek funding and create an analogous institution so the influence operations research community can further collaborative research on shared interests that cannot be addressed with existing models. Drawing from such models, industry would be a primary funder, but governments and philanthropic donors could also contribute to encourage independence and balance.

Introduction

Over the first half of 2020, Carnegie's Partnership for Countering Influence Operations (PCIO) convened a series of working meetings with academic researchers and representatives from the technology sector to explore joint research aims and the challenges to pursuing them. These meetings were structured around a working paper drafted by PCIO partners at Princeton University's Empirical Studies of Conflict Project.[1] That paper analyzed how existing institutions have failed to facilitate academia–tech industry collaboration and explored institutional models that have enabled research collaboration between the defense community and academia.[2]

These working meetings produced a clear and multifaceted picture of the problem. Academic structures incentivize talented researchers to focus on work that advances tenure-track careers but do not necessarily reward descriptive studies that could help civil society and industry to counter influence operations. Academics' access to data is spotty, as platforms differ in the kinds of data they release and in their structures for granting privileged access to specific researchers. Legitimate privacy and intellectual property concerns hinder efforts to grant greater data access. Other challenges—such as uneven financial resources for research, disciplinary silos, and even antagonism between academic and industry partners—continue to plague academia-industry collaboration.

Many of these challenges are not unique to the field of countering influence operations, and a range of institutional models have been used in other areas to help foster research collaboration. Among these, federally funded research and development centers (FFRDCs) offer a particularly relevant and promising model (see box 1). An analog to FFRDCs for influence operations—what this paper refers to as a multi-stakeholder research and development center (MRDC)—can support research efforts across both sectors to the benefit of society.

BOX 1

Federally Funded Research and Development Centers

FFRDCs are research institutions funded by the U.S. federal government that "provide federal agencies with R&D capabilities that cannot be effectively met by the federal government or the private sector alone."[3] Throughout their nearly seventy-year history, FFRDCs have developed a reputation for producing objective, independent, and credible research on issues of importance to both the public and private sector. Their prestige and influence on a wide range of sensitive national security topics have attracted talented researchers and allowed FFRDCs to emerge as trusted interlocutors between academia and the federal government.[4]

FFRDCs emerged at a time in U.S. history when scientists, engineers, and academics mobilized to support science and technology research in the post–World War II era. Consensus had emerged that properly harnessing scientific research would be necessary to support U.S. strategic aims.[5]

A few key factors have enabled FFRDCs to succeed in this role. They are operated with federal funding on a not-for-profit basis by contractors (including universities and other not-for-profit organizations), typically through renewable five-year contracts.[6] Federal regulations prohibit FFRDCs from competing for contracts outside the scope of their research mandate, meaning that the organizations offer a degree of commercial interest protection to those that share sensitive data with the organization.[7] Contracting guidelines protect the independence of research conducted within the organization, meaning that technically capable researchers can maintain the credibility of their work while still gaining access to classified government or sensitive corporate information.[8]

FFRDCs accomplish five essential tasks that would lay the foundation for long-term collaboration between academia and industry on influence operations:

1. They facilitate sustained funding for long-term projects.
2. They provide a venue for developing shared research agendas and a mechanism for executing studies.
3. They create conditions that help build trusted, long-term relationships between industry and academic counterparts.
4. They offer career opportunities for talented researchers to produce credible work.
5. They guard against inappropriate disclosures while enabling high-credibility studies with sensitive information that cannot be made public.

Beyond these five essential tasks, PCIO identified three specific research areas of particular interest to both academic and industry partners: how influence operations spread across platforms, the effects of influence operations, and the impact of interventions against influence operations. An MRDC could advance work on all three topics by reducing transaction costs, allowing close collaboration and iterative testing, enabling projects with sensitive data, and securing deeper buy-in from platforms and researchers. The benefits of an MRDC are summarized in figure 1 and described in more depth in the following sections.

FIGURE 1

Adapting the FFRDC Model to Counter Influence Operations

Goal for counter-influence operations research community	Requirement for securing collaboration	Relevant capabilities of FFRDCs	Role of an MRDC
Encourage long-term collaboration	Financial resources	Multiyear funding provided by government agencies	Multiyear funding provided by platforms, universities, governments, and/or foundations
	Research career opportunities	Permanent professional staff	Permanent professional staff
	Shared research agendas	Multi-stakeholder models for determining research priorities	Joint participation by industry and academia in determining projects and priorities
		Clear contractual guidelines on procurement, contracting, and limited functionality	Clear contractual obligations of researchers and platforms; MRDC limited to research activities
	Conditions conducive to building trusted relationships	Neutral site for enduring engagement between defense companies, government officials, and researchers	Neutral site for enduring engagement between platforms and academic researchers
	Process for securing sensitive information	Review process for shielding classified information from publication	Data protection standards written into contracts; review process for shielding sensitive corporate data from publication
Study how influence operations spread across platforms	Longitudinal cross-platform data	Secure facilities for storing classified information	Secure systems or facilities for sharing sensitive social media company data
		Security clearance processes for vetting researchers with access to classified information	Vetting process for researchers with access to sensitive company information
	Commercial interest protection	Regulatory prohibitions on competing against for-profit entities	Contractual limitations on use of company data and bylaws prohibiting for-profit work
Measure the effects of influence operations	Access to data from a variety of sources	Secure platform for working with data and routine engagement with a variety of data sources	Secure platform for working with data and staff to reduce the transaction costs of negotiating access to other data sources
	Continuous data access	Maintained through long-term contracts	Earned through trusted relationships and demonstrated adherence to security protocols
Design and study the impact of potential interventions	Information on policy rationales	Trusted relationships between defense community and government researchers	Trusted relationships between social media companies and MRDC researchers
	Continuous monitoring and testing	N/A	Regularly collect and update quantitative data to facilitate hypothesis testing and design of intervention strategies

Fundamentals for Research Collaboration

Financial Resources

FFRDCs rely on federal funding to perform research that supports basic government aims. An analogous institution focused on influence operations would foster cooperation primarily between academia and industry (rather than government).[9] The first and most obvious source of funds for such an enterprise would be social media platforms. They often have ample business profits to support public interest research and in some cases already have funded analogous research.[10] Philanthropic organizations and universities also could be sources of funding, perhaps along the model of academic research centers that operate outside the departmental tenure-track model.[11] (University centers are in turn often funded by wealthy individuals and/or corporations.) Federal funding would remain an option but may require congressional authorization.

Whereas platforms and other supporting organizations currently fund specific, often short-term projects on an ad hoc basis, a FFRDC-like structure could seek and provide funding that would enable researchers to work on long-term projects and avoid the capriciousness of changing attitudes and political developments that could affect where money is directed.

Shared Research Agendas

FFRDCs provide a range of models and precedents that can help guide the development of shared research agendas. Collaboration models range from participation of various stakeholders on advisory boards providing nonbinding feedback to professional staff to more intensive models. In the case of the RAND Civil Justice Institute, for example, the process for determining research agendas involves "a balanced group of plaintiffs' lawyers, defense lawyers, judicial officers, insurers, [and] representatives from other industries and from consumer and labor groups" that "advises the staff on the development of new research projects."[12] At RAND's Arroyo Center, a governing board member sponsors specific projects and maintains ultimate responsibility for approving research design, ensuring data access, and reviewing publications for accuracy.[13]

The specific process by which an intermediary organization determines research agendas, hires researchers, and selects a board of advisers will be one of the most contentious challenges for any institutional model and is outside this paper's scope. The challenge is not unique to an MRDC but characterizes any effort to conduct cross-sector research, a hurdle that will require a careful balancing of interests with clearly articulated processes for executing research agendas. (Indeed, in other sectors, these challenges often lead to the creation of multiple research institutions with competing models, talent, and so on.)

Conditions Conducive to Building Trusted Relationships

One of the greatest successes of FFRDCs is their ability to nurture trusted relationships among academic researchers and defense community counterparts. FFRDCs provide an enduring, neutral venue in which researchers and their industry counterparts can collaborate with one another.[14] As such, they can enable productive collaboration by potentially reducing (or at least setting aside) some of the mistrust among stakeholders and allowing researchers and industry to identify shared priorities and work toward common goals.

Research Career Opportunities

Permanent, full-time research staff operate FFRDCs. Adopting the same practice could allow an MRDC to attract talent capable of producing credible, peer-reviewed work outside of traditional academic positions.[15] Providing the chance to develop a career within such an institution could mitigate the tenure-clock pressures that make temporary secondments and time spent pursuing operational or descriptive research a difficult sell. The organization could allow researchers to maintain credibility through contractually stipulated limitations on the influence that platforms (and other funders) could exert over the publication of findings. Multi-stakeholder funding—through academic and/or philanthropic organizations—would further mitigate credibility concerns by minimizing reliance on platform funding and potential conflicts of interest. Pooling of platform funds into an aggregate budget would obviate the perception and/or reality that a particular platform was influencing the outcomes or products of specific researchers or teams.

Means of Reviewing Research and Withholding Sensitive Information From Publication

A peer review process is critical for ensuring the credibility and quality of research. Engaging industry in the process in a limited manner can also support the legitimate need to shield sensitive data from publication. FFRDCs have built rigorous, systematic processes for peer review of work with sensitive information, including classified data and sensitive health information.[16] An MRDC could do the same by contractually requiring peer review of research and by conducting that peer review process at arm's length from the companies. The review process should include industry, both to verify that publicly released studies do not include sensitive corporate information and to comment on findings. And just as the Government Accountability Office includes its response to agency comments in its reports, MRDC reports should include industry comments and a response to them. These steps would further enhance the credibility of research and help shield sensitive information from publication.

Specific Research Questions

The MRDC model is well suited to support the specific types of research needed in this space. Academic researchers and industry partners share an interest in investigating many topics, including how influence operations spread across platforms, what effects they have on beliefs and real-world behaviors, and what impact potential interventions may have on users. Answering questions in each area presents distinct challenges and requires different levels of collaboration. In the following section, these areas of research are listed by increasing order of difficulty and the level of collaboration they require. Each successive set of research tasks should build on the coordination mechanisms described for the prior set of research aims. This way, platforms and researchers can develop trust and enhance processes while addressing more straightforward questions before turning to ones that are more complex and require deeper collaboration.

How Influence Operations Spread Across Platforms

Influence operators use multiple platforms to spread content, but research to date has not adequately answered how influence operations move across platforms.[17] Filling this gap requires improved infrastructure for data sharing among platforms and with scholars. Having an MRDC to support this infrastructure would reduce transaction costs for platforms and researchers by enabling cross-platform data sharing while protecting commercial interests.

Longitudinal cross-platform data and data security standards: An intermediary organization like an MRDC could provide data security guarantees and reduce transaction costs involved in working with data from multiple organizations. Such an organization could address privacy and commercial concerns that plague data-sharing efforts. It could also facilitate the creation of standard data formats for sharing cross-platform data.

An MRDC also could address data protection concerns by providing a secure facility and technology platform where researchers can work with data.[18] Providing a secure system that prevents leaks would also partially address commercial concerns by ensuring that sensitive data do not fall into competitors' hands. A means of vetting researchers—whether through security clearance–like processes, nondisclosure agreements, noncompete clauses, or a combination of these—before allowing them to work within the secure facility can further reduce risks.[19]

In addition, an MRDC could help solve other operational challenges that are difficult or time-consuming to solve on an ad hoc basis.[20] For example, platforms differ in the types and amounts of data they release publicly, meaning researchers may have access to an incomplete set of data necessary for tracking an operation across multiple platforms. The internal structure of social media platforms may

also necessitate working with multiple teams within a single company to obtain access to all the relevant data. At Facebook, the Product and Content Policy team works with specialists dedicated to managing search functions or misinformation. This team is separate from those that work on child safety issues or cybersecurity policy, but data from both teams may be essential to researchers studying influence operations.[21] An MRDC could address these obstacles to operational data sharing. Rather than have each researcher take on the task of securing data access individually, an MRDC could oversee the process and then distribute access to its professional researchers on a project-by-project basis. Such an arrangement would also relieve platforms of the responsibility to vet external researchers and the liability that comes with providing data access.

An MRDC could also enable cross-platform research by establishing standards for determining the scope of data sharing and a workable format for data storage.[22] In such a model, contracts between the organization and platforms would stipulate data requested and their format. This arrangement would also allow cross-functional platform teams to better work with internal data, making such a practice operationally valuable to the companies. As with data security, delegating this task to an intermediary institution would allow it to be completed in a more efficient, enduring, and rationalized manner rather than on an ad hoc, project-by-project basis.

Commercial interest protection: By assuming responsibility for apportioning data access, an MRDC could also secure the commercial interests of platforms. The Federal Acquisition Regulation prohibits FFRDCs from competing for commercial contracts in their area of research. This stipulation clarifies that FFRDCs are not competitors but instead are nonprofit public interest actors with few incentives to betray the trust of industry partners.[23] By writing similar terms into its contracts with both companies and researchers (akin to including noncompete clauses in employment contracts), an MRDC could mitigate concerns related to employee use of sensitive business information for profit in other contexts. In the defense community, as Jacob N. Shapiro and his co-authors put it:

> The [Institute for Defense Analyses] has dealt with proprietary information from multiple companies competing for multimillion and sometimes multibillion-dollar contracts. . . . Firms are comfortable sharing proprietary information with the [institute] because they knew the organization has strong security protections in place and that its contracts prevented it from monetizing any information it receives.[24]

A key design question emerges in considering both data-sharing and commercial interest protection: Should an intermediary organization be designed to provide data access only to in-house, professional staff, or should access be expanded to external researchers working on ad hoc projects?

The first option—a closed model in which only the employees of an MRDC have access to data—provides strong protections. By limiting the number of people with access to sensitive platform data, platforms and researchers can first test data-sharing arrangements on a smaller scale and build trusted relationships. After demonstrating the feasibility of data sharing in this manner, an MRDC could later explore an open model in which external researchers can apply for access to data and undergo a vetting process without necessarily becoming employees of the organization. Arrangements for academics to undertake short-term visits (while on sabbatical or in their allowed time for outside activities) could further expand collaborative progress.

The Effects of Influence Operations

Platforms and academic researchers also want to know how exposure to and interactions with online content changes user attitudes and behaviors—that is to say, the real-world effects of influence operations. This type of research faces the additional challenges of securing access to information not held by platforms (information that is instead held by other private companies, government organizations, or nonprofit organizations) and building infrastructures and relationships that allow for long-term engagement on specific projects.

Access to data from sources other than platforms: An MRDC can help expand access to data sources. To understand how exposure to online content affects voting behavior, for example, researchers might want access to voting records, exit poll information, or deanonymized survey data—information that they can then attempt to match with user identity. Similarly, to understand the impact of exposure to health misinformation on user health, researchers may want access to data used in medical or public health studies or other data collected by public health authorities. Project contracts could specify the data necessary for undertaking a research project, and the MRDC could approach external sources on an ad hoc basis to secure additional data. The credibility of the organization as a neutral intermediary, coupled with demonstrated secure data-sharing processes, could help facilitate such partnerships.

Continuous data access: The FFRDC model also provides a site of enduring engagement for project teams studying influence operations longitudinally and can help facilitate continuous data access. As platforms engage in A/B testing, researchers studying the impact of influence operations would need access to regularly updated data, rather than the kind of snapshot data provided under the status quo. By serving as a site of enduring engagement, an MRDC could enable platforms to develop the kind of trust necessary to agree to this level of data sharing.

Impact of Potential Interventions

The set of research aims requiring the most significant collaboration involves sharing information about interventions to counter the spread of influence operations across platforms.[25] Platforms have tried various policy changes and content moderation strategies to counter influence operations, but the results of these efforts are often unknown outside the platforms. Furthermore, researchers and platforms lack an adequate evidence base for evaluating the results of these efforts. Platforms sometimes provide public rationales for macro-level policy changes or aggregate information about content moderation.[26] Still, the information related to specific content moderation decisions (for example, whether to remove a page because of connections to a hateful or criminal organization) are often opaque, as are the consequences of these decisions. In some cases, reporters have highlighted that malicious operators simply recreate pages and profiles when they are removed from platforms.[27]

Information on policy rationales: By promoting trusted relationships between researchers and their industry counterparts, an MRDC could help facilitate the sharing of qualitative information regarding content moderation decisions. To comprehensively study the impact of potential interventions, researchers would benefit from understanding the contextual factors driving a decision's timing and the details of rollout and operationalization. Through repeated interactions with industry counterparts in a neutral venue like an MRDC, the sharing of nonquantitative information could ideally become a routine, low-risk practice that informs research on the impact of potential interventions. Sharing this kind of information would, in turn, benefit the platforms by providing them with a credible external source for evaluating their actions and a broader evidence base for choosing effective courses of action and avoiding counterproductive policies.

Continuous monitoring and evaluation: At the most collaborative level, platforms and researchers might co-design and test strategies for countering influence operations. An MRDC could allow researchers to continuously monitor and test interventions' results to evaluate their impact and effectiveness, using insights from these trials to inform future strategies. Similarly, project co-design would allow MRDC researchers to work with platform employees to act on research findings in real-time rather than wait for research to be fully published, an important benefit when dealing with information threats that can quickly evolve. Given the particular sensitivities and coordination needs of this specific research aim, industry and academic partners may need time to build the kind of trust underpinning such coordination.

Looking Ahead

Adopting this collaborative model would raise new questions and challenges and require a patient, iterative process of trial and error. For example, FFRDCs have succeeded by enabling collaboration among a particular set of actors, each playing a distinctive role. When platforms take on the government's role of funding and contracting with researchers, will the model still succeed?

Having all the details figured out is not a prerequisite for moving forward. FFRDCs earned their credibility as neutral intermediary institutions through decades of iteration. This paper has attempted to outline a scalable solution for promoting research collaboration on influence operations with the hope that, by trial and error, industry and academic partners can adapt the FFRDC model to find a workable format for cross-sector collaboration.

To that end, academic and industry partners should begin scoping what a first attempt at an MRDC might look like. The effort should start with at least two of the social media platforms determining a means of funding such an institution—either by directly funding a set of projects or creating a trust that would work out the specifics of such an arrangement. Parties would then need to scope the characteristics of the institution and the contracts written between researchers and platforms. These contracts would specify conditions under which the new institution and its staff could engage in work for hire, allow the companies to solicit specific studies, and establish a peer review process. The contracts would also include some kind of nondisclosure agreement and data-sharing requirements.

Over the last several years, progress has been made on some of these fronts, especially data sharing, proving it is possible to create the conditions necessary for successful collaboration. Twitter, for example, hosts a publicly available archive of the state-sponsored information operations its content moderators have removed from the platform. The company also allows specialist researchers to request access to an unhashed version of the data contained in the archive.[28] Facebook acquired a tool called CrowdTangle that enables "content discovery and social analytics" and has opened use of the tool to academics and researchers.[29] And earlier this year, a team of researchers at the University of Texas at Austin, New York University, and Facebook announced a partnership for election-related research in which Facebook researchers will share some sensitive information with a small group of academic researchers. The agreement allows the company to "check that papers do not violate legal or privacy obligations" and restricts raw data access to Facebook employees, but the agreement's terms include significant transparency requirements and safeguard academic independence.[30]

Alone, however, such efforts are insufficient. The scale of the challenge is simply too vast. A robust, enduring structure is needed to facilitate such work. Researchers and industry partners should look to the history of FFRDCs and begin building the structures necessary for successful collaboration. A new organization would deepen the capacity for the kinds of work being performed by status-quo institutions. More importantly, it would open doors to new avenues of collaboration and cross-sector support for mutual research goals. In the interest of bringing the community countering influence operations together, this seems a viable and worthwhile path forward.

About the Authors

Jacob N. Shapiro is a professor of politics and international affairs at Princeton University whose research covers conflict, economic development, and security policy.

Natalie Thompson was most recently a research assistant with the Technology and International Affairs Program at the Carnegie Endowment for International Peace.

Alicia Wanless is the director of the Partnership for Countering Influence Operations at the Carnegie Endowment for International Peace and a doctoral researcher in war studies at King's College London.

Notes

1 Jacob N. Shapiro, Michelle Nedashkovskaya, and Jan Oledan, "Collaborative Models for Understanding Influence Operations: Lessons From Defense Research," Carnegie Endowment for International Peace, June 25, 2020, https://carnegieendowment.org/2020/06/25/collaborative-models-for-understanding-influence-operations-lessons-from-defense-research-pub-82150.

2 Ibid.

3 Marcy E. Gallo, "Federally Funded Research and Development Centers (FFRDCs): Background and Issues for Congress," Congressional Research Service, December 1, 2017, 1, https://fas.org/sgp/crs/misc/R44629.pdf.

4 Shapiro, Nedashkovskaya, and Oledan, "Collaborative Models for Understanding Influence Operations."

5 Ibid.

6 Ibid.

7 Ibid; and Gallo, "Federally Funded Research and Development Centers," 3.

8 Shapiro, Nedashkovskaya, and Oledan, "Collaborative Models for Understanding Influence Operations."

9 Gallo, "Federally Funded Research and Development Centers."

10 "Facebook Supports Research on Misinformation and Polarization With $2 Million Commitment," Facebook Research (blog), February 24, 2020, https://research.fb.com/blog/2020/02/facebook-misinformation-polarization-rfp-two-million-dollar-commitment.

11 For example, the Center for Security and Emerging Technology at Georgetown University receives funding from the William and Flora Hewlett Foundation, the Open Philanthropy Project, and the Public Interest Technology University Network to "[produce] data-driven research at the intersection of security and technology, providing nonpartisan analysis to the policy community." See "About Us," Center for Security and Emerging Technology, https://cset.georgetown.edu/about-us.

12 RAND Social and Economic Well-Being, "RAND Institute for Civil Justice Board of Advisors," RAND Corporation, updated October 2020, https://www.rand.org/well-being/justice-policy/centers/civil-justice/about/board.html.

13 The Arroyo Center is the FFRDC tasked with supporting the Department of the Army's research. There, regulation has "establishe[d] a governing board of Army leaders known as the Arroyo Center Policy Committee," which meets twice a year with the Arroyo Center's professional staff. At least one senior leader from the army sponsors projects undertaken by the center and "has responsibility for helping to formulate the project, providing access to needed data and other information, monitoring its progress, reviewing its publications for accuracy, utilizing its findings, and implementing its recommendations." See RAND Army Research Division, "Arroyo Center Policy Committee," updated October 2020, https://www.rand.org/ard/about/policy_committee.html.

14 As Shapiro and his co-authors describe the predecessor organization to the Institute for Defense Analyses, part of its mission was "to approach its tasks from an impartial perspective that transcended the various branches of the armed forces." See Shapiro, Nedashkovskaya, and Oledan, "Collaborative Models for Understanding Influence Operations."

15 Gallo, "Federally Funded Research and Development Centers," 3; and Shapiro, Nedashkovskaya, and Oledan, "Collaborative Models for Understanding Influence Operations."

16 Ibid.

17 For example, Graphika has detailed a Russian influence operation that included accounts on Facebook, Instagram, Twitter, YouTube, Medium, Tumblr, Reddit, Telegram, Pinterest, and other blogging sites. See Ben Nimmo, Camille François, C. Shawn Eib, Léa Ronzaud, and Joseph Carter, "Further Exposures of Russian Military Assets Across Platforms, 2013–2020," Graphika, September 24, 2020, https://graphika.com/reports/gru-and-the-minions.

18 Though not a feature of FFRDCs specifically mandated by law, such an arrangement has precedent in other federally funded institutions that allow for data access. Federal Statistical Research Data Centers (FSRDCs), for example, are designed specifically to provide researchers access to deanonymized, disaggregated Census Bureau data. See Shapiro, Nedashkovskaya, and Oledan, "Collaborative Models for Understanding Influence Operations."

19 Ibid.

20 Though previous efforts—notably, Social Science One—have attempted to address the issue of data sharing, they have suffered without institutional resources like those provided by FFRDCs or FSRDCs that can help reduce operational frictions and provide opportunities for collaboration. For example, privacy concerns hindered the release of the originally promised URL dataset, but the protections afforded by an FFRDC or FSRDC for managing data in a secure manner—sometimes a secure facility—and the vetting of potential research staff can help overcome some of these limitations. See "Public Statement From the Co-Chairs and European Advisory Committee of Social Science One," Social Science One, December 11, 2019, https://www.socialscience.one/blog/public-statement-european-advisory-committee-social-science-one.

21 Facebook Careers, "Product & Content Policy at Facebook: How We Build Safe Communities," Facebook Careers (blog), February 20, 2019, https://www.facebook.com/careers/life/policy-at-facebook-how-we-build-safe-communities.

22 Institutional models like the U.S. National Institute of Standards and Technology (NIST) have helped create common languages for technical information in other contexts. Emulating these efforts could help influence operations researchers determine the appropriate scope of data and a workable format for sharing it. NIST has worked with industry partners on a variety of cybersecurity issues, releasing in 2014 a cybersecurity framework that "provides a common language that allows staff at all levels within an organization—and at all points in a supply chain—to develop a shared understanding of their cybersecurity risks." NIST's Cybersecurity Center of Excellence (NCCoE) also convenes industry, government, and academic experts to address cybersecurity challenges: "Using standards and best practices, the NCCoE and its collaborating partners demonstrate how to apply secure technologies to accelerate the adoption of cybersecurity and improve the security posture of businesses." For more on NIST's cybersecurity framework, see: "Cybersecurity Framework," National Institute of Standards and Technology, accessed November 19, 1010, https://www.nist.gov/industry-impacts/cybersecurity-framework; and for more on NCCoE, see National Cybersecurity Center of Excellence, "Projects," National Institute of Standards and Technology, accessed November 19, 2020, https://www.nccoe.nist.gov/projects.

23 Gallo, "Federally Funded Research and Development Centers," 3; and Shapiro, Nedashkovskaya, and Oledan, "Collaborative Models for Understanding Influence Operations."

24 Shapiro, Nedashkovskaya, and Oledan, "Collaborative Models for Understanding Influence Operations."

25 The question of potential interventions can be divided into two areas that could be addressed with different models. The first—testing in the abstract—could more easily be accomplished by properly aligning incentives to pursue social science research. Some of this research has already been conducted by academic researchers. For example, psychological research has examined the role that "prebunking"—that is, "preemptively warning and exposing people to weakened doses of misinformation"—can play in enhancing resistance to misinformation. Other studies have examined the impact of misinformation labels on user

perceptions of unlabeled content. See: Jon Roozenbeek, Sander van der Linden, and Thomas Nygren, "Prebunking Interventions Based on 'Inoculation' Theory Can Reduce Susceptibility to Misinformation Across Culture," *Harvard Kennedy School Misinformation Review* 1, no. 2 (January 2020): 1–23, https://misinforeview.hks.harvard.edu/wp-content/uploads/2020/02/FORMATTED_globalvaccination_Jan30.pdf; and Gordon Pennycook, Adam Bear, Evan T. Collins, and David G. Rand, "The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings," *Management Science Articles in Advance* (February 2020): 1–14, https://pubsonline.informs.org/doi/pdf/10.1287/mnsc.2019.3478. The authors thank Kamya Yadav for drawing their attention to this important research.

26 Twitter, for example, published an update in May 2020 to explain how the coronavirus pandemic affected their misleading information policy. See Yoel Roth and Nick Pickles, "Updating Our Approach to Misleading Information," Twitter blog, May 11, 2020, https://blog.twitter.com/en_us/topics/product/2020/updating-our-approach-to-misleading-information.html; and platform transparency reports: Google Transparency Report, "YouTube Community Guidelines Enforcement," Google, accessed November 19, 2020, https://transparencyreport.google.com/youtube-policy/removals; Facebook Transparency, "Community Standards Enforcement Report," Facebook, accessed November 19, 2020, https://transparency.facebook.com/community-standards-enforcement; Twitter Transparency, "Transparency Reports," Twitter, accessed November 19, 2020, https://transparency.twitter.com/en/reports.html.

27 According to an internal memo that Buzzfeed obtained from a former Facebook employee, efforts to remove Honduran influence operations resulted in the removal of thousands of profiles and assets. However, the operators had recreated their campaign within two weeks and the author of the internal memo wrote, "A year after our takedown, the activity is still live and well." See Craig Silverman, Ryan Mac, and Pranav Dixit, "'I Have Blood on My Hands': A Whistleblower Says Facebook Ignored Global Political Manipulation," Buzzfeed News, September 14, 2020, https://www.buzzfeednews.com/article/craigsilverman/facebook-ignore-political-manipulation-whistleblower-memo.

28 Twitter Transparency, "Information Operations," Twitter, accessed November 19, 2020, https://transparency.twitter.com/en/reports/information-operations.html.

29 Facebook for Media, "CrowdTangle for Academics and Researchers," Facebook, https://www.facebook.com/formedia/blog/crowdtangle-for-academics-and-researchers.

30 Talia Stroud, Joshua A. Tucker, Annie Franco, and Chad P. Kiewiet de Jonge, "A Proposal for Understanding Social Media's Impact on Elections: Rigorous, Peer-Reviewed Scientific Research," Medium (blog), August 31, 2020, https://medium.com/@2020_election_research_project/a-proposal-for-understanding-social-medias-impact-on-elections-4ca5b7aae10.

1779 Massachusetts Avenue NW | Washington, DC 20036 | P: + 1 202 483 7600

CarnegieEndowment.org