

B R I D G I N G
THE DIVIDE:
TRANSLATING RESEARCH
ON DIGITAL MEDIA INTO
POLICY AND PRACTICE

ISSIE LAPOWSKY

COMMISSIONED BY



**KNIGHT
FOUNDATION**

INTRODUCTION

Jeff Horwitz was chasing down a tip on Instagram last year when he stumbled on a particularly troubling finding. As a technology reporter covering Meta for the *Wall Street Journal* and author of a book about the company, *Broken Code*, Horwitz has seen some of the worst content people post on Meta's platforms. But this Instagram account he found, which appeared to belong to a 14-year-old selling sexual photos of herself, stopped him short. Even more startling: Horwitz noticed Instagram was recommending additional accounts to him that seemed to be trafficking in similar content.

We've got a problem, he remembers thinking. Horwitz knew he'd need help not just scoping the size of the issue, but also navigating a sea of potentially illegal posts. So he turned to a group of researchers at the Stanford Internet Observatory. The group, which included Reneé DiResta, Alex Stamos and David Thiel, had deep experience studying Meta's recommendation systems, and Stamos and Thiel had spent years on Facebook's security team before joining Stanford. Working in collaboration, Horwitz and the Stanford researchers soon uncovered hundreds of accounts that appeared to be selling "self-generated child sexual abuse material" on Instagram and found that Instagram's recommendation engine was playing a key role in connecting suspected buyers and sellers.

In June of last year, Horwitz and a colleague wrote up their findings in a bombshell [story](#) documenting the "vast pedophile network" on Instagram. The Stanford team published their own detailed [paper](#) the same day. The response was explosive. Meta told the *Journal* it was setting up an internal task force to address the investigation's findings. Congressional leaders announced inquiries. And New Mexico's attorney general used the story to

support a [lawsuit](#) accusing Meta of enabling child sexual exploitation and mental health harms.

For Horwitz, the Stanford researchers' input was essential to getting that information out into the public as safely and robustly as possible. "There was clearly a gravitas and a skill set that we didn't have," Horwitz says.

The Stanford Internet Observatory, housed at the university's Cyber Policy Center, is just one of more than 60 institutions that the John S. and James L. Knight Foundation has supported since it set out to catalyze a field of research focused on digital media and democracy in 2019. At the time, the foundation's hope was to address the rising pollution of the on-line information ecosystem by launching something akin to a public health moment for the internet.

Since then, Knight Foundation has committed more than \$107 million in grants to fund research on a wide range of topics related to how information travels online, from the study of mis- and disinformation to antitrust enforcement, Section 230 reform and broader trust and safety issues. Together, these grants have supported the work of more than 800 people, including researchers, administrative staff and students. Knight hasn't been alone in this endeavor either. Other organizations, including Hewlett Foundation, Charles Koch Foundation, Craig Newmark Philanthropies, Omidyar Network and the National Science Foundation, have contributed significantly to the field with funding efforts of their own.

As this field has grown, so has its impact on the world outside the research community. In 2022 alone, Knight grantees published more than 900 articles, participated in more than 900 speaking engagements and testified before Congress 15 times. But as Horwitz's collaboration

with the Stanford team demonstrates, those numbers only tell part of the story. What they don't show is what happens next: How do people in positions of power actually respond to and act on that information? How has this field of research driven technology-related policies and investigations inside the government? How has it influenced tech companies' own content decisions and user guidelines? What front page news wouldn't have been possible without research findings to back it up, and how does research help advocacy groups rally public support around a cause?

The answers to those questions can be hard to track. Often, research travels a circuitous route through news coverage and closed-door briefings on its way to making an impact in the world. This paper—which draws on interviews with roughly 40 sources across government, media, the tech industry and civil society groups—is an attempt to trace that route. It finds that, in just a few short years, researchers in this field have played an instrumental role in driving legislative proposals, platform policies, product changes, media coverage and advocacy campaigns. While the impact of a select group of researchers, including at the Stanford Internet Observatory, is easy to pinpoint, the growth of this field has also had a cumulative effect in terms of setting the agenda for what government leaders, tech workers, journalists and advocates are paying attention to.

As Yasmin Green, CEO of the Google research subsidiary Jigsaw, put it, “There's more than one thing that publishing research does. Of course it can and should inform further research and action, but the other thing it does is spotlight topics in the public debate. That's an area where this research community has grown in not just size and sophistication, but it's also grown in clout.”

And yet, even as conversations for this report revealed what's working in the field, they also revealed what's not. Slow publishing timelines and research questions that aren't aligned with practitioners' needs remain significant impediments to impact. Meanwhile, the overreliance on personal connections at large tech companies can give select “superstar” researchers preferential access to data, while others—particularly researchers operating outside of the U.S.—wind up locked out. That has downstream effects on who gets funded, covered in the press and ultimately called upon by decision makers in government and industry.

For funders like Knight, these challenges are also opportunities. By elevating new, emerging voices in the field and investing in time-sensitive ways to synthesize research findings for public consumption, Knight can help maximize the impact of vitally needed research.

* Institutions that have received research funding from Knight Foundation are identified throughout with an asterisk.

PART 1:

How research reaches tech companies

Before Twitter transformed into X, its leaders had something of a confidence problem. The company was making ever more complex content moderation decisions as it grew, which was making Twitter's senior leadership deeply uneasy. "Executives at Twitter very consistently felt illegitimate about the decisions that the company was making," says Yoel Roth, the company's former head of trust and safety, who worked at Twitter from 2015 to 2022.

That became a challenge for Roth in the run-up to the 2020 U.S. presidential election, when his team was testing ways to curb the spread of misinformation on Twitter with interventions like "pre-bunks," which are strategically placed messages that expose users to known types of propaganda in order to build their defenses against it. The Twitter team measured the impact of these interventions in hopes of getting the green light to deploy them more broadly. But these internal studies failed to convince senior leaders, says Roth, who is now vice president of trust and safety at Match Group. "They didn't feel that they had validation of its effectiveness," he says.

So Roth's team began a review of all the available studies related to pre-bunking. Researchers at the [University of Cambridge](#) had already conducted experiments demonstrating that when people are exposed to common misinformation

techniques, it has an inoculating effect, making them less likely to believe that type of misinformation when they encounter it later on. Roth credits that research and other academic work with helping convince Twitter's senior leadership to use pre-bunks widely during the 2020 election cycle. "We were able to say: Look, there's the theory. Here's our own findings. This is clearly the right path forward," Roth says.

To have a direct impact inside a company, research must be both specific and actionable, tech industry sources said. This is true not just for designing new product features, but for platform policies, too, says Dave Willner, who previously led trust and safety and content policy teams at OpenAI, Airbnb and Facebook. When Willner was first developing hate speech policies at Facebook, he says most of the research defined hate speech in terms of broad value judgments that were impractical given the way content moderation works. "All of the stuff out there was like, 'Does this offend the dignity of the human person?' I definitely cannot teach an outsourcer to use that to make decisions," Willner says.

But when he met Susan Benesch, executive director of the Dangerous Speech Project and a faculty associate at Harvard University's Berkman Klein Center,* he found a framework that clicked. Benesch had outlined specific hallmarks of what she calls "dangerous speech," like, for instance,

the use of [dehumanizing language](#) that compares people to vermin. The recommendations felt to Willner like a set of instructions that content moderators around the world could follow. “It was like, ‘This is a different kind of can opener. I can use this to open cans,’” says Willner, who is now a non-resident fellow at the Stanford Cyber Policy Center.*

Benesch’s work heavily informed Facebook’s hate speech policies, as well as how the company classified hate groups that should be banned from the platform altogether. And, years later, when Willner was leading community policy at Airbnb, he says these principles proved useful all over again when Airbnb was trying to block members of hate groups from booking housing for the Unite the Right rally in Charlottesville, Virginia.

Beyond the impact that individual researchers have had on tech companies, there are also signs that the field writ large has pushed the tech industry in new directions. For one thing, tech companies have hired more researchers who study social media internally—though those research teams have also been heavily impacted by recent industry layoffs. And while data access remains a major impediment for researchers studying these platforms, companies have undoubtedly built new data access tools, like, for example, the [Meta Content Library](#), largely in response to the research community’s public and private clamoring for data. (It should be noted, however, that Meta has simultaneously announced plans to [shutter](#) another important transparency tool in Crowdtangle.) The emergence of this field has also birthed a growing number of workshops and conferences where researchers and tech workers get to learn from one another. Several tech industry sources cited these convenings as critical to their ability to tap into expertise outside the Silicon Valley bubble.

Still, while this field’s growing influence on tech companies at both an individual and collective level is obvious, substantial barriers remain.

Almost every tech industry source interviewed for this paper, for instance, highlighted the persistent problem of research timelines. When academic papers take years to publish, their findings are often obviated by the fact that the platforms have already changed substantially.

Another challenge is the persistent lack of access academics have to platform data. Despite some recent progress toward transparency, as mentioned above, there is still so much vital information locked inside tech companies that outside researchers can’t see, requiring them to use methodological workarounds that insiders view as flawed. Roth of Twitter pointed to frequent research findings on bots as an example of this. Trying to spot bot networks on Twitter without access to data like IP addresses, Roth says, is like working “blindfolded with one of your arms tied behind your back.”

One way researchers have sought to address the data access issue is to collaborate with tech companies on research. While there’s value in that approach, those partnerships can also raise tricky questions about data privacy and researcher independence and are susceptible to the whipsawing whims of company leaders. Rebekah Tromble, director of the Institute for Data, Democracy, and Politics at George Washington University,* learned this lesson the hard way. In 2018, Twitter [selected](#) Tromble to lead a team studying the “health” of conversations on the platform, then quietly withdrew promised data access, effectively derailing the project Tromble and her team signed up for.

These partnerships also don’t guarantee quick turnaround times. When Facebook announced its unprecedented 2020 election research project—which was led by Joshua Tucker of New York University’s Center for Social Media and Politics* and Talia Stroud of the University of Texas at Austin’s Center for Media Engagement*—the researchers expected their findings to come out in 2021. The complexity of the project resulted in the

[first four papers](#) being published just last year, with many more still to come.

Even when research partnerships proceed exactly as planned, they are still an imperfect solution to the data access issue. These partnerships are, by their nature, selective and rely on personal relationships between tech executives and top researchers at well-resourced institutions. It's these researchers who get access to proprietary data, leaving others out. "When data is provisioned selectively to researchers that are handpicked by the company, or a funder associated with the company, that is very problematic," says Lauren Wagner, Meta's former head of product marketing for open research and transparency and a current fellow at Berggruen Institute.

This isn't just an issue of equity. Several sources noted that the lack of relationships with international researchers in particular—and the lack of funding going into international research—means that tech companies are often getting the least outside input in the places where they have the least internal expertise. This issue is exacerbated by the fact that, often, tech companies are most responsive to research that gets covered in U.S. media or that gains the attention of U.S. policymakers, leaving vast swaths of the world under-scrutinized.

As a result, several tech industry sources identified the need for more research funding to flow into international organizations and institutions. For those organizations that are unable to directly fund international work, designing in-person events with international researchers in mind may be another way to build more personal connections to people working at U.S.-based tech companies.

Others recommended more investment in third-party labs where representatives from both academia and industry can come together to work on targeted research questions or design experimental staging environments that mimic how platforms actually operate. Emily Saltz, a senior user experience researcher at Jigsaw, highlighted the Cornell Social Media Lab's social media simulation platform, Truman, as well as tech tools created by Chris Bail of Duke University's Sanford School of Public Policy,* as examples of useful staging environments that already exist. The goal, Saltz says, should be to ensure that researchers are designing experiments that are more directly applicable to existing tech platforms and to prevent researchers from having to reinvent the wheel each time they conduct those experiments. "Once you get closer to the industry side," Saltz says, "you see how crucial all of these elements of realism are to understanding what's happening."

PART 2:

How research reaches government

Within his first year in office, President Biden announced the formation of an “information integrity” working group tasked with coming up with a government research strategy for studying mis- and disinformation. At the time, the country was overwhelmed by misinformation about the 2020 election and COVID-19, and the goal of the working group was to generate a plan for government agencies looking to fund research on the topic. Over the course of a year, the working group, which included representatives across government agencies, worked in close consultation with experts who already study the information ecosystem, to [map out research priorities](#).

It was wonky work. The final product—a roadmap for government funding of research—isn't exactly the kind of thing that makes national news. But for Christina Ciocca Eller, the White House's former assistant director of evidence and policy and a current assistant professor of sociology and social studies at Harvard, the research community's input was essential. Even if exercises like that don't immediately translate to tangible policy outcomes, they do organize the government's attention on a set of issues, which in turn can influence policy down the line. “It's a much longer process that starts way farther upstream,” Ciocca Eller says.

For government officials, tapping into the collective wisdom of the research community through

high-level discussions like these is often as useful, if not more, than any individual research publication or finding. This is true in the U.S. and abroad. The E.U.'s Digital Services Act, for example, requires very large online platforms to provide data to vetted researchers. According to one European regulator interviewed for this report, that provision was informed by extensive conversations that took place in one [transatlantic working group](#) on hate speech and online extremism. The working group was spearheaded by the University of Pennsylvania's Annenberg Public Policy Center.

Erie Meyer, chief technologist for the Consumer Financial Protection Bureau, also pointed to the [Public Technology Leadership Collaborative](#), which is run by Data & Society,* as another useful venue where government officials and researchers are able to come together for deep discussions on complex topics. “It's this really in-the-weeds thing where academics can disagree with each other in front of you to help you understand an issue,” Meyer says.

Government officials stressed the importance of actually hiring academics internally, even if only for brief tours of duty in government, because they can act not only as translators of research, but also as bridges to the broader research community. “Bringing outside expertise from academia to help with discrete projects

has been tremendously valuable to our work,” says Alan Davidson, administrator of the National Telecommunications and Information Administration (NTIA), which hired Ellen Goodman of Rutgers Institute for Information Policy & Law* to lead the NTIA’s work on artificial intelligence accountability. “In addition to being a leading expert herself, she knows the other great minds we needed to bring into the fold to inform our report,” Davidson says.

The NTIA isn’t the only agency taking this approach. Northeastern University assistant professor Laura Edelson recently wrapped up a stint as chief technologist in the Department of Justice’s Antitrust Division. Olivier Sylvain, a professor at Fordham School of Law,* spent two years as senior advisor to the chair of the Federal Trade Commission. And when the White House Office of Science and Technology Policy was developing its Blueprint for an AI Bill of Rights, acting director Alondra Nelson brought in two leading ethical AI researchers—Suresh Venkatasubramanian, now of Brown University, and Sorelle Friedler, of Haverford College—to lead the work. Several government sources also cited fellowships—including those run by the Institute of Electrical and Electronics Engineers and TechCongress*—as being essential links to the research community. One former DOJ official described TechCongress in particular as “a really powerful pipeline.”

While these formal channels for bringing in outside experts are important, government officials also emphasized the value of informal briefings with researchers. Officials and staffers at all levels of government said they’re far less likely to scour the latest literature on a given topic than they are to pick up the phone and ask a researcher to get them up to speed quickly or provide them with feedback on draft legislation and policy proposals. Better yet if those researchers come to those conversations with realistic ideas for how policymakers and regulators can respond to that research.

One former Republican Senate staffer, for instance, described calling on Kate Klonick, an associate professor at St. John’s University School of Law,* to mark up a discussion draft of child safety legislation. Though the draft was never introduced, the staffer says Klonick’s input helped guide the senator’s office toward more privacy-protective approaches. The staffer said these briefings are crucial given how few resources congressional offices have and how little time they get to produce new legislation. “If those are meaningful pieces of legislation, they each require such a nuanced understanding,” the staffer says. “If I can send around a discussion draft of a bill that hasn’t been introduced, and I can ask for feedback within a week or two-week timeline, I will get really substantive, technical and detailed feedback.”

Still, some of the challenges that tech workers identified in working with researchers were also of concern for government officials, including research timelines. While government agencies and Congress certainly don’t move as quickly as tech companies, they do tend to have very narrow windows of opportunity during which to act, sometimes driven by shifting political will and sometimes dictated by statutes or executive orders. Research often can’t keep up.

Several sources also highlighted the fact that lawmakers tend to rely heavily on research that gets media coverage. That’s partly due to the fact that many Congressional staffers don’t have subscriptions to academic journals and partly due to the fact that congressional offices, in particular, are short-staffed. Press coverage provides a shortcut. “They are just finding the scholars who have tons of followers and who get quoted in the *Washington Post*,” says Anna Lenhart, a former senior technology legislative aide to Democratic Representative Lori Trahan and current staffer at George Washington University’s Institute for Data, Democracy & Politics.*

Lenhart says this creates a “superstar problem,” in which government officials are repeatedly

calling upon the same small group of researchers. Not only does that mean government officials aren't getting the broadest range of perspectives, it can also have an impact on the wider field of research. "It creates this world where academics have to choose: Are they going to play that game *and* try to publish *and* take care of their kids *and* teach?" Lenhart says. "The ones that get really good at influencing get really good at getting money, and it becomes an incentive thing where they keep doing influencer work, and guess what suffers? Teaching. And their rigorous research."

On the other end of the spectrum, some government sources said they were surprised at how seldom researchers proactively reach out to them to promote their findings. "This happens way less than it should," said one Democratic Senate staffer, who described the disconnect between research and government as "mostly a delivery issue." The staffer encouraged researchers to identify which congressional committees are most relevant to their work and email legislative staffers working on those committees, particularly in advance of key moments like hearings or bill markups. Another Republican House staffer confessed that the first thing they do when they need to get up to speed on a topic is search their inbox for experts' emails.

And yet, getting staffers' attention requires research findings to actually be responsive to government officials' questions. Too often, they aren't, government sources said. Unlike in the tech industry, where specificity is key, often, government agencies and lawmakers are looking for meta analyses that summarize what is known broadly about any given topic. That's not always the work that is being pursued or rewarded in academia, says Alondra Nelson, who, since departing the White House, is now a sociologist and professor at the Institute for Advanced Study.

"Sometimes what academics are asking is so in the weeds and narrow," she says. Not only that, but even the best research on harms caused by

technology can fall short of suggesting meaningful policy reforms. "We get trained very well on critique," Nelson says. "But how do you go from knowing what's wrong to knowing how to make something good?"

Ciocca Eller said researchers and funders should consider setting research agendas "based on what the government actually says it needs information about." Since 2022, federal agencies have been publishing their so-called "[learning agendas](#)" outlining their priorities, but Ciocca Eller says they're "unbelievably underutilized" by researchers.

What makes it difficult to quantify this field's collective impact on policy is the fact that, as Ciocca Eller noted, real policymaking that goes beyond dead-end legislative proposals is a slow and grinding process, where there often isn't a one-to-one relationship between a specific research finding and a particular policy outcome. Still, there are indicators that the government sees value in the development of this field, including the fact that the National Science Foundation has also funded research in this space.

Of course, the opposite is also true: As the field has grown and earned the attention of government officials, it's also inspired substantial backlash, particularly among conservative lawmakers. Over the last few years, researchers, including a number of Knight grantees, have faced politicized attacks from the right over their work related to election and COVID-19 misinformation. They've faced subpoenas, death threats and frequent charges of enabling government censorship. These emerging legal and personal threats have not only made it harder for researchers to be public about their work, they've also added further obstacles for government staffers on both sides of the aisle looking to collaborate with researchers, particularly regarding issues of online speech.

One Republican House staffer says the politicization of the research community has limited, for instance, who they're able to call to testify on even

uncontroversial tech topics. “We either have to find conservative, or at least moderate or apolitical academics, to cite or to rely on, or we need to have a case for somebody’s sheer expertise and unique value in their field such that we can argue overlooking other sorts of comments or positions they’ve taken,” the staffer says. These dynamics have even made Democratic staffers wary of openly relying on researchers, said Lenhart. “Even if you’re just trying to get technical information, the fact that you’re talking to academics can look like you only talk to people from the quote unquote left,” she says.

These obstacles to collaboration with researchers are more than just a matter of perception. A string of lawsuits, including [one recently argued](#) before the Supreme Court, have sought to frame academics studying misinformation as state actors engaging in unconstitutional censorship and have attempted to legally bar some researchers from working with the government. In this environment, providing legal support for researchers, as Knight Foundation and others [have already done](#), is crucial.

Perhaps the biggest challenge of all for researchers trying to make an impact in Congress, of course, is the fact that Congress hasn’t done much legislating on issues related to tech and democracy. State legislatures, on the other hand, have forged ahead, passing laws related to child online safety, social media content moderation, privacy and more. But these state-level offices are often operating with even fewer staff and less expertise than congressional offices. Matt Perault, director of the University of North Carolina’s Center on Technology Policy,* has written extensively about [state-level tech policies](#). He says there’s an “inverse relationship” between the actions states are taking and the amount of external research being directed at them.

“No matter your political persuasion, there are groups that provide expertise to Congress, but they’re providing expertise for a system and structure that’s not actually going to pass legislation,” Perault says. “There are lots of states seeking to be active, and that increases the value and importance of expertise funneled into these policy debates.”

For researchers, that knowledge gap presents an opportunity to work more closely with state legislators who are actually passing tech-related bills into law. It also provides an opening for institutions like Knight to fund organizations that are looking beyond the federal government and are focusing instead on translating research into policy in the states.

PART 3:

How research reaches the media

As the 2020 election approached, Brandy Zadrozny was looking for ways to keep track of the chaos. Zadrozny, a senior reporter for NBC News covering online extremism and misinformation, has built a career investigating the internet's most dangerous conspiracy theories, and it was clear that this historic election—carried out in a pandemic and largely by mail—would be flooded with them.

At the same time, researchers at the University of Washington's Center for an Informed Public* were beginning their own effort to identify and counter misinformation related to the election. Zadrozny had worked closely with the Center's researchers for years, including its director Kate Starbird, and asked if she could be a fly on the wall in the research team's Slack channel. That access would give Zadrozny a more thorough view of the election misinformation landscape than she could ever get on her own. With some guardrails in place, the University of Washington team agreed, and even used the opportunity to publish their own study on journalist-researcher relationships. The partnership paid off, giving Zadrozny and her NBC colleagues a live window into viral rumors as they emerged. One story in particular [catalogued](#) 20 different misinformation narratives that the researchers had identified and explained how these narratives, taken together, cemented the false idea that the election was stolen. That story, Zadrozny says, was “solely”

possible because of the research Starbird and her team were doing. “At the time, we didn't have at NBC the tools and the know-how to monitor the fire hose that was Twitter,” Zadrozny says. “We were just on TweetDeck trying to follow our weirdos.”

One of the biggest impacts that this field has had on the media is that it's given reporters the ability to contextualize their reporting in ways they might not have been able to do before. A decade ago, it wasn't unusual for reporters to write news stories based on a few anecdotal examples of problematic posts. Now, they're far more aware of the limitations of that approach and are producing more rigorous investigations as a result, frequently in collaboration with researchers. “I just don't get as interested in stories about somebody who just used keyword searching to find ten bad things,” says Casey Newton, founder of Platformer News. “I think we're seeing less of that kind of journalism, in general, these days.”

Zadrozny said researchers in this field have also acted as an important “gut check” on reporters' instincts, including her own. She points to one recent [story](#), for example, in which Joshua Tucker, co-director of the NYU Center for Social Media and Politics,* warned her against implying that misinformation changes election outcomes, since research has not found that misinformation has a direct impact on the way people vote. “That was

great,” she says. “I put a speed bump in my own story, but an important one.”

Several journalists interviewed for this report said collaborating with researchers is often more valuable than covering researchers’ papers outright, in part, because they aren’t bound by long peer-review timelines. These collaborations also give journalists a way to report around research findings in order to bring the story to life for readers. And, as was the case for Horwitz’s reporting on Instagram, partnering with researchers on investigations can wind up making those investigations stronger. “It’s not just that there’s an academic name associated with it,” Horwitz says. “The work itself is clearly better and more robust as a result of other people being involved.”

For reporters, one major benefit of working with researchers is that they tend to have access to investigative tools and expertise on investigative methods that newsrooms don’t, and reporters often lean on those tools and techniques in their own work. Bloomberg technology reporter Davey Alba described how Megan Brown, a PhD student at the University of Michigan’s School of Information, proposed [analyzing](#) Community Notes on X to assess how effectively the platform was combating misinformation. Reporters also sometimes tap tools that have been built within academic environments, including [Hoaxy](#), which was developed by Indiana University’s Observatory on Social Media.* Reporters have [used the tool](#), which visualizes how information spreads online, to track, for instance, how a small number of social media accounts churned out hundreds of pieces of propaganda about the humanitarian group the White Helmets in Syria.

Still other times, researchers find ways to share discrete, time-sensitive findings with reporters that don’t require peer review. Zeve Sanderson, founding executive director of NYU’s Center for Social Media and Politics,* cited one [story](#) in which his team provided data to the *New York Times* documenting the “spillover effect” that YouTube’s

election misinformation policies had on other social media platforms. The finding would later appear alongside others in a peer-reviewed paper, but Sanderson says, “We thought that was an important finding to get out into the world.”

Research papers can also inspire reporters to build upon or replicate researchers’ findings. Nitasha Tiku, a tech reporter at the *Washington Post*, said her recent co-authored [investigation](#) into how AI amplifies stereotypes was based heavily on earlier work by researchers including Abeba Birhane, a senior fellow in trustworthy AI at Mozilla; Sasha Luccioni, an artificial intelligence researcher at Hugging Face; and Pratyusha Ria Kalluri, an AI researcher and PhD candidate at Stanford University. Another [paper](#) by researchers at the Allen Institute for AI prompted Tiku and her colleagues to publish their own investigation on the [“secret list of websites”](#) that make chat-bots work.

Tiku, who has been covering the tech sector for more than a decade, said researchers’ work has increasingly become “invaluable” to her reporting, particularly as the technology feeding our information ecosystem becomes ever more complex in the age of generative AI. “As these systems that we write about get more and more impenetrable, and companies get less transparent, and the technology gets more sophisticated, it is really hard to keep tabs on what’s happening,” she says. “You can’t just keep [companies] accountable by being a really shrewd, on-top-of-it individual reporter.”

While most of the journalists interviewed for this report were hungry for closer collaboration with researchers, their appetites for writing about research papers outright varied. For some larger national newsrooms in particular, coverage of research papers has to clear a higher bar, because it’s competing with all the other news of the day. Often, pitches about new research papers are more successful at tech-focused publications or in newsletters that require a constant

stream of fresh insights. “We write three newsletters a week, so we’re always excited when somebody tells us they’ve done a bunch of work that we can write about,” says Newton of Platformer.

Even so, the research still has to be timely. Just as long publishing timelines can undermine research’s usefulness to tech companies and government officials, the same is true for reporters. “A lot of times, by the time good, solid research comes out, the conversation has already moved on,” says Ashley Gold, who co-authors Axios’ tech policy newsletter. Gold noted that she has covered a number of papers published by the University of North Carolina’s Center on Technology Policy* because they tend to align well with the news cycle.

For this reason, when asked what research has been most influential in their reporting, the journalists interviewed for this report repeatedly pointed to work by the Stanford Internet Observatory,* the University of Washington’s Center for an Informed Public* and New York University’s Cybersecurity for Democracy project—all of which have found ways to share at least some of their findings in real time or near real time.

Journalists also raised concerns about the extent to which researchers rely on tech platforms giving them access to data and how the pressure to play nice with tech companies may influence the research itself. Some expressed concerns that the companies’ chokehold on data could steer research agendas toward topics companies are more comfortable with and away from ones they’re not. Another related critique is that academic researchers are often asking questions that can’t be answered by the research design. For instance, even the most well-crafted experiment that lasts only a few months can’t definitively determine whether social media makes political polarization worse, because social media doesn’t exist in a vacuum. As researchers themselves often note, such experiments don’t

account for the cumulative impact of social media platforms over time or the impact of traditional media and the broader public discourse.

That’s one reason why Horwitz says Facebook’s internal research, which he wrote about at length in his book *Broken Code* and in the *Journal’s* series “The Facebook Files,” has been so telling. It tends to focus on discrete product features and examine their effects, rather than taking on broader societal questions. In one [2016 internal presentation](#) he covered, for example, a Facebook researcher found that 64% of the time when people join extremist Facebook groups, they were led there by Facebook’s own recommendation tools. “That’s very useful,” Horwitz says. “It indicates Meta is definitely recommending extremist groups, as opposed to being like: How has social media changed politics?”

Understanding research methodology can also be a challenge for reporters. Journalists tend to gravitate toward research that not only has clear findings, but that also clearly articulates how the researchers arrived at those findings. Craig Silverman, who covers platforms and disinformation for ProPublica, says there’s often a “technical and knowledge barrier” that interferes with reporters’ ability to cover new research. Silverman said he’d welcome dedicated workshops where researchers can walk reporters through their studies and help make journalists “better consumers of that research.”

As in other sectors, journalists’ are highly reliant on personal connections to surface new research. Given the sheer volume of studies that are produced both by academics and by advocacy groups, having a trusted source vouch for research can signal to reporters that it’s worth their time and attention. And yet, several reporters noted that recent changes at Twitter, now X—including the exodus of some researchers and reporters—has made it harder to keep track of the research that their sources are surfacing. “It’s a problem because there was

this community that existed there that was very helpful to finding the researchers and findings that were influential,” says Khari Johnson, a tech reporter for CalMatters.

Justin Hendrix, CEO and editor of Tech Policy Press,* echoed this sentiment. “I feel I’ve lost a little bit of the peripheral vision that I had because of Twitter,” he says.

This gap highlights the important role that well-connected researchers in this space can play in elevating the work of their lesser-known colleagues. It also presents an opportunity for organizations like the recently launched Knight-Georgetown Institute to surface the most compelling emerging research in ways that are less reliant on social media promotion. The Institute’s goal is to uplift the latest research on information and technology and translate it into actionable resources for journalists, policymakers and others. “We are aiming to broaden the pool of researchers whom policymakers at the federal and state levels call on for expertise, and whose research results are used to inform policy making,” says Alissa Cooper, executive director of the Knight-Georgetown Institute.

Despite their critiques of the field, perhaps the clearest evidence that journalists consider this research to be vital is their shared concern about rising political attacks against researchers in this space. Several reporters interviewed for this paper expressed anxiety about the chilling effect these attacks will have on the research community—and the downstream impact that will have on their own industry. “If those researchers are being muzzled because of these lawsuits,” says Silverman of ProPublica, “that’s going to have an impact for journalists, as well.”

PART 4:

How research reaches advocates

In the fall of 2020, Will Duffield, an adjunct scholar at Cato Institute,* was preparing [testimony](#) for a congressional hearing on internet radicalization. He was looking for scientific evidence to support his central thesis that social media algorithms aren't primarily to blame for rising extremism. He found what he was looking for in a [paper](#) by Kevin Munger, an assistant professor of political science and social data analytics at Penn State University, and Joseph Phillips, who was then a PhD student at Penn State and is now a postdoctoral research fellow in at the University of Kent.

The paper challenged the idea that YouTube recommendations drive people into alt-right rabbit holes. By dissecting the existing research on the topic and contributing new data to the discussion, the researchers argued that, in fact, social media doesn't drive people toward extremist content; it just creates a vast supply of it—a supply for which there is already existing demand. “Clearly this suggests a social media reflective of increasing radicalization, rather than one that impels extremism via algorithm,” Duffield wrote in his testimony.

For Duffield and other advocates working in civil society, the growth of research on digital media and democracy has provided the crucial hard evidence they need to make what might otherwise have been purely ideological or moral arguments. For Cato, that argument centers on the idea that

cracking down on extremist speech can also punish innocent speech.

Research is unlikely to radically alter where advocacy groups stand on these issues. But it can help support the arguments they make to lawmakers and to the public. “A huge part of what we do is getting people's hearts pumping and their blood flowing and reaching them on an emotional level,” says Evan Greer, director of the grassroots group Fight for the Future. “But it has to be backed up.”

Recently, Greer's organization has been pushing back against child safety laws, including the Kids Online Safety Act, which would put new limits on kids' access to the internet. As part of that work, Greer says she's repeatedly pointed to [research](#) that was also cited in a report by the Surgeon General, which shows that online communities provide a crucial form of social support for LGBTQ kids. Greer says she could make that point based on her own lived experience, but research gives her the data to strengthen that argument. “The more research we have to back up what we say, the more forceful and righteous we can be,” Greer says.

At a higher level, Duffield says Knight has added substantial value to civil society groups by both hosting and funding convenings on issues related to digital media and democracy. He pointed to a recent [workshop](#) organized by the Knight First Amendment Institute,* which focused on

“jawboning,” a practice in which public officials use informal persuasion and pressure tactics to get their way. The workshop happened to take place the same day that the Supreme Court agreed to take up a high-stakes case related to jawboning, underscoring the prescience of the Knight First Amendment Institute’s gathering. “As much as the individual research funded is useful, I think [Knight’s] role in bringing researchers together, and providing for their collaboration or greater awareness of one another’s work is probably most impactful,” Duffield says.

Despite research’s utility to advocacy groups, the tech industry’s widespread funding of research is generally seen as a problem in advocacy circles. Some, like Tech Transparency Project, have gone so far as to publish a database of academic papers that were funded with Google’s backing. This funding can lead to accusations of researcher capture and wind up undermining research findings, particularly in the eyes of advocacy groups that are critical of the tech industry. “Any academic studies we look at that relate to tech that might be related to something we’re investigating, we always make sure to take with a grain of salt and investigate the funders behind whatever that research may be,” says Katie Paul, director of Tech Transparency Project.

Ben Winters, former senior counsel at the Electronic Privacy Information Center, also raised tech industry funding as an issue, though not an entirely new one. “We’ve always had this in every industry, of course, where there’s industry-funded research or really self-serving research,” he says. “It’s hard to separate the wheat from the chaff.”

Similar to government officials, many of the advocates interviewed for this report said researchers often aren’t asking the kinds of questions that they want answers to, including concrete questions about how platforms’ promises live up to reality. “Academic research is totally important, but it also doesn’t always meet the practical needs of activists,” says Jessica González, co-CEO of the media and technology nonprofit Free Press.

González says she’d like to see more attention paid to tracking, for instance, how effectively tech platforms are enforcing their own policies, how often they’re changing those policies, and whether they have certain fundamental speech policies in place to begin with. Often, advocacy groups are the ones doing this kind of work, but with a fraction of the resources that research institutions have. “At every convening I go to where researchers and activists are brought together, I always hear: What do you need? What do you want? And I’m always like: This is what I need. This is what I want, and it never gets accepted,” she says. “I think it’s because what I need and want is not a brilliant innovation. It’s actually just some dirty work.”

And yet, Duffield points out that there’s good reason to leave that work to advocacy groups, because research that is easily confused for activism risks being delegitimized and framed as further evidence that academics are part of a vast censorship regime. “If that’s what academic work becomes, if that’s what people see academic work as doing, state schools will be defunded for doing it,” he says. “Rightly or wrongly, you will be pigeonholed by that.”

The politicization of research, particularly related to online speech, remains an intractable problem in the advocacy world, too. While activists on the left may have their quibbles with researchers, on the right, activist groups focused on technology are often openly hostile to research that suggests in any way that platforms or the government have a bigger role to play in moderating speech. “The research tends to cluster around one specific vision of what good online speech looks like,” says Neil Chilson, a senior research fellow at the Center for Growth and Opportunity,* noting that he largely agrees with that vision. And yet, he says it can also blind researchers to the fact that other people may see that vision as dangerous. “I think that view is so ubiquitous in the research community that they don’t understand how other people could see that as a threat to their speech online,” he says.

There's no easy solution to that challenge, if it even can be overcome. But Chilson says one step in the right direction would be to create more opportunities for people who hold different views to come together and disagree in good faith. "It's about building a network of trust and building the context in which people are even willing to accept new ideas," Chilson says. "That's not a researcher problem. It's a networking and connecting problem."

CONCLUSION

It's undeniable that the growing field of research on digital media and democracy is already making a meaningful impact on the world in a short amount of time. As the examples in this paper have shown, researchers in this community are regularly shaping legislation and regulatory proposals, influencing tech product and policy decisions, driving national headlines, and galvanizing and strengthening advocacy campaigns.

And yet, while the impact of individual researchers is clear, the collective impact of all of the resources that have been spent on building this field are still coming into focus. Yes, the field has pushed for greater transparency from tech companies—and international laws that require that transparency—but whether more data will actually lead to a better understanding of these platforms and more effective fixes for their many flaws is still an open question. Yes, the field has helped educate lawmakers in Washington, DC, about digital media's impact on democracy, but Congress's will or ability to act on that education has, as yet, proven limited.

This means the field still has room to evolve, mature and prove its worth. Sources across tech, government, media and civil society all said they want more input from researchers—not less—and they want it from a broader diversity of voices. They want true collaborations with researchers who are asking questions that better align with real-world consequences, and they're hopeful that, as the field continues to grow, organizations like Knight will support effective channels for translating research into policy. There is also a broad awareness across these institutions that as new technologies like generative AI threaten to pollute the information ecosystem even further, the need for expertise on how to counter and anticipate those threats will only grow.

Ultimately, that is the goal of this research: Not to enact narrow changes to laws or generate news coverage, but to make the public more resilient to new information crises as they arise. Even those who had pointed critiques of the current field of research noted that the success of the research community should not be measured solely as a function of how it influences decision makers in other sectors. Research that helps inform the public conversation about technology's impact on society is critical all on its own, whether it can be immediately traced to a policy, a product, a news story, or an advocacy campaign. As Dave Willner of Stanford put it, "Socializing a working understanding of why things are broken will help us fix it, because it will reduce useless answers and increase useful ones."

Issie Lapowsky is a journalist focused on tech, policy and national affairs. Her work has been featured in the *New York Times*, *WIRED*, *Fast Company*, *Protocol* and more. She was previously a Knight-funded fellow at the University of North Carolina's Center On Technology Policy.