Cristina López G.:

Good afternoon, morning, and evening, wherever you are. Thank you for joining us here today for Databite #139, "Electionland Misinformation," featuring Ryan McCarthy from ProPublica. My name is Cristina López, Senior Research Analyst on the Disinformation Action Lab (DAL) here at Data & Society. I will be your host, supported by my team behind the curtain: C.J., Rigo, Angie, and Eli.

For those of you who don't know us yet, Data & Society is an independent research institute studying the social implications of data and automation. We produce original research and convene multidisciplinary thinkers to challenge the power and purpose of technology in society. You can learn more about us through our website at datasociety.net.

To begin, I ask you to join me in acknowledging where Data & Society was founded: Lenapehoking, a network of rivers and islands in the Atlantic Northeast we now refer to as New York City. Personally, I was born in the lands of Cuzcatlan, renamed El Salvador, but I'm based on Nacotchtank, now known as Washington, D.C. Today we are connected on the internet, a vast array of data servers and devices worldwide. This system sits on stolen land acquired under the logic of white settler expansion. As an organization, we uplift the sovereignty of Indigenous people and commit to dismantling the ongoing practices of colonialism and all its material implications on our digital world. You'll notice a link in the chat that directs to more information about occupied lands. If you haven't already, use the Q&A feature to share your location.

And now a little more about why we're hosting this conversation and about our speaker today. I've been thinking about the harms brought on by online disinformation for a few years now, and I actually came to this field of research during my previous job at a progressive media

watchdog, where I was deputy director for extremism research. And, at the time, the discussion was mostly centered around grappling with the ways that social media platforms and their amplification powers were used for extremists to organize and for the dissemination of harmful disinformation. I joined Data & Society about a year ago, and I've been working with the Disinformation Action Lab, and our team has been looking, not just at this and misinformation, but also lately—and this is something that I'd love to bring into this discussion later—thinking about the limits of the way that disinformation is defined and the constraints that this definition has introduced and the questions these raise for enforcement.

For the past year, our main focus at the Disinformation Action Lab has been in building capacity for a coalition of civil society organizations: providing them with research-based tools to respond to online disinformation—specifically, disinformation as it related to the 2020 census. Through this work, we've focused, not only on effective ways to reduce the amplification of misinformation, but also on harm reduction: considering first who would be harmed most by t he effects of potential misinformation. In the case of the census, it was helpful to consider issues beyond the true and false binary, and complicate the analysis to include questions about: Whose participation would be suppressed the most? Which communities are harmed more by undercounts? We then considered strategies to help inoculate these communities against potential harmful narratives.

Now, I want to turn it over to Ryan McCarthy, ProPublica's Electionland project editor and reporter, who covers voting, election security, and misinformation. Previously, he was the editor-in-chief of VICE News, where he ran its digital newsroom, built its audio team, and helped launch *Vice News Tonight* on HBO. He has also held editing and leadership positions at *The New York Times*, *The Washington Post*, and *Reuters*. After Ryan's introduction to the project, I'll kick-start a conversation and connect some of the dots between reporting on disinformation and researching it, before inviting some questions from you, the audience. Please add and up-vote questions using the Q&A feature at the bottom of your screen. Ryan, give us the scoop on what ProPublica's Electionland project is all about and some of your key findings.

Ryan McCarthy:

First off, thanks for having me—really glad to be here. It is rare to find the space and time to have these types of conversations during this election season. And I'm sure, to some degree, we're tired of this election, but [I'm] really happy to discuss this pretty serious problem in some of our work. I thought I would just start off with a little information on my background, since I think it actually did inform some of my work here, and is relevant to this discussion. You know, my career as a reporter and editor tracks pretty directly to the social media era: my first editing job was at a local newspaper, nearly 20 years ago, but my editing career really started in earnest a little bit later in the early days of *The Huffington Post*. And these were really the halcyon days of digital media and the Internet—where we were breaking some real news and trying to present digital media in a different way—and, of course, that led us to depend on, sometimes too much, platforms like Facebook and Twitter. And in the early days, that felt like it was solving a problem. It felt like it was a way to make media more immediate, more voice-driven, sometimes more efficient in ways that some of the larger companies couldn't keep up with. And that led me to big digital transformations at places like *Reuters*, *The Washington Post*, and *The New York Times*. I was in *The New York Times* when we rolled out digital metrics for the first time, and that was quite an uneasy transition. But, you know, the last eight, nine years of digital media in journalism—no matter how you slice it—have been reckoning with the forces of digital distribution for good and bad. And, you know, I've seen careers and companies built on this and seen those also ruined quite quickly. And so that led me to Electionland and ProPublica, where I thought we really had a chance to cover—not just the classic beats of election administration and how people go out to vote, which are changing like crazy during this pandemic—but also the ways in which there was a pretty large, and sometimes well-organized, misinformation campaign about what voting is. And, so that led us to this project in mid-July, where we decided to examine, as much as we could, the roots and extent of misinformation related to the election on Facebook.

We did that not just because of Facebook's history on election-related issues, and not just because of the president's continued and false assaults on voting by mail, and claims—even for elections that he won—[that] they were rigged. We did that because we knew that election officials and voters across the nation were really struggling to figure out what the facts were

about a very new and very unprecedented election. In a lot of ways, you know, the way America has voted has changed more in the last year than it has in the last ten or twenty at least.

So we set out in working with First Draft, which is a great social media and misinformation research organization, to really just see if we could dig into the landscape of what misinformation related to how America votes looked like on Facebook. And the findings that we came up with were, unfortunately, not that comforting. Let me just pull some of them up here.

The first big finding is that election-related misinformation—and we're defining "misinformation" as: "false or seriously misleading claims about your vote counting, about the ways in which you can vote, about the outcomes of election," and we also included in here, you know, "broad, unfounded conspiracy theories about stolen elections." We found that nearly half of the top fifty posts, as judged by engagement on Facebook, contained serious election-related misinformation. Examples included: claims that Democrats were stealing the election, misleading claims about how much fraud exists in the voting system, or, you know, claims that were in direct violation of Facebook's stated standards on election-related misinformation—so misrepresentations about whether your vote would count. In one case, we found a very popular conservative commentator literally saying that Barack Obama or Hillary Clinton would burn your ballots in their fireplace. Facebook took those posts down after we identified them, but the preponderance of Facebook stories—the most dominant strain of any single conversation in voting on Facebook, at least as judged by engagement metrics—was one that was seriously filled with misinformation. The second thing we found was that election officials are really fighting a constant battle—and have been before 2016—against election-related misinformation on a local level. So, you know, the California Secretary of State started an entire office and wing to deal with this; a lot of election officials on the state and county level have a direct line into Facebook and Twitter. They say the platforms just are more responsive than they were in 2016, but that they still need to proactively put out information to combat election-related misinformation. And some of that stuff is: "Hey, I went to vote and I was denied. There must be a conspiracy." And some of that is largely conservative networks of commentators pushing out unfounded claims about voting. They, frankly, in order to do their job—which is literally to facilitate voting

and grow voting—they have to fight these battles. And they've found that platforms like Facebook are not always willing partners.

The third thing that we found, which is interesting particularly in light of Facebook's recent actions, is that there's a relatively small amount of misinformation in election-related ads—paid political ads. We took a pretty heavy look at Facebook's ad archive to try to see if the kind of claims that we saw in organic posts—claims of stolen elections, conspiracy theories, burned ballots, you know, claims of *the next civil war* coming from voting by mail—to see if they were really out there and if people were putting serious money behind them. And the truth is we only found a smattering of paid political ads that contained misinformation. And Facebook has been, including recently, pretty upfront that it will limit election-related ads after the election in an attempt to avoid the problem of 2016—but our research says the problem of 2020 is not, in any way, similar to the problem of 2016. This is an organic content problem.

You know, the fourth thing I already hit on. A lot of the misinformation was not just about "Your polling place will not open because Democrats will push an all vote-by-mail platform," which was never true. A lot of it was wild conspiracy claims—and in some cases racialized claims and racialized memes that turned people of color into the face of voter fraud. In Facebook's very narrow definition of what "voter suppression" is and what "hate speech" is, these were not deemed to be violative content—which certainly was not something that civil rights groups and a lot of activist groups found comforting. But we found it was very easy to push these types of claims, and those types of groups were really convenient scapegoats and really unfair scapegoats for these problems. And the last thing I'll say: in addition to Facebook's relatively narrow definition of "voter suppression," they really largely look at things that are saying "Actually don't go out and vote," or material misrepresentations about when and how you can vote. The language is broad but they are actually applied pretty narrowly: and if you press Facebook on what voter suppression means, you'll find that their definition is wildly different from [that of] activist groups—including civil rights groups—how they define voter suppression.

The last thing I'll say is a minor finding that we had which is proven to be true from Facebook's history: that by the time they act to stop some of this stuff—to either take it down, or slow the

spread of misinformation on the election—it's often too late. You know, so we found one claim from a purported Trump voter saying that her vote was denied. It turns out it was completely made up—the state that it referenced was not even holding voting that day. Facebook and its fact-checking partner PolitiFact checked that piece of content, but by the time a fact-checking label labeling it false was applied, the video had more than 3.7 million views—it was shared more than 170,000 times. And that's just what we know of from CrowdTangle, which is Facebook's public analytics tool.

So, I think I'll leave it there. Our work really was based in a midsummer's view on what Facebook was doing about election-related misinformation. They've taken a number of actions since, which we can talk about. But I think the worrying thing is the problem seems central and sort of endemic, and the caveat on this research is, of course, we have a very limited set of data on what the world of Facebook is—and it's largely through CrowdTangle. The most engaging Facebook—which is, you know, a metric determined by a sum of shares, likes, and reactions—tells us that Facebook has a serious problem with election-related misinformation that has, in fact, I think poisoned the well on its ability to get reliable election-related information out there. And I think one telling thing is [that] after we wrote the story, they did announce, and have announced, a series of linking and labeling policies, where they actually push people away from Facebook into state or local voting sites to try to get them information on that. And we can debate whether or not that's sufficient—I think it's good to see that they are doing that—but I do think it is, in part, a mission of the problem, that they are actually saying: "Facebook users, step away if you want to find real information about voting." So I'll leave it there, if that's good.

Cristina López G.:

Thank you so much. I think that's amazing table-setting for conversation. I actually want to start our discussion specifically with social media platforms and the related issues of responsibility *and* accountability. Because we've often heard leadership at platforms—whenever their feet are on fire—say that they don't want to be in the position of editors, or that they don't want to be in the position of being the arbitrators of what's true and what's false, and they say their

responsibility really stops at facilitating the right to free speech of other people. And so when you actually consider the harms of misinformation that come from the amplification and reach—and how much of that is actually facilitated by platforms—we actually land back on the question of accountability. And, oftentimes, excellent reporting, excellent election reporting exists, but it really cannot compete, in terms of speed and reach, with just the way that election disinformation travels. So, I think I want us to talk a little bit about the responsibility of platforms in this scenario, and maybe talk a little bit about what changes they did after your reporting—maybe that could shed some light on the ways that we can hold them accountable.

Ryan McCarthy:

Yeah. So, I think it's good, I'll just start with what's happened since July when we dropped this article. On the day we published it, they did announce a broader linking policy, essentially linking to nonpartisan voting-related sites underneath politicians' posts. Since then they've sort of taken a number of escalatory actions that are similar to that: they've banned militarized language in talking about poll watching—you saw the president's son call for an army of poll watchers—and they were late to that, and those messages were certainly circulated widely before they did that. They have banned election ads after the election. They have announced policies to put "non-neutral labels" pushing back on things that are actually untrue. But in other cases, they have yet to take serious action: the post that President Trump put up imploring people in North Carolina to go vote twice—I think their reaction to that was not what election officials would want—and, you know, it certainly traveled a lot farther than, say, the [North Carolina Attorney General's warning](#) saying voting twice is illegal.

[20:12]

I want to go back to the responsibility issue and the accountability issue because I think this is actually pretty important. The idea that social media companies are just a mirror to society and just a reflection of their users' interests is one of the original lies about social media. You know, I think we need to make sure that we understand as users that the people who build these products are making affirmative choices—and those affirmative choices affect what you see. And

we can quibble about whether or not Facebook is a media company or some weird, new platform/publisher, but the truth is the impact they have is similar to that of a publisher. They are making choices that will incentivize certain kinds of actions and penalize others. So I think that is the first ground rule we have to think about when we think about Facebook and Twitter and YouTube. The second thing is—and I am not the first one to say this—but I think another one of the original misconceptions about social media is that these are social media platforms first. Facebook, in any case, earns 95% of its revenue from advertising—this is an advertising business, you know—and I think we don't look at Exxon and call it a people-moving business, you know? I do think there is a bit of marketing spin when you imbue an advertising platform that lets you have free social media features with democratic qualities—or qualities that we should normally invest in civic institutions or in other things that contribute to democratic society. So I think holding these platforms accountable as businesses, advertising businesses first—and ones that are making choices that we know can affect the way people see the world and can affect democratic institutions and sometimes undermine them—I think is key to keeping them actually accountable. Because if we let platforms see themselves as central and inseparable from democracy, I think they've won.

Cristina López G.:

I think that's a really, really good point. And a lot of people have been into the narrative that they are not really editors or that they are neutral to whatever happens and reflect what's going on. And I think that, largely due to how much of 2016, and the story that came out of that election, was centered on disinformation, they came out this year, a year in advance, specifically—Facebook in October of 2019—and they built a number of election integrity policies ahead of this election. And they were looking at what you can call the "low-hanging fruit:" [policies] that were seeking to reduce inauthentic behavior, they were—very narrowly, but at least—focusing on voter suppression. But I kind of want to lead the conversation here towards the limits of the disinformation definition, because, as we understand it, it is: information that is false, that has the potential to harm, and that is spread with ill intent. However, specifically in the context of the election, you might have seen that the information that we're seeing

amplified—it goes beyond necessarily just the true or false binary. You see entire narratives, entire storylines—some of which are playing up true elements—but that have been completely decontextualized in efforts to manufacture doubt. Like, weaponizing electoral reporting, but in order to hurt the perception of legitimacy in the electoral process; in different words, the biggest issue often is the framing and bad faith that occurs. And I'm thinking about the rampant speculation that you mentioned in your intro about potential voter fraud via absentee ballots, or accusations that a lot of practices that are actually legal in a lot of states and that have nothing nefarious about them—like ballot harvesting—how those are equated sometimes, in memes and content, as equivalent to ballot-box stuffing. So, I kind of want to expand a little bit on something that I think would be really interesting to consider. From your position as a journalist covering the election, how do you grapple with the risk of having your story being amplified for nefarious purposes? Or decontextualized to advance narratives that you actually have very little power to combat, and that—seen with the very narrow scope of what the platforms are considering to be election integrity policies—would likely go under the radar?

Ryan McCarthy:

I think that's such a really good point: both about the idea that a lot of the problematic content is not just binary true or false, but blends some serious false elements with some stuff that's actually pretty true; and that the idea that as a reporter—even if you're the most straight-down-the-line person—your information will be weaponized. On that last part, you know, last week in partnership with *The Philadelphia Inquirer*, I wrote a story about local election offices being inundated with duplicate mail-in ballot applications—largely because Pennsylvania is growing its vote-by-mail operation in an unprecedented fashion—voting groups are sending out unsolicited ballot applications to people not knowing whether or not they've received one prior, and, in some cases, state websites are inaccurate or just totally misleading and causing people to request more ballot applications. And so this is not a story which says our election is doomed or voter fraud is rampant. This is a story about administrative failure—or administrative struggles, right—because Pennsylvania's election, I want to be clear, has not failed. But election offices were, in some cases, working 24/7 processing duplicate

applications—90% of some of the applications that they were processing were duplicates—which is just a silly sort of administrative struggle. But, as soon as we put that out, our story got twisted by conservative commentators who—instead of "nearly 400,000 mail-in ballot applications being rejected by Pennsylvania"—took the "applications" out. The difference here is, when you want to vote by mail, you usually have to fill out some forms proving to your state or county that you're eligible to do it—but usually it's a rubber stamp—especially for states with no-excuse voting. But conservative commentators removed that "applications" part of it, I think intentionally, to try to cast doubt on voting by mail. And those things largely weren't fact-checked because the fact-checking operations for the platforms just don't have the manpower to keep up with everything. And they circulated widely—thousands and thousands of retweets—and there's sort of nothing, really, you can do about that. But I do think the problem where you have a claim that circulates widely online, one that says "All vote by mail is fraud" and then mentions one legitimate, infinitesimally-small, incident of mail-in fraud is a problem of proportion and framing and contextualization that the platforms really, really struggle with. And I think that the only way to solve that is to either do a network analysis, as other people have suggested, or look at some of the dominant narratives: look at the meta-narratives that run through multiple conversations and multiple commentators to try to say: "Is there a way we can turn down the volume or disincentivize stuff that has a significant misleading element to it?" And I think that takes the conversation beyond individual takedowns and more [towards] looking at the health of social media systems as a whole.

Cristina López G.:

Right, that's a really good point. Maybe turning it into an analysis of: Who is harmed the most by this? Who would find it really, really hard to vote if presented with what is likely to get the most engagements? And I actually want to focus a little bit on the harm question, because a huge causality of the increasingly fragmented media landscape—and of the business model dependence on platforms for traffic—has been local media. Local newspapers have disappeared at alarming rates—there are areas that are actually news deserts now, and they are all over the country. And I read somewhere that you can say that the local reporting well is drying while we

deal with a deluge of this information that's created at the national level and that has huge funding and that is often attempting to pass as local information, as we've seen in recent reporting. And I wanted to bring this up because I really want us to focus a little bit on the harms and kind of move away from the abstraction that it is "truth" or "reality" that is harmed by political disinformation. Because I think it's important that we learn from this and, leaning into your reporting and your knowledge of this beat, we should focus on the victims: Can you talk a little bit about who is harmed at the local level? Let's discuss, for a little bit, those local election officials and those voters without enough quality information that might find it really, really hard—or might be too terrified—to go vote.

[30:40]

Ryan McCarthy:

Just to zoom out on the problems: the stat that flies around a lot from some of the media analysts over the last 12 or 15 years is that newspaper employment has fallen faster than employment in coal in America—in coal mining. Which is a really sad state of affairs for the ability of local communities to get reliable and factual information about the basic functions of government. And, prime of those concerns is really election-related concerns—especially in a year where, you know, some states are growing their share of vote-by-mail ballots from 4% to close to 50% overnight. It's really crucial that people sort of trust the information coming out from their local community. And the truth is—there really isn't another outlet for that other than newspapers right now. I think there was a thought that social media companies would lead to new business models that could fill the local reporting gap: if that's going to happen—and I'm dubious that it will—it hasn't happened yet. So when you talk to local election officials about the problem of misinformation, about the problem of saying: "Look, I have to spend millions buying PPE and printing mail-in ballots, and I don't want those to be wasted." And those folks will say this to you regardless of what side of the aisle they are on. And the local news—to the extent that it hits their community—gets swept up in these national mega-narratives. So, you know, in Paterson, New Jersey, there was a case of election fraud where some local candidates for office there seemed to engage in some ballot harvesting. Ballot harvesting is legal to a point—you can

collect other people's ballots in New Jersey, as long as you're not doing it for more than three people, and as long as you're not an unauthorized sort of transporter of those—but it seems there was some illegal behavior in here on a limited and local level. But, of course, there just really aren't enough local news resources to really delve into that in-depth, and so, there's really only been surface explorations of the problem in Paterson, New Jersey. And that's what happens: folks who have, maybe, preexisting views about voter fraud seized on this and tried to turn it into a national story—without the type of scrutiny, detail, context, and nuance that would come if you had sufficient local reporting resources to handle it there. And that's happened with a number of cases of election-related misinformation—and, obviously, that problem is not constricted to election-related misinformation.

And so a lot of times you'll have national news organizations, in some cases like ProPublica, parachuting into local communities to fight this problem. And I do think we can get tied up in knots about how or whether social media companies have destroyed local news—I don't think that's, frankly, a useful discussion at this point. I do think, in information deserts or news deserts, social media companies just have the potential to do so much more harm, because they are not competing against any other local news resources.

Cristina López G.:

That's a really, really good point. And, because I want to turn it to a little more positive or empowering side of the conversation, we've talked a lot about the role that platforms have to play and how they should definitely have policies against disinformation: ideally, we'd like to see more consistent enforcement, ideally, we'd like more transparency—so we know what we don't know, and we know whether the measures applied so far are working, or if they are just another PR stunt. But one of the things that I think we don't talk about enough in discussing a problem that is a *networked* problem, like disinformation, is how it still is a networked solution, and everyone has a role to play. And you kind of eluded to this earlier in your introduction when you mentioned the partnership that ProPublica had with First Draft, and there is definitely a role for scholars, for civil society, for journalists—even journalists that are spread super thin in the election—in combatting disinformation. Let's talk a little bit about how is the experience of

partnering up with civil society and independent researchers to try to tackle this problem? Because case studies like this are incredibly valuable for other situations of disinformation.

Ryan McCarthy:

Absolutely. So, I'll talk a little bit about what ProPublica has done and is doing in the future, and then try to get at what's left on the table—at least from my view. So ProPublica runs a project called Electionland which is now in its third election, I believe—started in 2016—in which local newspapers can partner with us and we get tips from a massive group of civil society organizations, election protection organizations, and we feed the tips out—after checking them—we feed those to local newsrooms and they report on voting problems in something like real-time around the election. This year we've expanded Electionland a little bit with some grants that we've offered to local newsrooms to do investigative, more long-lead, stories about elections. We partnered with the *Tampa Bay Times* this year, *The Philadelphia Inquirer*, Georgia Public Broadcasting, and an organization called Wisconsin Watch (The Wisconsin Center for Investigative Journalism) to do that. And, in investigating the election story, we worked with First Draft—which is a really fruitful partnership—and they have access to data that, frankly, would take us a lot longer to look at.

But in terms of what's left out there in solving a networked problem—and I think it's an important point that's not about individual takedowns or individual bad actors—this is about a *networked* and sort of *incentive-based* problem. I think, if this is a network problem, we need to know what the network is. So we need to not just know some basic information about engagement—we need to know about what kind of posts are actually leading to takedowns. Facebook is not transparent about its takedowns: it brags about some, it does others under the cover of night. We need to know about the ways in which different actors share or coordinate which are, in some cases, counter to Facebook's rules about authentic behavior. We know that memes can spread really quickly—[but] we don't know how that spread actually works, because Facebook won't let us see any of that data. Twitter won't either. You know, one idea, there, is to try to look at: "Are there Facebook pages or Twitter pages that are sharing the same content—identical content—with absolutely no alterations within the span of a minute?" or

something like that. That, I think, would let us get a little bit clearer about what problems we are actually trying to solve here. Because you could hire an infinite number of content moderators to solve some of this—frankly, I think Facebook and other platforms should hire more to live up to their actual promises—but until you know the sometimes-inauthentic and sometimes-just-odd ways in which information groups together through networks and flows, and until you know the full influence of that, I don't think you can really solve this problem. So, if platforms like Facebook *were serious* about this problem, I think they would open up to that idea that this is not a serial problem, this is a problem where you need to change the incentives and the real flood of information—and how that's propelled through networks.

Cristina López G.:

Yeah, that's a really good point. And we learned very similar lessons through our work in partnering with civil rights organizations—combating something that had, I would say, not even close to the volume that we saw with election disinformation—but disinformation that was targeting the 2020 census. In some ways, there's a lot that you can do in terms of petitioning social media platforms and suggesting what you think could be "best practices" in terms of how you want to see those civic integrity policies enforced. But, at the end of the day, you don't have the capacity to figure out where the coordinated behavior is coming from or who is even reaching when you lack so much transparency in terms of what happens under the hood at social media platforms. One thing that we found was, at least, within the realm of what you can do—from the civil rights or civil society organization side—was to to think of their audience and to try to flip the analysis into: "Who would be harmed the most by this narrative in the wild, or this potential strand of disinformation?" And to prepare in caring for *your audience*, for your little group online, and trying to figure out where this little group gets their news—and trying to protect that. Because, trying to combat it from the other side, it does seem [to turn into] a game of whack-a-mole really quickly.

[40:50]

Ryan McCarthy:

One thing on that before we move on which I just wanted to mention: you know, a lot of the political science research and just the academic research says any barriers that you put up to someone voting or getting information about voting can actually turn them off and decrease the likelihood that they vote. So I think platforms operating from that perspective would be really useful, and I'm not sure they are operating from that perspective.

The second thing is: platforms should care about the overall institutional creep and, I would say, belief in the system long term. And this is just something that political scientists worry about with election-related misinformation—even if some of it is borderline factual and not worthy of a direct immediate takedown. Stuff that proliferates on platforms and constantly fights against the reliability of the electoral systems or the reliability of our civil institutions—that may not have an acute immediate damage to people's ability to turnout—but, over the long term, that's going to degrade our belief in society. And I think that is a slow creep problem that platforms really should be incentivized to fix. And I'm not sure they are even thinking about it, to be honest.

Cristina López G.:

We thought about it in the same way with regards to the 2020 census. When we think about things in just the tight binary of "this is true, or this is false," that does leave out the slow drip, I would say, that has lasting impacts that we can't really measure right now, and that is: people's perception of legitimacy in the system, and institutionality, and core values that make a democracy work. And I agree with you, I don't know that it's being thought about from that angle within social media companies—and even in some newsrooms, I don't think.

Ryan McCarthy:

I think that's totally right about some newsrooms, as well. That long term problem—it would be great to get some thinking going on that— about our information systems. Because I worry about that all the time, and I have no idea what I can do to effect change on that.

Cristina López G.:

Let's turn to some of the questions that have been left on the Q&A. One of them says they

understand the need for social media companies to be accountable to stop the spread of election disinformation, but: "I am also worried about the possibility of allowing social media companies to wield too much power over the content shared on their platforms. For example, if social media was popular in the early 2000s, what would have stopped these companies from censoring skepticism about the existence of WMDs during the Iraq War?" I do think about this a lot.

[McCarthy laughs]

Ryan McCarthy:

 I think—look, I think it's interesting. If you look at some of the ads that companies like Facebook—well, Facebook in particular, actually—runs on places like *Axios* or media sites, they actually are pretty open that they welcome Internet regulation. And, certainly, they are spending money to lobby on it. But I think this is where we need to realize that Facebook has over 2.7 billion users: to some extent, as I said, it is not a mirror to society, but if we don't put our own values on top of it, Facebook will insert its own values in there. I do think we should be worried about mass censorship—I just don't think that's the problem right now. And we need to plan for that possibly becoming a problem, but I think reliable information—the preponderance, or the relative lack of reliable information on platforms—is a more urgent problem to what's going on right now. But I also—look, Facebook is so vast that it's hard to generalize. I think in America, the problem is about reliable information. In other places,state-sponsored inauthentic media companies may be a bigger issue; hate speech may be a bigger issue. I know Facebook leadership tends to take a very Western-centric look at the biggest problems and tends to underinvest in problems in the rest of the world out there. So, I don't know if that answers the question. I do think they are both serious problems—but, at least in America right now, the acute problem is misinformation, in my view.

Cristina López G.:

I also think that the fear of the censorship issue is often what allows a lot of social media companies to let themselves off the hook. You brought up something that is very, very near to

my heart—because I am from El Salvador, where social media companies are used *by the state,* often, to have amazing reach without any competition. And, honestly, it is the fear of censorship, [that] keeps social media companies from actually intervening or dismantling entire operations of inauthentic behavior which are being used to suppress actual journalism. So, it is an issue that is worth paying attention to without letting social media companies use that narrative to their advantage.

We have another question: "It seems like many journalists view their job as essentially exposing problems and points of failure but, as you mentioned, this kind of reporting can easily be weaponized to undermine the legitimacy of the election. Would it be useful for journalists to, instead, think about their work as providing alternative narratives that can combat harmful ones like voter fraud? Or is "Man Votes by Mail Legally and Easily" just never going to be a story?

[Laughs]

Ryan McCarthy:

I think this is a real problem, right? I go back to the classic definitions of journalism as expressed by the people who have been around forever and are, you know, more accomplished than me: Seek the truth, first of all. Hold power to account, second of all. And third, explain the world. And I think depending on those three qualities—depending on what percentage you dedicate to each of them and what proportion—you come out with different answers. I think the truth is: journalists kind of need to use all of those. And the idea that we should be afraid of reminding folks that our system works is kind of silly, and probably reveals the bias towards the new or the salacious, or stuff that's more accountability-focused in journalism. I think it would be great if *The New York Times* or ProPublica or something like that ran an article that said: "There were some problems with the election—but it largely worked." And I hope that they do run that, because I don't think you can get around the idea that you have to fight against the growing powers of misinformation and propaganda in society right now. And my guess is there is probably a 50-50 chance that *The Post, The Journal*, or *The New York Times* runs a story which essentially says: "Our election system basically worked." I hope they do—and if they don't,

we can write some letters.

Cristina López G.:

I would imagine it would go viral. But, from the perspective of journalism, you can definitely remind folks that, you know, voting by mail is boring and easy—and it works.

You mentioned earlier—this is another question—that Facebook has a really narrow definition of "voter suppression" misinformation. How do you talk to organizations or people with similar conceptions of misinformation about expanding or reframing their understanding of the issue?

Ryan McCarthy:

I think just going back to what I was saying—and this is not just me saying it—I think the civil rights folks who participated in [Facebook's civil rights audit](#) earlier this year and some of the groups that were leading the boycott: they argue that Facebook takes a too-narrow definition of voter suppression—such that it's actually suppressive. And I think this idea of stuff that really calls into question the basic legitimacy of our electoral system fits into that definition. Facebook thinks that's allowable political speech; and, on a sort of black-and-white principle level, I think they are right. But the problem becomes when that kind of political speech dominates its platform, and is spread quickly, and in some cases in a coordinated fashion, and in some cases just really takes over the conversation. So, I think when you're talking about those meta-narratives, I think you really have to remind them of a sense of proportion—again, going back to the voter fraud issue. It is vanishingly rare that voter fraud happens in any form, including in voting-by-mail, but fraud in voting-by-mail is *modestly* more common than it is in-person election fraud. That said, you are more likely to get struck by lightning than to ever be a perpetrator of vote-by-mail fraud. And so, I think proportion and context is really important when you're talking about this sort of stuff, but, you know, again—to hint at something Cristina said—proportion and context don't travel well on social media. In fact, stuff that doesn't have that travels faster.

[51:24]

Cristina López G.:

Again, going back to the issue of: "Who is harmed the most?" If social media platforms took the perspective of the harmed group more often, rather than the perspective of the group that is likely to [spark] backlash and mount a social media campaign about censorship when policies get enforced, I do think that a lot of the damage could be curtailed, at least.

This looks like an easy one, I guess: "What public policy solutions would you suggest to stem the flow of electoral misinformation?"

Ryan McCarthy:

Super easy. *[Laughs]* I don't know. Look, I think I'm a little bit uncomfortable proposing policy solutions as a journalist. I think the concepts behind it, though—some of which we've already touched on—first off, platforms like Facebook need to live up to their promises. If they say they are going to ban militia-related content, you know—profiles with militia in the titles need to be taken down. If they say they are serious about state-sponsored action in other countries—they should get serious about that. There is increasing evidence, and continued evidence, that they are just not hitting their own goals—and that's not to say they should catch every problem on their platform—I think the biggest complaint is "Do more of what you say you're going to do." But in terms of actual policies, I think we need to know about the health of the network, going back to what we were saying earlier. And I think we need some actual metrics and research done by third parties about how healthy Facebook's discussion on certain topics is, with a particular mind to things, like Cristina says: "Who could be most harmed by this?" If I have never voted by mail, and if this is my first vote, or I'm in a voting group that has faced historical barriers and very direct intentional barriers to voting, how could I be harmed by this kind of information?

And I think approaching it from a regulatory standpoint like that would be good,but I don't think we can do that without seeing more of what is actually happening on there. Kevin Roose of *The New York Times,* with which we're all familiar, has this automated bot about the most popular-by-engagement posts on the platform in any given day—I think that's good. Facebook has, at times, responded saying that "Engagement isn't reach." In my experience as someone

who's helped run digital brands, engagement and reach are highly correlated. Not perfectly, but highly correlated. That's a basic question we need to know: if the discussion on Facebook is a lot healthier than we're seeing—or if there are other ways to measure this—we should know that. The evidence, both from the impact that we can tell in talking to people and from reporting—election officials, voters—is not good. The evidence on what we know of engagement metrics is not promising either.

Cristina López G.:

I agree with you. I think that in the issue of platforms being the ones advertising—you mentioned *Axios*—but platforms are advertising and spending a lot of money on a campaign that basically asks to be regulated: "Come and regulate us, then," "We welcome regulation." But it is really hard to do when you don't have data that would allow you to figure out which solutions are the ones that would work. We know very little—in terms of the solutions that they have enforced, themselves, so far—we know very little about how successful they are, because we don't know if presenting voters with neutral information or links to a different neutral nonpartisan site that will show the right way to vote—we don't know the click-through rate. And we don't know if folks see those as annoying pop-ups, or if they see anything that would divert them away from the platform as a valid or unmeddling way to deal with an issue like an election.

So I do believe there is regulation that would ameliorate some of the harms of the problem. I don't know that platforms have been forthcoming enough in presenting enough information to know that the regulations that could be available are those that would also be successful and valuable.

We actually are almost at time, so I would like to leave some time for you to tell us if there's anything that you'd like to leave with us today—even if it's a wish list—or anything that you'd like to see happen in the little time we have until the election.

Ryan McCarthy:

Yeah. There is this cliche called the Election Administrator's Prayer, which is: "Whatever

happens, don't make it be close." I know a lot of voters feel that way and a lot of election officials feel that way. I think there is a good chance that we avoid a misinformation disaster, this time around, in an election-related disaster, this time around. I guess my wish list would be that both of them involve, like, big, messy, varied institutions: one which is a classic civic institution, and one that's sort of become one. I hope that we solve the problem a*s related*, right, because I think we could have the most fully functioning election system in the world—and I would argue we don't just yet—but we could have that and still have these strains of propaganda and misinformation swirling around, and we wouldn't trust it or believe in it, or rely on it to yield fair, democratic outcomes. I think one of the lessons of the social media era, just to put a cap on it, is: the information problem is a civic problem. And I hope that whatever sort of policy discussions or regulatory discussions or individual choice discussions happen on the voters or users of these platforms keeps that in mind: that there is, in our digital spaces, a civic duty that we all have and the platforms we rely on have.

Cristina López G.:

And I would also hope that we move a little bit towards considering the perspective of who's harmed, and also considering the implications of the slow drip: what that could be changing in terms of how the doubt that is being manufactured. Or just planting that seed of doubt: how that could affect the perception of a lot of potential new voters who, perhaps this is their first election—and they might not feel very encouraged in continuing to participate in the system. And you might be just leaving the choice of deciding who represents us to the few people who are not yet cynical about the system—who haven't seen enough memes about the election—to get that perception of legitimacy eroded. So, I would hope for more conversations to be had at this level—not just *within* social media platforms—but I know that definitely civil society is having them. I know that scholars have been having them for a while, as well as newsrooms. It has become a little bit of a new challenge, and I do hope these conversations continue outside of just the election.

And I want to thank everyone for joining us, and thanks again to Ryan McCarthy for sharing your expertise today. And for the time you generously shared with us in preparing for this

discussion to be the success I believe it's been. Check the chat window for hashtags to keep this conversation going on social media. You know where to find us and reach us. Please complete the short three-question survey before you leave. And just a reminder that this event recording and the resources will be posted on Data & Society's website soon, that is datasociety.net, where you can also find our research and programming. So, we hope you'll join us again for a future program. Thank you so much, and take care.