Transcript: 1/25/2021 (checked by Malvika Jolly)

Databite No. 141: Digital Technology and Democratic Theory

Lucy Bernholz, Robyn Caplan, Rob Reich, Seeta Peña Gangadharan, and Archon Fung

Recorded: 01/21/2021

[4:39]

Robyn Caplan:

Hello everyone and welcome to Databite No. 141! We are celebrating the launch of *Digital Technology and Democratic Theory*, a new book published by the University of Chicago Press. My name is Robyn Caplan and I'm a researcher here at Data & Society. For the visually impaired—we're trying a new thing—I'm a white woman with long brown hair so you can form a better image of who is speaking with you today. I'll be your co-host, supported by my team behind the curtain—CJ, Eli and Rigo—as well as our partners at Stanford PACS, co-producers of today's event. I'm also one of the authors of this volume today, though, so I'm going to be doing double duty and I'll be showing up later on the panel as well. I have the pleasure of sharing the stage with Lucy Bernholz and Rob Reich, two of the editors from the collection, as well as two contributing authors, Seeta Peña Gangadharan, and Archon Fung. A link will be posted in the chat where you can read more about our guests.

For those 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.

The Stanford Center on Philanthropy and Civil Society is a global interdisciplinary researcher and publisher of the Stanford Social Innovation Review. Stanford PACS develops and shares knowledge to improve philanthropy, strengthen civil society, and address societal challenges. By creating a shared space for scholars, students, and practitioners, Sanford PACS informs policy and social innovation, philanthropic investment and nonprofit practice.

We're going to be spending the next hour together—actually I think a little bit more than an hour—so I want to all get grounded. If you're joining from a computer, use the features at the bottom of your screen to participate. So you can ask and upvote questions via the Q&A function. You can use the closed captioning function for subtitles. You can view links and prompts in the chat window—and we tend to share a lot—so you can make sure you're staying updated there. And we also want to note this event is being recorded and will be shared afterwards and you can follow the conversation on Twitter with the hashtag #DigTechDemTheory and we'll share that hashtag in the chat as well.

We like to start all of our events with a digital land acknowledgment. Data & Society was founded on the ancestral territory of the Lenape people—a network of rivers and islands in the Atlantic Northeast we now refer to as New York City. Today, we are connected via a vast array of servers and computers situated on stolen land. We acknowledge the dispossession of Indigenous land by the data-driven logic of white settler expansion, and uplift the sovereignty of Indigenous people and data. We commit to dismantling modern-day colonialism and its material implications on our digital worlds. You can visit native-land.ca and use the Q&A feature to share information about whose land you are joining us from.

So now I want to turn this over to one of the editors of the volume, Rob, to get us into the context of today's event. Hi Rob!

Rob Reich:

Hello, everybody. Thanks so much Robyn. And for the visually-impaired: My name is Rob Reich; I'm a faculty member at Stanford University; I'm a middle-aged white male professor. Stanford students—I think most students—sometimes refer to themselves as “rising” sophomores, or “rising” juniors, “rising” seniors—I like to refer to myself as a declining professor, declining egghead. I'm happy to be here today despite that self-description.

I want to offer just a very brief framing idea for this entire conversation. Of course, on all of our minds these days is the question about whether it was appropriate for the various social media platforms to *deplatform* now-former President Trump? Whether or not the banning of Parler to cloud access or to the app stores was indeed a good decision? That is just the most recent version of many of these questions that have beset all of us for *at least* the past decade.

And what I want to suggest that this book allows for us to do is to bring a certain type of sober, mature framework to thinking about questions at the intersection of the digital technologies that have revolutionized our lives and the democratic societies in which so many of us exist.

We want to bring a non-polemical orientation to the conversation in which we're not advocates of "liberation technology" or techno-utopians, but neither are we committed to the idea that Facebook is a destroyer of democracies and that tech companies in Silicon Valley and elsewhere do nothing more than hijack our attention for bad.

One quick note about how this volume came together—because I think it's important to register this—many edited volumes are nothing more than a stapled-together collection of workshop or conference papers, or occasionally a kind of celebratory volume of a particular scholar.
By contrast, this volume was done in a different way: we editors spent six months putting together a cast of contributors who then assembled for an eighteen-month-long series of workshops. We met on three separate occasions for multiple days across disciplinary lines with the idea that we would try to explore where different disciplinary orientations—from political science to philosophy to history to sociology to communications and media studies, and including, as well, engineers—where they were common concepts and common ideas at these intersections between digital technology and democratic theory. There was no editorial agenda about what conclusions to reach, but rather we sought through these workshops to come to some common views. This book is not, however, a manifesto. It represents the different conclusions of different scholars that are the product of these cross-disciplinary conversations.

So I hope with that you get a sense of the volume as a whole, and we'll get a much closer view of a few of the contributions to the volume now. So I turn the floor over to my colleague Lucy.

[11:25]

Lucy Bernholz:

Thanks, Rob. Thanks, Robyn. And welcome everyone. I'm Lucy Bernholz, I'm a middle-aged white woman with gray hair, red glasses and there is a beautiful picture of redwood forests behind me. And I'm calling in from the land of the Muwekma Ohlone people who are the original stewards and caretakers of the land that Stanford occupies today, and pay my respects to their elders past and present. Indigenous people are present and thriving despite our occupation of their ancestral lands, and I acknowledge their existence and continued existence.

Let me pick up where Rob left off about the book. You know, books take a long time to write and sometimes you get lucky as to when they come out. A lot of the topics and themes, as Rob mentioned, that we were wrestling with for the last few years are now very much on people's minds. I'm a historian by training. I don't think of myself as a democracy theorist, but one of the things that we tried to do with the way we went about putting this book together, and the final volume, is to think of it as an examination of these intersections between democracy theories and digital technology—each of those spaces taking each other very seriously and calling for expanded and intersectional approaches to theorizing about democracy in the digital age.

I think of this as: "No more democracy theory without digital dependencies."

[Laughs]

And then, when you think about the digital technology side of things, there are a lot of different logics for the digital technologies that we use and are surrounded by. The mass market ones are

largely derived from what I think of as an extractive capitalist logic and a surveillance state logic, but there are alternative logics out there and there are technologies built on them. And, particularly, as we think about where we are today—there is a lot of talk about new infrastructure— I think it's a particularly opportune time to engage with the potential and possibilities of different logics for digital technologies.

The book is by no means comprehensive in scope. Actually, as I look back at it now, I'm not sure that one volume could ever be that. It covers topics from governance by distributed infrastructure to the role of associational life; the role of journalism to the kinds of new skills we all need to thrive. The particular authors who joined us here—Seeta, Archon, Robyn, and myself—look at the public sphere, corporate governance, silence,e exclusion, and refusal. and the role of community power, but the book covers a broader range of topics than that.

And finally, it lays out a research agenda for more work—for scholarship that we hope the book will inspire. And, particularly, the value of interdisciplinary scholarship and also a little bit of a sense—I don't think of myself as a democracy theorist now but I've played one on TV for the last several years and, certainly, my own work is much richer for that—and I think there's room for people from many disciplines in that conversation.

[14:56]

So I'm going to ask each of our panelists to take about six or seven minutes to introduce the substance of their chapter—to give our listeners something to chew on. I ask you to jump right in with the Q&A and we'll be turning to the Q&A after a couple of rounds of questions from me.

So let me start with Archon. Archon, you and your coauthor Josh Cohen, in your chapter of the book, take on the role of the public square in democracies or the public sphere. What aspects of digital technology come into play in your chapter and in this idea of the public sphere? And what big questions did you address in the chapter?

Archon Fung:

Great. Thank you very much, everyone, for joining us this afternoon and thanks, Robyn and Lucy and Rob, for bringing us together and having this great discussion at this moment in which this topic is, I think, on a lot of people's minds.

So our main question—we began the chapter—what motivated us to write the chapter was: To think more clearly about all of the many criticisms of social media today. Especially criticisms that social media is bad for democracy, right? And so, what we set out to do in this chapter is gain some more clarity about what exactly these complaints are—what are they about? And so how we do that is that we aim to do two things:

First of all, develop a kind of normative yardstick for what a good functional public sphere—that is a media sphere of news, information, and communication—looks like in a democracy. If you're in the Middle Ages or in the 20th century or in the 21st century: What are those normative standards and values? And we asked: Good for what? Good for democracy is our main question. Right. What kind of public sphere is good for democracy—because we believe that a successful democracy requires supporting informational conditions. And so we try to lay out those goalposts in a general and careful way. So that's the first thing—construct this normative yardstick.

The second thing, I used the yardstick is to roughly measure the quality of the public sphere in two different eras: the first era is the mid 20th century postwar mass media sphere, and then the second public sphere that we want to use the yardstick on is our current environment of social media-infused digital information.

And, in application—in using this yardstick—roughly what we find is that the mass media public sphere did some things *better* than the current sphere, such as seeking truth and for people to communicate civilly. But the digital public sphere does some things better than the mass media public sphere: it offers opportunities for a much wider range of views and for many people—especially those without resources and connections—to say things, to put their views out into the public sphere and for other folks to respond to them. And so, just to give a little more detail, I want to spend the rest of my time, the few minutes giving people a sense of what that yardstick is.

Our yardstick has two pieces: a set of rights and opportunities—and you might think of that as the structure or the affordances of the public sphere—and the second category, the second part of the yardstick, is norms and dispositions. That's how people themselves, roughly, behave in the public sphere. So the first category has five things and the second category has three.

Let's just talk about the rights and opportunities—and where do you get your rights and opportunities from? Some are from platform design. Some of it is from laws and regulation. Some of it is actually from how other people treat you in the public square or the public sphere.

And so we think that there are five important things: One is a right to free expression—which is kind of a first amendment value—you shouldn't be just shut down because of your point of view or what you say. The second—and this is an important one—is that everybody, each individual, should have good and equal chances for expression, for saying things in the public sphere. So, for those of you who were living and communicating before the internet was widely diffused—people who had a column in the newspaper had much better chances to get their views out there in the public sphere than me as a graduate student—I would write these angry op-eds, submit them, and they would all get rejected. So my chances for expression in that public sphere were very, very low. Even though I had a college degree and stuff.

The third criterion is access. So everybody participating in the public sphere should have access to reliable and instructive information. And so, in the internet world, I think there is access to a much wider array of reliable information than there was in the mass media public sphere, but our access might be reduced a little bit because there is a whole lot of noise too.

[20:17]

The fourth criterion—the fourth part of the yardstick on the rights and opportunities—is diversity. As someone entering the public sphere in a democracy, I should have the opportunity to hear a wide variety of views on matters of public importance. Here again, I think that the social media public sphere does much, much better than the mass media public sphere—partially because most of the speakers in the mass media public sphere were professional journalists who had gone to the same schools, been socialized according to the same norms and behaviors. And, in the United States and Europe, some of their main sources were *government sources*. So the sources were largely homogenous and the speakers were largely homogenous in their viewpoints, whereas now, as we all know, there is a wide, wide variety, diversity, of views on whether or not global climate change is happening or not. On whether or not we should wear masks or not. On any conceivable subject there is a wide diversity of views.

And finally [the fifth is] communicative power: when I enter the public sphere I should be able to do things, find other people with ideas similar to mine, and join with them, and maybe organize a March for Black Lives, or—if I think the vote has been stolen—I might want to organize with other people in a peaceful protest to complain about that. A digital public sphere, or any public sphere, ought to provide the opportunities for people to exercise communicative power.

The second part of the yardstick is really, really important. The first part is about laws and platform designs and algorithms, if you like, in today's world. But the second part is about *how*

we behave—a lot about how we behave as users of the digital public sphere or being in the public square. And many of us in the 1990s thought that the internet would be very, very good for democracy; we had kind of quasi-utopian or bordering-on-utopian views about how liberating it was going to be for everyone.

And, I think, many of us now think that was just a mistaken view. And a big part of that mistake is I think we underestimated how awful we, as human beings, are when interacting on a platform. Never underestimate how awful you are or I am, because, if you do, you'll come to this utopian view about how great a free internet is—which it hasn't been so far.

And so, if the public sphere is going to be *good*, we have to be *better*. And so these are three ways in which we think people interacting on the public sphere need to regulate their behavior and be better: The first is we should have a high regard for the truth. We should try to separate the truth from things that aren't true and not amplify falsity. And the 20th century public sphere was actually much better at this—because journalists were very well-trained and made it part of their professional mission to seek truth and expose falsity—and they got pretty good at it. They didn't get everything right. They got a lot of stuff wrong. But that was part of what they were trying to do, unlike many users of the digital public sphere today—it's maybe more important for a lot of us to feel good or to tweet things that we agree with than it is to be truth-regarding.

The second normative disposition that we should have—the normative regulator that we should each have when we're on Twitter or Facebook or whatever—is we should be oriented toward the common good. We should be saying things that try to, at least, move all of us forward—even people who really, really disagree with us. And I can see a little more about what that looks like in the Q&A.

And then, finally, the third thing is civility: we should treat people civilly. Not in the sense that we should be nice to them or polite to them—maybe we shouldn't—because maybe they're not good people or they're out to hurt other people. But what civility means is that we should treat

them with enough respect to listen to their views, and maybe we reject them, but at least try to understand what they're saying and then to explain ourselves to them. That's what the norm of civility means is: treat someone else as somebody, initially, who you take seriously enough to try to actually figure out what they're saying, and if you think it's just not true or you have a different perspective, that is fine. But listen. And then also you have an obligation to explain—you can't just show people the hand. Showing people the hand would be a violation of the civility norm.

That's the normative yardstick of rights and opportunities and norms and dispositions. We think that there's a lot of ways in which the current digital public sphere falls short of those norms and dispositions and opportunities and rights. And we think it's really important that all of us—all hands on deck—government, citizen, platform companies get to work reconstructing our digital public sphere in a way that conforms to that normative yardstick. Thanks a lot.

[25:42]

Lucy Bernholz:

Thanks Archon. I see there are a number of questions coming in already. So, no hesitancy from the audience—which is great—please keep dropping them in. Let me turn now to Seeta.

Your chapter takes a very different approach and looks at communities and their interactions with digital technologies focusing specifically on issues of exclusion and inclusion, silence and refusal. How does technology fit into your analysis and what big questions did you grapple with?

Seeta Peña Gangadharan:

Thanks Lucy—and to our hosts at Data & Society and to my co-panelists. Let me just begin by saying—for the visually impaired or those joining in by audio transmission only—I'm a very skinny Filipino-Indian-American woman with black cat-eye glasses and a long black ponytail, sitting in essentially a storage closet in north London.

[Laughs]

To answer your questions, Lucy, I actually want to talk about the big questions, and then sort of back into the question of technology and how I've conceptualized it in this chapter and provide some additional context to the chapter. So what initially motivated me for this chapter were questions arising from the work that I do with Our Data Bodies. Our Data Bodies is a research and organizing collective and, at OBD, our praxis extends from a recognition that technology intersects all sorts of struggles for racial, economic, and social justice. And so we work in between these domains, as well as traversing typically outside of those realms—digital rights, privacy, surveillance, or issues of fairness, accountability, and transparency in data driven systems.

And for five years, so between 2015 and 2020, Our Data Bodies did a participatory and collaborative research project—some of which is reflected in this chapter—interviewing folks in the most systematically-oppressed neighborhoods, or groups that are systematically oppressed in vulnerable neighborhoods in Charlotte, Detroit, and Los Angeles. And what struck me in 2018, I think, when we first began talking as a group—seems like a really long time ago now—and what continues to strike me today is that there are myriad ways in which individuals and groups of people are resisting, rejecting, and refusing the tyranny of technological systems over their lives.

And this, I argue, in my book chapter represents a form of digital exclusion. And here I'm thinking—and for those of you unable to see the screen—I have a slide up that says “refusal.”

And if you look closely at the image, it's a collage of different individuals expressing themselves, protesting—a lot of pictures collaged together of protests in the background in those letters that make up the word "refusal."

What I argue is there a type of refusal of technology that is a form of digital exclusion—it's sort of self-exclusion from digital systems or sociotechnical systems. And, importantly, it's a *generative* type of exclusion from these sociotechnical systems that can get people reimagining the role of technology in their lives. The ways in which technologies do—and *should*—govern their lives.

So, whether in Los Angeles, Detroit, or Charlotte, over and again in the interviews that we did on the ground, we heard multiple people thinking and talking and confronting technological systems, confronting data-driven systems, confronting the processes of data collection that they encounter so many times in a day, in a week—and doing this confrontation in a way that helped them survive, helped them feel dignified, helped them set the record straight, gave them a sense of control. And, in some cases, take direct aim and agitate against the institutions aiming to force sociotechnical systems upon their lives in profoundly unjust ways.

[30:59]

My chapter grapples precisely with that. It's sort of the meeting ground of this participatory research and organizing project that we did for five years and taking that experience—and the questions that were arising in that work—and putting that in conversation with how we think about technology in relation to this larger project of justice and this larger project of democracy. So the major questions are: How do we see or view these acts of technological refusal, whether they're individual acts of refusal or organized responses or collective acts of refusal against technology in political terms? How do we see this as a civic action or almost a civic duty or a new form of civil disobedience with consequential political reverberations?

There are two additional questions that also factored into or shaped the chapter which is: How do we center experiences of members of marginalized communities in this discussion of the relationship between technology, justice, and democracy? Because I think one difference—at least in terms of how I orient myself to the public sphere or understand the public sphere—is with recognition that there have been systemic forms of exclusion since the ideal or historical public sphere was conceived. That we've had systemic exclusions from the public sphere. And that often these groups—that *are* marginalized, that *are* exploited, that suffer from systemic violence, that are misrecognized on a daily basis—that they don't really factor into how we're thinking about information and our communication flows. Or even, beyond information and communication that is, you know, these two terms that really, anchor much of the thinking around democratic theory and mediated technologies of communication or media of communication.

The second important question was: How do we center the experiences of members of marginalized communities? And then the third question, I think, is centered around this idea of: How do we square when I'm talking about—this agentic form of technological refusal, this generative form of digital exclusion—with broader public debate and broader public policy initiatives that have typically characterized digital exclusion as a pejorative thing. So when we think about digital exclusion, we're typically thinking about: “I don't have access to the internet,” which is definitely an important and valid concern and one that has been absolutely highlighted during the pandemic. There are other forms of digital exclusion that aren't only about access to broadband or availability of broadband. There is digital exclusion that has to do with skills and know-how—a literacy-related type of digital exclusion. There are exclusions that have arisen with concerns about the ways in which digital technologies track us or target us—sort of this commercial targeting—and so privacy concerned are amplified and a form, potentially, of digital exclusion. And all of the surveillance concerns that can be marginalizing and can be impactful, especially on members of marginalized communities, creating a sort of digital exclusion. And all of these kinds of digital exclusion—these other four digital exclusions—are very negative. On top of that, we have a digital exclusion that I think is perpetrated by the sheer wealth of these big

technology companies and the economic geographies of these companies and that creating different kinds of exclusions in society. And so, if we think about those kinds of exclusions, they're generally negative ones—they're pejorative, right? We think of digital exclusion as a negative thing. But what I'm trying to do in this chapter is really recover this idea that Nancy Fraser, who wrote a lot about the counterpublic sphere earlier in her career, thinking about: "How do we ensconce ourselves into groups where we can see ourselves, and where we can see what's happening, articulate our concerns, and then express ourselves? Express our discontents and, ieally, eventually, turn that communicative power into political power." Being influenced by that, I have used this idea of technological refusal as, again, this force of political possibility. And, in that sense, I'm a firm believer that technology in the space of—when we think about marginalization and when we think about injustice and democracy, I understand technology in a very constructionist way: that technology isn't just a variable that *enters* the equation of democracy, it's something that is part of a larger set of institutions, practices, norms, social hierarchies. And that really affects and can exacerbate some of the existing problems that members of marginalized communities face—and yet they still can refuse.

So I'll end there and we can move to Rob.

Lucy Bernholz:

Thank you Seeta. The chapters and the key questions on dignity and power are very, very rich. And to show the breadth of the book, Robyn, your chapter looks at questions of power in a very different way: you look at different approaches to corporate governance and decision-making about content moderation, how those activities relate to democracies writ large. What do you want people to learn from your analysis?

Robyn Caplan:

Hi, everybody. I first wanted to say how amazing this is—this book was such a long process—this was three years in the making and it was a lot of traveling back and forth, and a lot of going back to Stanford specifically. So I'm so just happy that Data & Society could host this event for everybody. It just makes me very, very happy. As, also, the most junior person in this group by far, so it was a very overwhelming experience, at times, and it still is—to share the stage with these incredibly brilliant scholars, so I'm just very grateful to be here.

So I will be honest: I sat down to write this talk yesterday and it was really tough to go back to all this material that was so grounded within the context of the last four years. This book is very much spawned by questions of how digital technologies have shaped democracy, particularly as it related to the Trump era—I think that was really the impetus. I think our job now is to understand how many of the issues that gave rise to Trump—White Supremacy, political polarization, inequality—transcend the current digital era and also where we position platforms and digital technologies as digital actors. And I had the privilege of moderating a panel this week with Tressie McMillan Cottom and she's been influencing my thinking about my work actually all week. She noted that part of what we've seen in the last few years is, in many ways, an extension of how whiteness and power has functioned historically. We've also seen an expansion in the boundaries of acceptable political discourse in ways that are, frankly, undemocratic. And part of our job now is to redefine those boundaries—and platforms have a role in that.

On that note, I'm just going to give a bit of background into why I wrote this piece for the volume. The chapter is titled "The Artisan and the Decision-Factory: The Organizational Dynamics of Private Speech Governance," and it's actually based largely on a white paper I wrote for Data & Society, and released in 2018—which I actually never presented as part of Data & Society—so this is the first time, called "Content or Context Moderation." And this work is part of a broader research project that I'm still engaged in that looks at how platform companies are looking to network and distribute responsibility for content policymaking amongst a broad range of stakeholders including media organizations and fact-checkers, civil society organizations, academics, and users.

And this is a very funny chapter to have in this volume, I think, because it really gets into the

details of content moderation—which is an odd thing to pair with really big democratic theory But what I hope chapters adds is just some acknowledgment of how these theories and goals are interpreted and shaped within the context of the constraints and goals of organizations, in this case platform companies.

The chapter begins with an acknowledgment of something that it felt like the world learned in the last two weeks: platforms are private companies and, therefore, the First Amendment does not apply. I take time to consider in the chapter why this has been such a hard thing for people talking about platforms to accept because of how dominant a role they're playing hosting the public sphere—which is what Archon's chapter speaks to—but also because of the optimism and hope we had for these sites early on as being able to fundamentally transform democratic participation. And then moving on from an acceptance that these are public spaces—privately owned and governed—the chapter then goes into a more detailed examination of these companies and how they develop and enforce their content policies.

So I took a comparative approach in this work: I interviewed people who worked in policy across ten major platforms that host speech—so really all of them—I don't even know if I would be able to list them off the top of my head. Facebook, Reddit, Patreon, Vimeo—from the very, very big to very, very small. My real focus is trying to understand how these companies dealt with the types of speech that tend to adapt and shift very quickly online, so hate speech and false information My goal was to understand what parts of mahjong policies certain types of platform companies are good at and what types they're bad at, and how they kind of engage in this boundary-making. This work was also intended to understand how platforms of various sizes, of different business models, balance their desire to scale with the real need to learn the specific cultural contexts—and sometimes the languages of—where they're operating in order to make and enforce content rules. It was inspired, in large part, at the time, by a really fantastic piece of journalism done by ProPublica in 2017 where they discovered that Facebook's particular approach to content policy—which emphasized *scale*—protected white men but not Black children.

And what I found was a remarkable variation in how platforms form their policies. I pulled out

three different strategies that I found were spoken about most in my interviews. So the first was an artisanal case-by-case approach, and I don't use "artisanal" because I'm a Millennial currently sitting in Brooklyn, I use it because this is the word that they use and so I wanted to highlight it—and because it was always contrasted, in my interviews, with the notion of "industrial content moderation." And that's what's done by the major platforms, like Facebook and Google, where they have hundreds of thousands of people doing this work all around the world. And that is contrasted with the community reliance strategy that we often associate with platforms like Wiki media and Reddit.

And what I found was that as companies scale they move from this "artisanal" model—where content policy workers are kind of able to individually consider each flag and concern—to the industrial—which tends to emphasize a broad standardized approach with the goal of automation. And what I wanted to convey was that, as we think about regulation in this space, we need to carefully consider what type of content governance we want to incentivize: Are we going to build policies inspired by the approach taken by companies like Facebook—that has a bias towards developing global standards—or are we going to incentivize a more case-by-case approach that requires platforms to invest a pretty significant amount of resources into how they do it.

And I do think it's interesting to note that platforms can do many of these strategies at once. So the recent deplatforming of Trump was a very context-led moderation example that was also very informed by shifting power dynamics. It's important to see *when* platforms are using these different strategies and why.

To answer Lucy's question, I want to return back to what Tressie said earlier this week: "Platforms, like all of us, need to be engaged in redefining the boundaries of acceptable political discourse." For this work specifically, I think it's important to consider the strategies they're engaging in while they do this boundary work—how many resources they're devoting to it—and, in relation to that, how much of a role we expect them to play in shaping the public sphere compared to other institutions?

But I also think it's really important to consider the values that are driving their approach. And what I'm finding is that, increasingly, they're trying to mimic other institutions as a way to build back their legitimacy, as their novelty has worn off. So, in some other work I've been doing, I've been analyzing how platforms are in their public rhetoric often trying to position the work that they're doing in relation to how governments makes rules—so they're using terms like "precedence," building their own Supreme Courts, referring to their own little mini-legislatures. They're also, at times, trying to build off of the norms of media policy, emphasizing values like diversity and localism and representation. And they're also trying to reinsert expertise, touting how they're working with academics, civil society organizations, and media. So they're doing a lot of experimenting right now, in terms of what values are really going to drive their weight and the way that they want to operate going forward.

And what's really interesting about it—what's really funny about it—is that they're desperately trying to draw from the same institutions they were really intent on disrupting. So, in that sense, they can't just keep trying to prove themselves to us and gesture to inclusiveness and pluralism; they also really need to let people in in a real way and not just continue the harms of these other institutions.

Part of the job that we need to do—and I think that Archon really said this as well—is to continue theorizing, but maybe with a bit more skepticism this time about the role that they should be playing in the public sphere to guide this kind of "value development" on the part of platforms: So what do we expect from these companies—in the same way that we have with media in the past?

[46:28]

Lucy Bernholz:

Thanks, Robyn. So there's an example of the breadth of the book—the volume and the different chapters—and Robyn got us started bringing this all to the current moment. Again, it takes

several years to write a book. I want to ask each of our panelists to just comment on the following and then we'll go to the questions and answers that are flowing into the Q&A.

How does your chapter speak to what's happening in the U.S. and around the world today. Is there anything you might reconsider if we were starting this project right now?

Seeta, I will start with you. And then Robyn, you touched on this but you may have more to add. And then finally, Archon, before we go to the audience questions. So Seeta?

Seeta Peña Gangadharan:

Thanks for that. I mean, the events of the last four years have been tremendous and I think that, looking at where we are today in relation to where we were three years ago when we started engaging in this process, I was certainly more optimistic about this concept of technological refusal—this idea that *we can collectively respond*. I mean, there are individual responses that people are practicing and those practices are really important—and, at the same time, we're starting to see an emergence of practices of refusal emanating from different cities and municipalities around the country, so with respect to facial recognition or predictive policing or other kinds of technology.

What's interesting in the last three years is we've also started to see the rise of things like worker walkouts and high-wage-earning tech workers really trying to generate some movement around—whether it's protesting Google's contract with the Pentagon and Project Maven, or protesting Amazon's contracts with ICE—we've seen some pushback within companies.

Some of that is genuinely exciting, I think, and, at the same time, what I'm also seeing and sensing is that it almost feels like there is a co-option or an attempt to coopt this idea of technological refusal by the platforms—by the tech companies. "Yes, let's do participatory AI!" or "Let's do people-centered, human-centered technological design."

And those are important and I don't want to discredit them, and I acknowledge that within, you

know, the tech industry, there is a wide swath of people and some of them are doing really good work—it's not just extractive or co-optive. But, at the same time, I think that is drawing some energy and attention away from the very important work that's happening on the ground as communities are—members of, especially marginalised communities—are trying to grapple with: How do you intersect in the governance of these technical systems that are brought into public safety, that are brought into public education, that are brought into other realms of life—and it is really difficult to disentangle.

So even in Detroit—where most of the empirical focus of my chapter lies—I was optimistic there was going to be additional momentum and really exciting pushback against facial recognition within the city of Detroit and, potentially, in the state of Michigan. Now, there has been a lot of tension and a lot of pushback by the police and law enforcement and politicians, as well, and it's a tough struggle—so that is ongoing and it's continuing andI have a vision of how long that struggle is going to be. And it's going to be really tough.

[50:52]

Lucy Bernholz:

Thank you. And you bring up also the pervasiveness of the technology as well, as it moves off of our screens and into our streets, as it were, making both sites of pushback and resistance and incursion more pervasive.

Robyn, anything you wanted to add to your comments about how you might reconsider what you looked at if we were starting the project today?

Robyn Caplan:

I think that question is very funny for me, because I'm still very much engaged in this project. The chapter that I did for this book was really the beginning of a lot of this. I think a lot of the

work that I'm trying to do is really trying to understand how platforms are saying they're distributing power, distributing policymaking across a broad range of actors. I think, weirdly, what I would do a bit differently is—I tried to take a pretty good account of where I sat in that as an academic that does work on platforms—but I think I would have been way more conscious of that very early on, especially as somebody who sits in-between academia and civil society. So that's one small thing I would change. The real challenge for me is that I started this work in 2015—kind of *just before* this issue really blew up—and the big challenge has been trying to do work on an issue that's continuing to unfold. So I don't want to say, if I were to do anything differently, I probably wouldn't have done this project.

[Laughs]

Maybe. I don't know. It's a tough time to be a media scholar right now, especially one that does work on this area.

Lucy Bernholz:

Thanks, Robyn. Archon, same question to you: anything you would reconsider if we started today?

Archon Fung:

Yeah, that's great—I'll just answer briefly so we can get to the Q&A. One is platform censorship. I guess, at the beginning of the project, I would have been much more opposed to platforms deplatforming Donald Trump—I'm somewhat more sympathetic to that now. But on the flipside—I haven't actually talked to anyone that is active on Parler—but I imagine that the experience of Apple and Google and AWS yanking Parler, I have to imagine it must feel something like a Russian dissident [and] the state of Russia breaking up your printing press. Or, if you're a Hong Kong dissident, the government of China shutting down your social media account or yanking your website. So, I hadn't really considered the censorship dimension and

both sides of it. And I'm particularly worried about the—you know, Donald Trump has plenty of other ways, he has *good access* to the public sphere, no question about that, with or without Twitter—but I'm thinking about suburban Georgia regular person that thinks that the election did go another way, and is talking with eight friends about that, and all of a sudden their platform is shut down by the biggest companies in the world and they have to get hosting in Russia. Right, I mean, how do we think about that exactly, right?

And then the other thing that follows on that is the urgency. I was thinking: "Eell, okay, the platform thing, they're going to do what they're going to do—they're going to establish moderation councils, they're going to play around with it. But now they're making big, big consequential decisions all on their own that are dramatically different from two years ago or four years ago. So I think the urgency of getting some public grip around reconstituting platform decisions which constitute the public sphere is of much greater urgency, even, than how I was thinking about it when we wrote this.

[55:24]

Lucy Bernholz:

Thank you. So turning to the questions from the audience—and please keep populating the Q&A. I'll start at the top of the upvoted list and then I'll meld a few, because we have about fifteen minutes or so here.

The largest question that Danny Mendelsohn asked, and a lot of people supported, has to do with how each of you is even defining a public square—and you come at it from such different directions. "Can privately-owned companies be a public square? Or is there just a mismatch at that level that needs to be rethought?" I'll let any one of you that wants to jump in jump in first.

Archon Fung

Yeah, I'll jump in. I think this is a really important question. There's a bunch of ways to interpret it but, just by way of clarification, the public sphere—like how we learn about things, how we make decisions about whether to vote for this candidate or that candidate, how we decide to be an environmental activist or a racial justice activist and get that information—has been dominated by private sector actors for a very, very, very long time. Since the 18th or 19th century! The public sphere has been constituted by private companies *forever*. I mean, not for literally forever, and there are some exceptions, like BBC and PBS are publicly owned—but those are small parts of our informational environment. And, I think this is not where Danny is coming from. but one thing you've heard recently is: "Oh, you know, Twitter and Apple and Google and Facebook deplatforming Donald Trump or anyone else is absolutely no problem, because they're not governed by the First Amendment. States should take no action—they're not states, they're companies—so, of course, they can run their business any way they want."

From the perspective of our chapter—that is totally missing the point. What we want to do is develop norms and regulations that constitute these big platforms, and everybody else who is making the critical infrastructure on which we all depend, good for democracy—not bad for democracy. And for me, the property rights of Apple, Facebook, Google are like—twelfth—on my list of considerations of what is important, whereas different democracy considerations are 1–11. So that is really, really important to understand.

Lucy Bernholz:

Seeta, Robyn, anything on that? Seeta?

Seeta Peña Gangadharan:

I was going to say, Robyn, you should go first. I'm really curious to hear what you have to say, and then I can wrap it up.

Robyn Caplan:

Yeah. I think it's a very important point. And I think that it's always worth it. Because, right now, why people are bringing up the First Amendment—especially in the last couple of weeks with Trump—is because the conversation is focused so specifically on Section 230 reform and what should be done about that, and really what are the specific rights that these companies have versus don't have. And I completely agree with Archon. We don't want businesses to be denying people things on the basis of their gender or their sexual orientation and we don't really want platforms to be doing that as well. I think when it comes to Trump, specifically, and [in the case of] a lot of the people who got kicked off in the last week, it's a difference between allowing viewpoints to exist online and allowing calls to violence to exist online.

And one of the more interesting things that you could see was that there is an account that was called like @SuspendThePres or something like that—it was an account that was basically retweeting everything Trump would say and seeing how long it would take for him to get suspended. And he got suspended three times over the course of Trump's presidency for what he said, and Trump did not get suspended at all. And I think that's also a big complaint that people have online—that platforms have these rules but they're applied to different people differently. And, a large portion of the time, the powerful are the ones that the rules do not apply to. Which is really turning this thing on its head. In media policy, you think we should have higher standards for public figures and not lower; and, for platforms, much of the time we have—at least with platforms because of the way they're constructing their concept of "newsworthiness"—they've really made lower standards. So that's what I think people are really concerned about is the fact that you have these platforms that are saying—well, really, really they place a huge emphasis on procedural fairness, it's been a big thing and I think part of that is that some of them are working with Tom Tyler to develop some of these ideas—but they really want to have these fair rules, but they're not. They're not. And platforms are doing this really interesting thing where they're reinserting the gatekeeping function in a lot of different ways around very particular topics—like COVID and the election, they did that as well. But a lot of the

platforms they've also been working with media companies, because they have financial partnerships with specific users because they have different financial partnerships. It is not an even field, by any means. And whether or not you get your content taken down—and then, whether or not you're able to get an appeal through—is *highly dependent* on your relationship with that company. And that is dependent on power. And those are my thoughts.

Seeta Peña Gangadharan:

I kind of agree with that. In the interviews that we did for Our Data Bodies we found a couple of people mentioned, especially in Detroit: "I'm not really sure why these posts are being taken down," or in Charlotte after the uprisings: "I'm not really sure why certain things are not appearing." And the ability for those individuals to actually winnow their way into some unknown platform, bureaucracy, or labyrinth—like, that is just not going to happen.

I guess, just to come back to the question—I think I mentioned in my comments—the public sphere has never been or is not usually that welcoming of members of marginalized communities. So there are both advantages that you can point to, with respect to these large platforms, you could say that some of the uprisings over the past year were really aided through [platforms]—and this was the case with BLM and Black Lives Matter in the post-Ferguson era—there is something very genuine and important about being able to get your perspective out there.

What I worry—and think that we have historical examples to learn from—is that platform companies have a power that *isn't only about speech* but that gets *entangled with speech*. And so, we should be asking ourselves and really thinking: How did we get to here? How did we allow this to happen? And just thinking about how platform companies sold themselves on speech—that was the PR line for twenty years, right? I mean platform companies, Teleco providers before then, right? Something has really changed. It isn't only about speech, it's about speech power and a type of economic power that I think we haven't seen before.

Lucy Bernholz:

Thanks, Seeta. I also think I'll use the moderator's prerogative here: it's also very much—and we're seeing this more and more and the pandemic highlights, I think, in very powerful ways—that framing it entirely around speech is also not getting us towards solutions. And the technology, itself, the companies we need to be talking about have left the screen also. We're talking about companies that build tools very much designed to digitize our physical spaces and to affect the way we organize and assemble and can even get together as a two or three or a bigger group of people. So I agree with you very much there. I think we're not seeing all that we need to see by focusing solely on speech.

There are some very big questions in the Q&A so I'm going to try to meld them a little bit. But I want to start with one that starts with direct focus on you, Seeta, about this idea of refusal and how do we theorize access to refusal? Who can refuse and what must be the prerequisites for that? And are there ways in which it gets so entangled with the broader question of access which you mention? And I'll start with Seeta and then, if Robyn and Archon have thoughts on that as well.

Seeta Peña Gangadharan:

That is such a great question, and a really tough one because there have been moments in my journey as a scholar where I've thought we should be learning and we should be thinking and we should be researching communities that have very little access. And it seems like, for many years, we kind of left that behind—at least in the United States—we don't know as much as we need to know. And when we do research it, it's usually from this lens of: You need that access. And I'm not denying that at all, right?It's impossible to be homeschooling a kid right now without access to reliable broadband. There is no question about that.

But I do wonder, and certainly in the work that I've done with Our Data Bodies and organizing

communities in general, I think part of what we try to think about is where have we organized already? What are we already doing to meet the needs of our communities?

And that's not to say that I support only a communitarian understanding of what democracy ought to be. It's just that we can learn a lot from communities that haven't had access—and from communities that have deliberately shut out certain kinds of access. And, that, I think can inform our own consumption habits, where we put emphasis—it's almost a question of *when* do we need that access? And when do we need to refuse it? And do we actually have the right to refuse it? And what is enabling that right of our refusal. That is how I would answer that question.

Lucy Bernholz:

Archon, Robyn? Nope. Okay.

I have three more minutes, so let's see how quickly we can get into some of the questions. There is a very specific question and I'll ask it to all three of you which is: If there are examples of platform design, technology design, moderator tools—and I'll add in here governance processes—that are actually achieving any of the goals each of you is interested in?

In other words, are there good things we can look to? Are there good examples of where, maybe, a better relationship—a more democratic, equitable relationship—exists between communities and certain technology tools? The specific question from Matthew asks: “There is a design trade off between breadth and ease-of-use readability but given your expertise and the breadth of domains you’ve touched on, I think it goes beyond that.

But any last comments before I turn this back to Rob?

Robyn Caplan:

I guess that speaks mostly to my work. I have always really struggled with this question, because I think with all the models that I studied there's trade offs. And I think, actually, the reason why I struggled with it more is that—the design I like best is the community-reliant form—it's really great to have communities building their own structures and doing their own norm building. Of course, there is a very urgent labor problem with how platforms use the labor of volunteers.

But I also think because community-reliant moderation—it ends up showing some of the problems with kind of creating an ideal anyways. Because, in most of those cases, the way that those platforms are governed—who has the time to actually do the moderation—it ends up actually reproducing many of the same power dynamics that exist offline. And so I think it's a problem to posit any of these as an ideal. That any of these ways—"if we can only perfectly design the perfect system, then we would solve all the problems." It's just not possible. And many of the problems that existed offline are the problems that exist online. So I have a real hard time answering this question, sorry.

Lucy Bernholz:

Thank, Robyn. Archon?

[1:10:09]

Archon Fung:

Yeah, that's a great question, and I like Matthew's formulation. I think that it's premature—how we should be thinking about it is not what are the tools or specific solutions—but what are the processes to get to some of these values that we're aiming toward. And let me just take the platform part of it. I think that in Facebook, YouTube, Twitter, there are many teams trying to optimize for many things. Like engagement, like sales, like user stickiness—many, many different dimensions. And so, I guess I would like to see: What if in each of the platforms we had a dozen teams and each of those teams had an ethicist, a data scientist, and a software engineer?

And they were doing AB testing? And they picked one of the many democratic values on my list—like truth orientation—and I bet you could get truth orientation in your YouTube algorithm or in your Twitter algorithm cranked up 20, 30, 40 percent in about five weeks if you just put your mind to it with those teams of people.

So the moderation, I guess, for me, is trimming at the edges of what you regard as false or fake news or hateful or harmful. But the whole iceberg is about the content itself: Are we being civil to one another? Are we orienting towards truth? Are we getting towards some kind of common good discussion?

And so what can we do as citizens—what can regulators do—but here I'm focused, particularly, on what platforms could do to up their democratic quotient *if they really tried*.I n the 90 percent of the iceberg that's under the water—not on the noisy, hateful, false Russian bot, troll farm stuff—but in the algorithm itself.

We know that 70 percent of views on YouTube are not organic. They're stuff that the algorithm feeds you, not on what you choose yourself—and any of our feeds are that way. So how can we increase the democratic quality of the feeds? It's just not a problem that people are trying to solve—but why not?

Lucy Bernholz:

Seeta, last word on that question.

Seeta Peña Gangadharan:

Yes. It's such a complex question. I don't know if I have a great answer except to say: all of these really hard questions that we're asking of now ought to be asked for technologies coming down the line. So I'm just thinking of quantum internet—let's have that conversation now. Let's think about the investment strategies, let's think about how this innovation is being not only

technically designed but marketed. How are we understanding the benefits of this innovation to society? We ought to be very meticulously capturing all of the questions that we've had and all of the learnings that we've had over the last 20–25 years, with respect to digital technology and this very different trajectory that we are now on from where I think we first started out with all of our optimising understandings and discussion about the information superhighway. We really need to be thinking about these questions and applying them to the things that are coming down the pike. So while there is a lot of urgency around thinking about platforms and platform governance—there is a lot more that is on its way. And a lot more that is already being put into place partly because of where we are with respect to the pandemic and how reliant we've become on technology.

Lucy Bernholz:

Thank you, Seeta. I want to echo that. I think one of the things that came up quite a bit, both through the process of writing the book and also Rob and our third coeditor Hélène Landemore had the opportunity to rewrite the introduction once the book was finished. So the introduction is probably the most recently-written part of the book—and I hope we made that point.

There was scholarship, largely coming from Black women and women of color scholars—twenty years ago that told us about the content moderation problems we'd be having and facing today. We shouldn't repeat that same mistake. We should be doing that work, absolutely, now on the pervasive nature of computing and its interactions in our social life.

I want to thank all of those of you that put questions in the question box and, in particular, just call out for the panelists to take a look at. I don't know what we're going do with this Q&A, we're not going to get to it all. But there is a really strong question, I think a wonderfully phrased question which is often sort of dismissed as naivety—the good newness and positivity of the internet—if we centered White Supremacy as a defining characteristic, at least in the U.S. where so many of these technologies are built, what kind of logic might that have led us build off of and reimagine what we built or where we're going, or at least how we understand it."

So I want to thank Seeta, Archon, and Robyn, and thank Data & Society. I'm going to pass it back to Rob Reich for the last few comments. Thank you so much.

[1:15:51]

Rob Reich:

Thanks for this terrific conversation—both the audience contributions and then all of the contributors here—thanks to Data & Society as well as Stanford PACS for cohosting the event.

I'll make one final comment by way of conclusion which is that my takeaway lesson from the 2+ years of thinking about this project and working on the book and now hearing this conversation, is what we come to at this moment is the urgent need for a multidisciplinary research agenda that brings together a variety of different scholars in the spirit of this book—but this is just a starting point. Because there are both normative questions, social scientific questions, and technical questions—and the intersection at the bare minimum of those three orientations—is what's necessary, in my view, to try to deliver ourselves a future where digital technologies will help democracies and citizens flourish rather than suffer.

So takeaway lesson No. 1 is the urgent need for a research agenda of the sort that this volume represents. And No. 2, I'll just mention that I think, we finally, indeed, are emerging—here in the United States, with the Biden administration—into a moment that a sober and mature stocktaking of where we stand is possible. So for those of us in the academy, the work that we can do involves not only the research and the research agenda I just mentioned, but also a curricular change so the technologists that get trained—perhaps especially in an institution like the one that I teach at, Stanford, which has provided so many of the engineers that populate the companies that have created and distributed the technologies we've been discussing—aren't just oriented towards questions of optimizing their technological designs alone. And, as Archon and others mentioned at the start, the idea of combining ethical, social-scientific, media and communications, STS—and plus engineering expertise—within companies as well as within

public agencies is also an urgent task for all of us.

We have so much work to do in front of us and I mean that as a big collective we. And I want to thank everybody for being a part of the conversation. I hope the book is of interest to anyone who takes the chance to look at it. And I will just signal, certainly from my point of view, my contact information is super easy to find online. I welcome any comments, thoughts, proposals. Contact me, I would love to be in dialogue with you. So thank you all.

[1:18:36]