#250: What's wrong with tech companies banning people? (Julian Sanchez)

Julia Galef: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef, and my guest today is Julian Sanchez.

Julian is a senior fellow at the Cato institute in Washington D.C., where he writes about technology, privacy, and civil liberties. He's also written for a wide range of publications including Reason Magazine, The Atlantic, The Nation, National Review, and I've been reading Julian's blog ever since I was in college.

I reached out to Julian to help me think about the issue of tech companies deciding to ban users for speech that is hateful, or false, or inciting violence. Most prominently, of course, there's Trump getting banned from Twitter, Facebook, Youtube and other platforms, but other examples in this category include the app Parler getting banned by Amazon Web Services, and earlier this year, companies including Youtube deciding to ban people for spreading misinformation about the coronavirus.

And many of the particular decisions made by these companies, I don't mind those people getting banned, but nevertheless I've been uneasy about this situation we seem to be in where a few tech companies seem to have a huge amount of power to determine what speech is allowed and what isn't.

So I turned to Julian – who I know to be an especially principled and nuanced thinker on this topic -- to help me figure out how concerned I should be. Here is my conversation with Julian Sanchez.

[musical interlude]

Julia Galef: Julian, I've been reading some of your Twitter conversations in the last week since Twitter and Facebook suspended Trump, and I've seen you make the point repeatedly that that is not a violation of "free speech." That these are private companies deciding who to host on their platforms; this is not a violation of the First Amendment.

All of which is true. I have no disagreement with that.

But in my mind, the question of whether the de-platforming of Trump and other right wingers is *legal* -- that's kind of the easy question. The harder question is whether that's good for society in the long run.

And I hadn't seen you comment on that, so I was just curious what your view is.

Julian Sanchez: Sure. I think you have to decide these things ad hoc. I would distinguish between the First Amendment and free speech. There are plenty of things that are consistent with the First Amendment and bad for free speech, but in this case, I have trouble getting, frankly, too upset.

I think in a sense, the premise -- the backdrop of the First Amendment is that the legal barriers to speech are very low. The ability to criminally penalize speech is very low. A whole lot of incredibly vile and sometimes even quite harmful speech falls within the ambit of the First Amendment, but then there are a series of other mechanisms that we use to push some kind of speech to the margins.

That's basically okay. That's what, in a sense, makes discourse sustainable and productive, that not every utterance and bleat is on equal footing.

Julia Galef: I like the word "bleat."

Julian Sanchez: I think in this case you have unambiguously bad consequences arising from a series of not just false, but completely meritless, claims about electoral fraud that I think at this point we can kind of comfortably say don't really need to be litigated for us to say that they're not true, and that are having obviously harmful effects. In that, people who believe that a basically free and fair election is rigged and engineered to deny the majority their choice has had the effect you would expect it would have.

When people believe that, that ordinary peaceful democratic politics is essentially a sham, is meaningless, and that they have no recourse, really, to expressing their political preferences through that mechanism... Many of them, at least, form the conclusion that only by direct action or force can they make their voices heard. We saw the result of that. Not just generally inspired, but more concretely and specifically directed, by the President.

I think what a lot of the platforms have said is that they make exceptions for political leaders from their normal rules that Trump and perhaps quite a few others under the normal rules that apply to every other user would have been booted off these platforms long ago. And so, in a a sense, the de-platforming we're seeing now is just the withdrawal of this special dispensation they were given, relative to other users. A special exemption from the normal standards.

And so, if they make that determination, I'm not worried. In significant part because if the question is "Well, what in practice is the effect on the ability of the President of the United States to get a message out, if he believes it is important to do so?" the answer is, in practice, really nil. It is not as though, absent the platform of Twitter or Facebook, the President is condemned to languish in obscurity, unable to make his perspective known to people who are interested in it.

I think about speech in some ways the way I think about heroin. I think it ought to be legal, but I don't necessarily think you ought to be able to find it in every corner deli. I'm pretty comfortable with a society where the law lets you get your fix if you're particularly dedicated to it, but this does not require us to make it easy to stumble across while you're shopping for candy bars.

Julia Galef: That's a fun analogy.

I know you said that these cases kind of have to be decided on an ad hoc basis, but I still want to make an effort to find some kinds of principles that we can agree or disagree about.

So, I guess Trump is a bit of a special case because he, as a political leader, can get his message out through a bunch of other methods. But other people who don't already have that fame or notoriety would have a harder time getting their message out. For someone for whom big social media companies are basically the only way to reach a large audience, would you be at all concerned about a precedent where the companies ban them at their own discretion?

Julian Sanchez: I mean, it depends on the reason. Within the assumption the platforms have, as a legal matter, the right to their own decision, the question is just: Are these good reasons? If a small person with few other options really just wants to just use racial slurs and call for race war and are told "Well, this is not the place for that" – yeah, I don't care.

Particularly, I think, in a sense, the government has this obligation to restrain itself from penalizing speech, to a significant extent without regard to the merits of the speech. Absent some very specific exceptions: inciting violence, and fraud. You can say something that is true and interesting, or you can say something that is false and loathsome, but unless it falls into one of these well-defined exceptions, it's not the government's place to make that judgment.

And I'm fine with the private platforms actually making, to some extent, judgements of merit and saying, in contrast to the government, "We will look to the content and say, if this is not just imminent incitement, but advocacy of violence, we don't think that has enough value for us to be required to rebroadcast it. If this is speech declaring the superiority of inferiority of particular racial groups, that is not of significant value to require us to amplify it." That's fine with me. I want for them to make their decision based on their judgments of whether this is a useful contribution to the discourse.

But also, in terms of making decisions about whether it matches the kind space they wish to create. When platforms say: "We've decided this is not going to be a space for pornographic imagery. It's not that pornographic imagery is per se bad and has no value. It might be very well to, say, go to a site for that, if that's what you want to see. And there's nothing wrong, perhaps, with wanting to see that. It's just that in this particular space, we want to create an expectation that that's not what you'll encounter here. As you're scrolling through your Twitter feed, you will not encounter a lot of photos of genitals. Nothing wrong with photos of genitals, but it's not what you're going to find here,"

... I think that's also fine.

Julia Galef: I notice that my intuition about that view changes a lot depending on whether I imagine there being ten different platforms, each of which has a 10% market

share, so to speak, of the audience... or whether I imagine there being basically just one platform that has a 99% market share.

Does your view hinge on the market share of the biggest companies?

Julian Sanchez: Yeah, to some extent, right? If we imagine a kind of very different world where the internet looks more like a giant AOL or Prodigy -- or the old French information service, Minitel, I think it was called – where there is really essentially one decision-maker, then I have a different attitude there. Because everything I'm saying is dependent, to some extent, on this model where there is a permissive baseline --

Julia Galef: How so?

Julian Sanchez: The permissive baseline set by the First Amendment. In the larger domain, the standard's extremely permissive.

And then you're going to have a series of more specific domains -- some maybe more mainstream, and some more marginal -- that will have a variety of different standards. Although there are probably certain things that there's going to be agreement on between a lot of different platforms, or across a lot of different platforms, that "We'd prefer not to have this here."

When you don't have that environment -- of, in a sense, the ability to see what kind of speech rules are congenial to users who have a range of options for expressing and consuming ideas to select between -- then that picture becomes less attractive. But I also just don't think we're remotely there.

Julia Galef: What would be the rough threshold for you, where you would start to rescind your argument?

Julian Sanchez: The problem with this is that it's difficult to compare market shares, in a sense.

This is why I think this conversation can be somewhat confusing. You have people saying things like, "Well, but Twitter is such a monopoly," or "Facebook is a monopoly." Which, to my mind, is just clearly ridiculous, if you're talking about the ability to express political claims.

It may sound more plausible if you're defining the market in which they have a supposed monopoly fairly narrowly. If you say, "In the universe of sites that do almost exactly what Facebook does," then yes, for at least English-speaking audiences, they are 90 something percent of the market. If you're talking about services that do almost exactly what Twitter does, then sure, in the English-speaking world, they have a huge segment of the market, even though the percentage of Americans actually on Twitter is not that high.

Julia Galef: Yeah, I mean, that definitely feels like an unfairly restrictive way to define the question. How would you define it?

Julian Sanchez:	Right. If instead the question is, "How many competitors in the domain of ways to get your view out, or your speech out, are there?" then none of these services look particularly dominant.
	Again, the president is maybe a too-easy case, but the fact that he's not on Twitter and Facebook is not any particular obstacle to getting his views out on any number of major cable television networks, and through the White House website, and through any number of other channels.
	Including having his words echoed by people who remain on those platforms. One thing to just bear in mind is the barring of a particular individual from a platform is not at all the same thing as the erasure of their ideas from that platform. There are a whole lot of people on Twitter and Facebook articulating things that Donald Trump is expressing, that he's no longer allowed to shout into the void directly.
	The question is: what is the functional ability, by one channel or another, of someone to get their ideas out into the ether, or into the discussion? And one standard here is, "Does that person have every venue that would have been available to anyone within most of our lifetimes?"
Julia Galef:	But isn't it different now because people have switched over, to some extent? Like maybe they used to read more newspapers or books or something, but now that Twitter and Facebook have come along, they've substituted for the newspapers and books, et cetera? So your options actually have shrunk compared to the past.
Julian Sanchez:	To some extent, but all of those legacy venues still feed into the new ones. So even if you're directly barred from Twitter or Facebook, the fact that you may still be able to get an op-ed published in the New York Times gives you a back channel into that discussion.
	So there are different questions. One is, do you as a particular individual have access to one of these social media platforms? But the other question is, via the other channels that remain available, are the ideas you're trying to express still part of the discourse on those platforms?
	And I think the answer is clearly yes. There's all sorts of ideas that are current in conversations on the major platforms that are not introduced via those platforms, and maybe best not being introduced via those platforms. There are popular podcasts that people listen to, that articulate things in a way that would be, frankly, quite tedious if you tried to make this into a Twitter thread.
Julia Galef:	Just to check the principle: Would your intuitions change if you imagined that the executives at all the top social media companies were conservative, and they had a rule that you couldn't be on their platform if you questioned the existence of God, or said anything negative about the military, or whatever?

Would you still feel like, "Yes, of course, that's fine. Not only is it their legal right, but I'm fine with it on principle, for them to be making those judgment calls about what speech gets amplified"?

Julian Sanchez: Right. I would, of course, disagree with that particular decision, in that sense that I would say "Well, I think you should probably have a different rule." But it wouldn't significantly change my view in any way.

But also, look, in part because I think there's an audience-disciplining effect here. Which is just to say, the endurance of these platforms is not foreordained. And a rule like "This platform is only for people who do not deny the divinity of Christ" would not be unrelated to whether that platform remains the dominant one for open-ended discussion.

There's a sense, I think, that some of these platforms are so enormous, and so locked-in by network effects because it's where everyone is, that they can effectively impose arbitrary rules and that will have no effect on their dominant position.

Julia Galef: And you don't think that's true?

Julian Sanchez: I just don't think that's true.

Look, ex-ante, if you said to me, "Think about the network effects argument -- do you believe Friendster can ever be dethroned?" I might have said, "Oh yeah, they're really locked in. I don't know how you'll ever get Friendster off the throne." Or MySpace.

Julia Galef: Well, how big were they, though? I feel like they were only ever about five percent of...

Julian Sanchez: But the snowball effect still should have made them a lot harder to displace than they ended up being. The truth is, they just ended up not being that hard to knock out of that dominant position.

Julia Galef: Maybe. I feel like there's some threshold of "percent of the country using a service," that once you cross that threshold, the network effects would be much harder to displace. As opposed to just the situation where the largest platform has 5% of the country, and even though they're the largest, it would still be much easier to displace because they only have 5% of the total country.

Julian Sanchez: Yeah, I'm not sure that's true.

One, I don't know whether the "installed user base" is necessarily the right metric. Some huge percentage of Americans have a Facebook account. I still have a Facebook account. I don't use or look at Facebook essentially ever. Maybe I check in once a month to see if anyone I went to high school with has had a baby.

It's got this sort of enormous base in terms of... it is sort of convenient, because it is the place where everyone is, so it is the place that's useful to check in and see whether, again, your old high school friends have had kids, because there's where you can be pretty sure all your old high school friends are.

But in terms of consumption of news information, or discussion, I'm just not using it for that at all. There are other places I go for that.

Julia Galef: Yeah. And it's true the younger generations have shifted over to Snapchat and TikTok and Instagram, and probably others that I haven't even heard of because I'm an old fogey.

Julian Sanchez: It's also just not mutually exclusive with using other things. When we talk about market share, in most contexts, we're talking about stuff that costs money, usually, and therefore to some extent...and this is true of social media too. It's to some extent displacing other alternatives, because time is not infinite.

But for the most part, when you say, well, "What supermarket do you shop at?" well, there would be one or two supermarkets. It wouldn't be that, "Well, I shop at one, but I'm also using 20 others." Whereas the fact that I still have a Facebook account doesn't really tell you anything about what number of other information sources I'm using.

In a sense, that's the important thing when we talk about the power of network effects. It may be that network effects are a reason why it's hard to displace Facebook for certain functions where it really is useful. That that's where everyone is. If I want to log in and make sure everyone I know is okay after some kind of national calamity.

But that's not really much of an advantage in terms of deterring people from joining other services that disseminate information. So, it's not that because Facebook and Twitter exist it was impossible for Discord or Twitch, or any number of other services that have grown enormous audiences in the last few years, to become popular by offering a slightly different way to do things. Or for online web discussion forms to continue to exist.

Julia Galef: Let me run another argument by you and see what you think.

It sounds like one of the cruxes here -- cruxes of disagreement between you and the people who are more concerned about the slippery slope or the dangerous precedent of banning right wingers from Twitter, one crux of disagreement is how foreordained one thinks it is which companies are the main platforms of discussion and discourse.

And you think it's really not that foreordained and not infeasible for other companies to spring up and capture a bunch of attention. And not all that infeasible for people who get banned from one platform to be able to get their

message out on another platform. And the people who are more worried maybe don't agree, or don't see things that way.

One argument that seems like it might undermine that is if you don't really think these decisions being made by the different platforms are being made independently at all. And I know no one's saying that they're *completely* independent, but if you think they're really all pretty dependent on the same pressure, then that would kind of undercut the value of having multiple different platforms to choose from.

So if you have a model where all of the platforms are responding basically to the same, let's say, mainstream liberal elite; the same people who work in the media or the same people who, I don't know, are influential in the liberal mainstream. Those are the people who are putting pressure on the companies. Maybe in part because those are the kinds of people who work at those companies. And so whatever that crowd of people has decided is acceptable speech is going to become the norm at most or all of the companies.

So I guess my question is two part: First, if that model were accurate, then would your views change? Second, how accurate do you think that model is?

Julian Sanchez: I would have more concerns, but I guess it depends what dimensions of uniformity we're talking about. If a lot of companies make similar decisions in the aftermath of the riot at the capitol because all of those people have a view that violence against democratic processes is a bad thing… I am, at some level, not that interested in that level of uniformity as sort of a harmful thing.

But yeah, if there's so much uniformity that there is a kind of good argument that there are other reasonable perspectives that are totally omitted from that decision-making process…. That reduces the value of having multiple loci.

Although again, I do think -- and I hate to sound like a central casting libertarian, but I do think this is one of those places where, in a sense, that the market will solve. In the sense that if the platforms are wildly unrepresentative of the norms of their audience, and the sense of what the audience finds acceptable, that is going to create pressure for alternatives.

This is one reason I have more reservations and more qualms about, let's say, moderation decisions "lower in the stack," that is, lower in the OSI stack, as it were.

Julia Galef: Oh yeah. Like what?

Julian Sanchez: In the sense that your internet connection is lower in the stack than the application you're running on your desktop. The operating system is lower in the stack.

This is a model that computer people use called the Open Systems Interconnection stack, in particular with respect to internet protocols, where

there's this set of different... there's the protocol for the ethernet cable that connects your computer to the wall, there's the TCP/IP protocol, and then there's a series of higher-level protocols that manage things like sending tweets or receiving webpages over HTTP. The lowest level of the stack, of course, is just the physical substrate, the wires or the fiber optic cables or the radio signals.

I think, in a sense, there is less to worry about the higher in the stack a moderation decision occurs. If you had ICPs deciding, "Well, we've decided these platforms are not going to be accessible," I would have significantly more concerns about that.

Julia Galef: That's a good distinction to make.

Julian Sanchez: In significant part because that tends to go with the lack of alternatives. Most people have -- even if their consciousness is dominated by Facebook and Twitter, if you get on the internet, there are thousands upon thousands of ways you can transmit or consume information. But most people have one or two realistic choices for their home internet. Although it's not just the one; there's usually a range of mobile providers at the very least.

But it's more concerning there, in part because there are fewer options, and in part because transitioning between them is more difficult. Changing mobile broadband providers is, as a rule, a lot more cumbersome than opening a new account on a new social site, which usually doesn't even require you to abandon your old social media sites. You can still have accounts on all the dominant ones and still open an account on a new platform, or install a new piece of software.

So if we think about Parler, I'm sort of still in the camp, I suppose, of "Well, if they're your servers, it's up to you what you host on them." I don't think anyone ought to be forced to republish or amplify content they find loathsome, or offer apps they think are harmful to the world. I don't think anyone should say, well, "Apple, you must make the Mein Kampf Race War app available."

Julia Galef: Please tell you just made that up as a hypothetical...

Julian Sanchez: I did, but I'm almost certain it's got to exist! An extension of rule 34 is that if you can imagine it, it exists as both pornography and as offensive speech, or just loathsome and hateful speech. These things exist in basically every variety. I'm sure somewhere there's someone dedicated to hating groups so obscure you're never heard of them.

Julia Galef: Hipster hate speech.

Julian Sanchez: So I kind of remain generally of the view that it's ultimately unwise, and to some extent, just sort of wrong, to compel companies to carry things they find really loathsome.

And that I think gets missed to some extent. We are, maybe very rightly, talking about the structural consequences of these decisions, but I do still think there is, at some level, a kind of moral dimension to this. Which is that people do not leave their values necessarily at home, in a box under the bed, when they go into the workplace or when they enter the world of entrepreneurship.

And so as a matter of respect for autonomy, if you have people who are operating or started a company and say, "We don't want this thing we've built to be used as a home for Nazis, because that's contrary to our values," then... there's just an argument totally separate from "Well, what are the structural effects of this for speech," that those people should not be compelled to override their own values.

It's interesting to see the disconnect between some of the people objecting to some of these de-platformings, and the attitude they had toward forcing Christian bakers to make cakes celebrating gay marriages.

Julia Galef: I noticed that comparison as well. Although, in that case, there are many more bakers than there are social media platforms, so I was a little more sympathetic to people who had a different view on those two things.

Julian Sanchez: Well, there are many more bakers, but I would bet that in most towns in America, you have fewer bakers capable of making large wedding cakes to choose from within a reasonably short drive, than you do options of social media outlets to get your speech out. In terms of how easy is it for you to choose among a number of alternatives, actually I think probably you have more social media outlets than you do bakers within at least easy accessibility, if you don't live in New York or LA or a very large city.

Again, the argument there was not just about these structural questions, but about the idea that people should not be compelled to do things that they feel betray their deepest moral commitments. Even if, like me, you find their deepest moral commitments pretty stupid and morally unattractive. So I think that's a dimension I don't want to totally omit here.

But with respect to Parler... so, one: looking at, for example, Amazon Web Services' reasons for saying "We don't want to host them anymore"... First, this is a decision, as we were saying, lower in the stack. Where Amazon Web Services hosts a whole lot of different platforms on their service. And so therefore, deserves more scrutiny and more qualms. In part because that's a decision that affects every user, and the speech of every user, on a platform with a very large and growing user base. Both those who are engaged in objectionable speech and those who are not.

So I think -- if not legally, then certainly in terms of our skepticism, or the scrutiny with which we greet that decision -- that justifies, I think, more criticism than we might offer if the decision had been made about an individual user.

But I think Parler does also sort of suggest, again, that the attitudes of the decision makers of the company being, past a certain point, out of sync with the norms of the audience and the user base, create an opportunity for new sites to emerge. So a site like Parler grew with, I think, pretty astonishing speed when a large enough chunk of the "legacy" or the incumbent social media base decided they didn't approve of the kinds of moderation decisions the leadership of those companies was making.

Rightly or wrongly. It's worth bearing in mind that the most shared content on Facebook -- consistently, the most lucrative and successful and widely shared content for years -- has been right wing content. It's Ben Shapiro and Dan Bongino and Donald Trump and Franklin Graham, and people like that.

And so it's a little bit mystifying that so many people appear to be convinced that these platforms are on a kind of relentless jihad against conservative speech, and determined to suppress it. They've done, I think, an astonishingly bad job of it, if that's the case, because they've been providing essentially a free mechanism for a lot of these speakers to be hugely more influential than they ever would have been before. So it is worth questioning the extent to which the picture that motivates the exit to another platform is accurate.

But this is the classic A. O. Hirschman model of "Exit, Voice, and Loyalty." You have a series of competing institutions, and the option of exit conditions the discourse within a particular institution. Your children attend a particular school, you have voice within the system in the PTA to meet with administrators and ask them to steer things in a particular direction, and the strength of your voice is to some extent amplified by the option of exit.

But also there is a point where, in a sense, the community maybe bifurcates, and so you actually create two somewhat different communities, exercising voice internally. So on the one hand, this creates a pressure to sort of expand the amount of voice afforded to the user base, in the interests of forestalling exit. But it also creates the option for separate communities that are more aligned in terms of their voice, and they're able to have somewhat different, but maybe more internally productive, conversations.

So on the one hand, I think there is reason to be a little more anxious about the de-platforming of Parler than the de-platforming of Trump. Because it is that kind of healthy mechanism, which is: Your users, or at least a substantial minority of your users, are out of sympathy with the norms guiding moderation decisions at the top of the incumbent platforms. And they're able to express that dissatisfaction by finding an alternative that's more in line with their norms.

And, one, I don't think they are doomed. We've seen a similar pattern with a platform like Gab. Part of the problem is, of course, when you sell yourself on the idea that you're going to be more permissive than the incumbent platforms -- you're going to allow what they do not -- it turns out that you then attract disproportionately people who want to say things that are unacceptable on those platforms.

And a lot of the stuff that's unacceptable on those platforms, it's not -- well, there's some conservative speech, but it is stuff that almost everyone finds repellant. You end up with, "Well, I just wanted a place where conservatives had a little more latitude, and I've found a kind of nest of literal Nazis."

So that creates its own difficulty in terms of sustaining a kind of viable external platform. You end up sort of needing them to be more permissive along certain dimensions, but still willing to do enough moderation to make the place palatable for people who are out of sync with, let's say, the leadership of Google and Twitter, but not anxious to spend their days on the Daily Stormer.

But those platforms continue to exist. Parler, even in its objections to what they say is breach of contract by Amazon, does not claim that without Amazon Web Services they're unable to ever exist. They say, "Well look, it's very inconvenient for us not to have 30 days notice to transition to a different solution." Because you just can't do that overnight.

But to some extent, this is a decision you make when you do external hosting rather than self-hosting. Sites that are, like Porn Hub, platforms and sites that host controversial content, or content that other companies tend to want to distance themselves from, will usually go with the option of self-hosting.

And you make a trade-off when you decide well, we're going to rely on the hosting services of Amazon for web services, or we're going to let Microsoft or Google handle our company's email. There's this trade-off. You get the benefits of the economy of scale and the expertise and the ease of internal scalability, of being able to grow your site -- by, one, taking advantage of that body of external technical expertise without having to build it in house. But also being able to scale from very small to very large pretty quickly, without having to literally buy more physical hardware yourself every time you've grown past the point you had previously allocated for.

That's a trade-off you make, but it does mean that you are, to some extent, in their power. You're using their stuff, and they may have the option to withdraw that permission. And there are pretty clear caveats in Amazon Web Services' terms of service, as well as the app stores, that define content they -- pretty broadly and sometimes vaguely, but -- content that they say, "You're agreeing, when we form this contract for us to host you, that you're not allowing your platform to become a vehicle for harmful content."

You can say, well what does that mean? If you look at the list in the litigation that Amazon provided, I don't think a whole lot of people would dispute that this is stuff it's pretty reasonable for a company to say they're not interested in hosting anymore. Just page after page of people talking about how hyped up they are for race war, and how the only solution now is to kill the N words and the Jews and various other groups that people are talking about, that C word and Stacey Abrams, that C word Nancy Pelosi, I can't wait to put a bullet in their head, they'll be good target practice for our beginners...

Julia Galef: Okay, I believe you, I believe you, stop, stop!

Julian Sanchez:	If I literally quoted to you, if I just spent five minutes... well, it would take longer than that. If I took 10 or 15 minutes reading out verbatim, and without my little edits, the content that Amazon cited in justifying their decision to terminate Parler's contract... and not just that it existed, but they said, "Clearly Parler has no viable plan for restricting this kind of content"... If I just read that list to you, there is no way you would put it on this podcast. There's no way, frankly, I think you would continue listening to me after about a minute. I think you would not want to hear it, and you certainly wouldn't rebroadcast it to anyone.
Julia Galef:	Yeah. Setting aside the question of how low we are in the stack, which I think is definitely an important distinction -- setting that aside, I'm much less concerned about de-platforming due to speech that incites violence, or threatens violence, than I am concerned about speech that is banned for being false.

In that latter category, I'm thinking of things like YouTube declaring, sometime last year, that they were going to ban misinformation about Coronavirus vaccines. And they defined misinformation as anything that disagreed with the WHO.

Which was very worrying to me, because the WHO has said... I don't know, they claimed masks weren't effective for the general public, and a bunch of other things that I would disagree with. It made me nervous that that was being held up as the standard for what was true or false. |
| Julian Sanchez: | Yeah, no. I agree. |
| Julia Galef: | That's the kind of thing where I worry a lot more. |
| Julian Sanchez: | I think that's absolutely right. I think a lot more self-restraint is sort of in order from the platforms when it comes to making determinations about truth. For a lot of reasons. In significant part, one, because we often just don't know with absolute certainty what the truth about any matter is. We may be very confident, but things that we have been very confident about in the past have changed.

But also, to some extent, the reason we can be reasonably confident about a lot of things has to do with the ability of counterarguments to be raised. Or of denials to be floated, or evidence against that position to be offered. So if you sort of remove the ability to offer the counterexamples, then to some extent, your confidence in your starting conclusion ought to be reduced, commensurately. It's no longer a proposition you're submitting to rigorous testing in the same way.

Although, different claims vary to differing extents in their dependence on mass media speech as a mechanism for testing it. With Coronavirus, I think it's a little iffier there. Because we have less certainty, certainly. It's, for example, a very bad idea to say, "Well, we're not going to allow people to become aware of a minority position held by maybe only 10% of physicians, |

but who have a different view about the efficacy of lockdowns," or something like that.

On the other hand, this is an instance where the wrong information can genuinely be immediately and directly physically harmful to people. And so, I'm a little more sympathetic about... if, for example, someone was touting the "bleach cure" for coronavirus on YouTube – or maybe not bleach, but some other quack remedy, that there's very good reason to think, not only is it ineffective, but in fact, very harmful to people -- that if people follow this advice, they and their families may die or become very ill...

So yeah, that's a question of truth, but it's also a question where you say, "All right, well, to what extent do you think the platforms reasonably bear some measure of moral responsibility?" At least if they see that a claim like this is getting millions of views. Where foreseeably, if even a small percentage of the people who are looking at this believe it, and try it, they will suffer real physical harm.

And on the one hand, yeah, I don't like the idea of Facebook and Twitter being, in general, arbiters of truth. But it's not that hard for me to think of particular circumstances where I think it would be frankly derelict for them not to say, "Nothing is a hundred percent, maybe there's some possibility this is right, but based on a fairly high level of certainty, and also the dire consequences to people if they act on this, we're going to take responsibility for shutting that down and saying, 'We don't want to be the cause of people poisoning themselves.'"

But yeah, that's a different scenario from, let's say, claims about electoral fraud. Again, on the one hand, you have genuine harm that is occasioned if very large numbers of people believe that peaceful democratic participation -- persuasion and casting ballots -- is no longer a viable way to express influence on the political system. Some percentage of those people will -- maybe reasonably, given that belief set -- come to view radical and maybe violent action as the only alternative. And that's a real harm.

On the other hand, electoral fraud is also the kind of thing where there genuinely may be information that is best uncovered via social media. That is, there really may be information individuals out there in the world have, that is relevant to our assessment of whether it is true that there was electoral fraud.

And so, maybe you're reasonably confident that there was not widespread electoral fraud. When I have that argument, I think it's fairly clear that these are meritless claims and that there was not fraud at scale.

But one of the arguments I would offer for that is: Look, it's just not plausible that something could be done on this scale. That you could coordinate across so many different states and so many different actors, in a way that's necessary to produce the sheer number of physical paper ballots to shift the result of a national election – so, tens of thousands in every state -- it's just totally not plausible that this would happen and you would be able to

successfully keep it secret. That word would not get out, that someone would not notice.

And so, to the extent you have the platforms deciding at some point, "Look, these are lies, this didn't happen, and because of the harms of people believing this, we're going to limit the spread of that information" ...that also, to some extent, kicks the legs out from one of the arguments I would offer for our ability to have confidence that these claims have no merit. Which is, it would just be so impractical – that people are able to share their experiences and their information about what they've seen in their jurisdictions.

And then, of course, that has to be tested. Because a lot of those people don't understand what they're looking at. One of the things that became clear in these affidavits... is that you had a lot of people who say, "I saw nothing but these disturbing fraudulent actions happening constantly," and then the election officials explain what they saw was not abnormal or evidence of fraud. They didn't see people checking signatures here, because the signatures were checked at a different stage of the process, and if they understood how the process worked, they wouldn't have been concerned.

But still, you have those claims tested, and when you've been able to hear all the things people were concerned about that they saw, and you go down the list and say, "Okay, yeah, they've had the opportunity to offer their evidence and it turns out that evidence is not particularly compelling," then, okay, you can be pretty confident that these claims don't have a whole lot of merit.

We can be pretty confident that large scale fraud is very unlikely, in a way that you couldn't be quite as confident if anyone who offered a claim like that were being shut down, and having their account deleted immediately.

[musical interlude]

Julia Galef: That was Julian Sanchez, senior fellow at the Cato institute.

And I do feel like I have more clarity now after our conversation about how to think about the power tech companies have over free speech. To summarize, we talked about three main factors that should affect how worried you should be about tech companies deplatforming people, and I think Julian and I agree about all three of these factors, at least directionally:

First, all else equal, you should be more worried about deplatforming "lower in the stack." So a company like Twitter banning users for their speech is less worrisome than a company like Amazon Web Services banning apps because of their users' speech, which is less worrisome than an Internet Service Provider deciding "We're not going to offer service to you because of your speech." The lower in the stack you go, the fewer options you have to switch to, so the more power a single company has over you and over the boundaries of acceptable speech.

Then, second: You should be more worried about deplatforming for "misinformation" than for hate speech or inciting violence. Because tech companies are really not equipped to be the ultimate authorities on what is true and what is misinformation, and also, as Julian was arguing at the end of the conversation, in some cases banning "misinformation" is self-undermining, because the whole reason we can be at all confident that it IS misinformation is that people have had an opportunity to discuss it openly. That's an interesting point I hadn't noticed before.

And thirdly, you should be less worried about deplatforming the more you believe that, as Julian put it, "the market will solve." That the norms of what tech companies allow you to say will never stray too far from the norms of the public, because if they do, that will give rise to competitors with norms people like better.

I think this third criterion is one where I'm less convinced than Julian – that it still seems to me that the power of network effects and the fact that a small group of tech companies are subject to the same set of pressures that may not be representative of the public as a whole – that those factors still leave me kind of uneasy about the situation we're in and not fully confident in the market's ability to undercut that. And my conversation with Julian did move me somewhat towards his view, but not completely. So I still have to think about that.

You can read Julian's writing at juliansanchez.com, and definitely follow him on Twitter – his handle is normative, N-O-R-M-A-T-I-V-E. I'll link to those sites, and to the A. O. Hirschman book he discussed, "Exit, Voice and Loyalty," on the podcast website.

That's all for this episode of Rationally Speaking! I hope you'll join me next time for more explorations on the borderlands between reason and nonsense.