Written Statement of

Ari Ezra Waldman
Professor of Law
Director, Innovation Center for Law and Technology
New York Law School

Before the

Committee on the Judiciary
Subcommittee on the Constitution and Civil Justice
United States House of Representatives

For a hearing on

"The State of Intellectual Freedom in America"
September 27, 2018
11:30 AM

Chairman Goodlatte, Ranking Member Nadler, and Distinguished Members of the Committee:

Thank you for inviting me to testify today on topics of vital importance to anyone with any sort of online presence—namely, all of us. My name is Ari Waldman, and I'm a law professor at New York Law School. I hold a JD from Harvard Law School and a PhD in sociology from Columbia University.[1] For the past 8 years, I have studied how technology mediates our daily lives. I study things like when and why we share personal information on social media platforms, how technology design manipulates our online behavior, how we can make online spaces safe for everyone, and the legal and regulatory environment in which data collectors, like Facebook, operate.

Let me note that although I research and write in this area, I stand on the shoulders of far smarter and more accomplished colleagues, many of whose work I will cite and rely on here.

My goal today is to help this Committee understand Facebook's editorial role, the dangerous implications of a world without content moderation, and the serious privacy implications of allowing Facebook and other social media companies to exist in a regulatory void.

I. What Should We Be Talking About?

A recent article in the magazine Wired made an important observation: "To communicate anything," the authors wrote, "Facebook can't communicate everything."[2] This has always been true of media platforms. Neither Salon nor the National Review can publish everything. Consider another analogy: Imagine if this hearing were conducted without rules or norms of appropriateness, where every member could scream into a microphone whenever she or her wanted and for however long she or her wanted, and every member of this panel could jump on this table, scream into our microphones at the same time and to whatever end. No one would learn anything, no one would get a point across. It would just be spectacle, and really bad, cacophonous, headache-inducing spectacle at that. Content moderation rules play similar roles. And they do not implicate the First Amendment in the traditional sense: We all may have a First Amendment right, subject to some limitations, to say what we want free of government intervention or censorship. But we don't have a First Amendment right to Facebook's amplification of our words.

My colleague Kate Klonick, an assistant professor of law at St. John's University School of Law, published an article in the Harvard Law Review recently in which she described how and why platforms like Facebook moderate content.[3]

---

[1] The views expressed in this testimony are my own and do not necessarily reflect the views of any of these institutions.

[2] Emma Grey Ellis & Louise Matsakis, *Diamond and Silk Expose Facebook's Burden of Moderation*, WIRED (Apr. 14, 2018), https://www.wired.com/story/diamond-and-silk-expose-facebooks-burden-of-moderation/.

[3] Kate Klonick, *The New Governors: The People, Rules, and Processes Governing Online Speech*, 131 HARV. L. REV. 1598 (2018).

First, the how. Content moderation is a complex ecosystem of technology and people. Moderation sometimes happens before content is published, in that time between you uploading a video and its publication. This is an automatic procress, using a data-trained algorithm that screens out things like child pornography, copyrighted material, or graphic violence.[4]

After publication, moderators either proactively remove content that violates platform rules, as with extremist or terrorist speech, or react when other users flag content as inappropriate. What happens then is a bit of a mystery because Facebook wants it that way. We don't know much about the precise guidelines Facebook uses to train its content moderators on how and when to remove flagged content. But we do know that Facebook employs a large layered team of people to do this work. Outsourced talent is supervised by more experienced content moderators, many of whom work outside the US or at call centers. The top level moderation happens back at headquarters, by lawyers directly responsible for content moderation policies and training. As Professor Klonick shows, these moderators work very much like a judge would: they are trained to exercise judgment based on rules set out by the platforms. And because these rules were created and promulgated by lawyers steeped in the American legal tradition, they reflect the free speech norms many of us learned in law school.[5]

Now we can move on to why platforms do this. They are, as you know, under no legal obligation to do so. But they have normative and financial incentives. Every platform designs values into its code. Some social spaces are meant to be overtly uninhibited, like the darkest corners of 4Chan. But even 4Chan has some rules. Other social spaces promulgate rules that make them family friendly, or safe for kids. Facebook has corporate values, too. One of its central values, and one to which CEO Mark Zuckerberg referred frequently during his Senate testimony two weeks ago, is to bring friends together. As a result, and in response to the manipulation of the platform by fake news sources, Facebook redesigned its News Feed algorithm to privilege and prioritize posts from our friends rather than from media or business pages.

The result is that lots of content gets filtered out, but no more so from the right than from the left. When victims of racist, homophobic, and sexist tweets and comments post those comments to call out the aggressors, it is often the victims that are suspended or banned.[6] Activists associated with the Black Lives Matter movement have reported just as many if not more take downs of images discussing racism and police brutality than any of the anecdotal evidence of suspensions or take downs on the right.[7] Facebook has a long history of taking down photos of breastfeeding mothers.[8] In 2014, the company suspended drag performers for using their drag

---

[4] *Id*. at 1636.

[5] *Id*. at 1638-43.

[6] Sarah Myers West, Nicolas Suzor, & Jillian C. York, *How Facebook Can Prove It Doesn't Discriminate Against Conservatives*, SLATE (Apr. 19, 2018), https://slate.com/technology/2018/04/how-facebook-can-prove-it-doesnt-discriminate-against-conservatives.html.

[7] *Id*.

[8] Lisa Belkin, *Censoring Breastfeeding on Facebook*, NEW YORK TIMES (Dec. 19, 2008), https://parenting.blogs.nytimes.com/2008/12/19/censoring-breastfeeding-on-facebook/.

names.[9] An advertisement for a book featuring an LGBT vision of Jesus was rejected. The artist, Michael Stokes, who is best known for his portraits of soldiers wounded in combat, has seen his portraits of gay soldiers taken down and his account blocked.[10] At a minimum, mistakes happen on the left just as much as they happen on the right.

Consider also what online social network platforms would look like without content moderation. Danielle Citron, a professor at the University of Maryland Francis King Carey School of Law, has showed how gendered cyberharassment proliferates when platforms like Twitter and 4Chan do nothing.[11] They become havens of hate that function to silence women's (and others') voices. There is nothing the drafters of the First Amendment would hate more than that! In my own work, I have seen how queer-oriented geosocial dating apps that ignore content violations like racist and transphobic profiles become havens for nonconsensual pornography, commonly known as "revenge porn," and cause untold emotional damage to victims.[12] As a result, marginalized populations are squeezed from both sides: too much or inartful content moderation takes down too many posts about queer or racial issues, leaving us silenced and invisible; too little moderation hands the platform over to violence, hate, and harassment that silences us anyway.

Why does this happen? Any content moderation that occurs algorithmically is subject to problems inherent in machine learning: biased data and an inability to understand context, for example. Data-trained algorithms that determine what we see on our News Feeds also can't tell the difference between two media articles of wildly different accuracy. They don't know off hand that an article claiming that a former presidential candidate sold weapons to ISIS is pedaling a false, click-bait conspiracy theory. All they know is that a herd of highly motivated users in tight, hyperpartisan networks are clicking on it, sharing it, liking it, sharing it again, commenting on it, and sharing it over and over again. To the algorithm—and to Facebook—this is great. Engagement is at the core of Facebook's business model because it allows the company to learn more about its users and charge higher prices for its targeted advertising tools. This problem, at least, is designed in.

When humans get involved, we are left with a couple of possibilities, none of which is an anti-conservative bias. Either the bias swings both ways or perhaps content moderation at the human level is more art than science, and Facebook doesn't have enough Rodins, Rembrandts, and Vermeers. Although Facebook needs to do better, the steps we know it is taking will not solve the problem because, as we know from Facebook CEO Mark Zuckerberg's recent testimony, the company wants artificial intelligence (AI) and data-trained algorithms to do more of the work.

---

[9] Vauhini Vara, *Who's Real Enough for Facebook*, NEW YORKER (Oct. 2, 2014), https://www.newyorker.com/business/currency/whos-real-enough-facebook.

[10] Bil Browning, *Facebook Apologizes for Banning Gay Photographer Michael Stokes Over Photos of Wounded Soldiers*, THE ADVOCATE (Sept. 24, 2015), https://www.advocate.com/arts-entertainment/2015/9/18/facebook-targets-gay-photographer-again-concerning-photos-wounded.

[11] DANIELLE KEATS CITRON, HATE CRIMES IN CYBERSPACE (2014).

[12] Ari Ezra Waldman, *Law, Privacy, and Online Dating*, 44 L. & SOC. INQUIRY __ (forthcoming 2019).

Mr. Zuckerberg frequently attempted to assure the panel of Senators that Facebook was working on artificial intelligence tools to address all sorts of problems, from fake news to online harassment. This makes sense, at least from Facebook's perspective. Facebook is a big company, and an even bigger platform. And, as Silicon Valley likes to say, AI scales, humans do not. That is, it's a lot harder (and far more expensive) to hire an army of people to effectively monitor and address rule violations than it is to write up an algorithm to spot offending content, train it on terabytes of data, and let it loose on the Facebook platform. AI is fallible, and not yet as good as you think it is. It doesn't understand context. It has trouble with idioms.[13] It can be biased, even racist.[14] An almost religious reliance on the capacity of AI and machine learning to solve Facebook's problems represents a vision with a big blind spot,[15] like when the platform asked users to send in their naked pictures so they can train their AI to spot nonconsensual nude image sharing.[16] More importantly, Facebook's problem isn't a scale problem. It didn't fail to pick up *some* misinformation, Russian agents, or data misuses because the Facebook platform is too big. It failed to identify them as bad conduct in the first place. And it did so because it operates without headwinds, friction, or regulation on how it designs its platform, gathers and manipulates data, and treats its users.

Content moderation on Facebook is part of a larger narrative about how the lack of even reasonable regulation allows Facebook to take a cavalier approach to our privacy, safety, and civic discourse. This was on stark display when Facebook allowed data on 87 million of its users to be accessed in violation of its Terms of Service.[17] So although the evidence isn't there to suggest a systemic bias when it comes to content moderation, there is evidence that Facebook, when left to its own devices, cares very little about the safety of our data. It only cares about collecting it. Reasonable steps must be taken to reign in Facebook's near unlimited power to violate our trust.[18]

The data Facebook has on us, the vast majority of which is collected passively, without us even knowing, is collected in a regulatory and legal void. We have no comprehensive data privacy law in this country. The privacy laws we do have apply in relatively narrow contexts, not applicable here. But leaving things up to Facebook is a recipe for disaster. Take a look at Mr. Zuckerberg's testimony again.

A common theme in his responses was his professed desire to give users more "control" over data. Mr. Zuckerberg used the word 54 times. All of the Senators at the hearing used it 11 times.

---

[13] James Vincent, *Why AI Isn't Going to Solve Facebook's Fake News Problem*, THE VERGE (Apr. 5, 2018), https://www.theverge.com/2018/4/5/17202886/facebook-fake-news-moderation-ai-challenges.

[14] Julia Angwin, Jeff Larson, Surya Mattu & Lauren Kirchner, *Machine Bias*, PROPUBLICA (Mary 23, 2016), https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

[15] BRETT FRISCHMANN & EVAN SELINGER, RE-ENGINEERING HUMANITY (2018).

[16] Madison Malone Kircher, *Facebook Would Like You to Send Them Your Nudes in the Name of Safety*, New York Magazine (Nov. 8, 2017), http://nymag.com/selectall/2017/11/facebook-wants-users-to-send-nudes-to-stop-revenge-porn.html.

[17] Issie Lapowsky, *Facebook Exposed 87 Million Users to Cambridge Analytica*, WIRED (Apr. 4, 2018), https://www.wired.com/story/facebook-exposed-87-million-users-to-cambridge-analytica/.

[18] *See* ARI EZRA WALDMAN, PRIVACY AS TRUST: INFORMATION PRIVACY FOR AN INFORMATION AGE (2018).

But it is not at all clear what more control means or what it will accomplish. If, by control, Facebook means more buttons to click, more confusing words to read in a privacy policy, and just more of the same tweaks we see from Facebook every time it breaches our privacy and apologizes, we are going to see something like Cambridge Analytica happen again. Privacy policies are confusing,[19] inconspicuous,[20] and no one reads them.[21] They are long[22] and difficult to understand.[23] Even privacy experts find them misleading.[24] And a litany of scholars and experts have shown that although privacy policies are at the foundation of our self-regulatory approach to data privacy and are supposed to provide us with notice of data use practices and the choice to opt out, they, in practice, provide neither sufficient notice nor choice.

II. Next Steps

On its own, Facebook is unlikely to take any measure that creates friction with user engagement. Clicks, or more broadly, engagement on the platform, is the lifeblood of the Facebook business model. The more we like, the more we comment, the more pictures we upload, and the more we simply browse the Internet, the more Facebook knows about us, the better it can train its News Feed and facial recognition algorithms, and the more money it can charge those who want to target us for advertisements. So, anything that increases engagement is good for the bottom line. Anything that gets in the way of data collection is bad for business. Hoping that Facebook will regulate itself to protect our privacy and safety is, therefore, a fool's errand.

There are, however, reasonable steps Congress and regulators can take to protect us against Facebook's worst privacy failures.

---

[19] Joel R. Reidenberg et al., *Disagreeable Privacy Policies: Mismatches Between Meaning and Users' Understanding*, 30 BERKELEY TECH. L.J. 39, 40, 87-88 (2015) ("[A]mbiguous wording in typical privacy policies undermines the ability of privacy policies to effectively convey notice of data practices to the general public.").

[20] Janice Y. Tsai et al., *The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study*, 22 INFO. SYS. RES. 254, 266-67 (2011).

[21] *See, e.g.*, George R. Milne & Mary J. Culnan, *Strategies for Reducing Online Privacy Risks: Why Consumers Read (or Don't Read) Online Privacy Notices*, 18 J. INTERACTIVE MARKETING 15, 15 (2004); Jonathan A. Obar & Anne Oeldorf-Hirsch, *The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services* (forthcoming), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2757465 [https://perma.cc/D7M2-JWSW].

[22] George R. Milne, Mary J. Culnan & Henry Greene, *A Longitudinal Assessment of Online Privacy Notice Readability*, 25 J. PUB. POL'Y & MARKETING 238, 243 (2006). Lorrie Cranor estimates that it would take a user an average of 244 hours per year to read the privacy policy of every website she visited. *See* Lorrie Faith Cranor, *Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice*, 10 J. ON TELECOMM. & HIGH TECH. L. 273, 274 (2012). This translates to about 54 billion hours per year for every U.S. consumer to read all the privacy policies she encountered. *See* Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 I/S: J.L. & POL'Y FOR INFO. SOC'Y. 540, 563 (2008).

[23] *See* Mark A. Graber, Donna M. D'Alessandro & Jill Johnson-West, *Reading Level of Privacy Policies on Internet Health Web Sites*, 51 J. FAM. PRAC. 642, 642 (2002).

[24] Reidenberg et al., *supra* note 19, at 87-88.

1. *Take design into account*. Technology design, or the architecture of a platform, is central to protecting privacy. Most data is collected behind a user interface, without us having to key in information. That means that how a product or website is designed has a significant impact on data collection and user safety. But design principles are often neglected, both in law[25] and on the ground when companies are making new technology products.[26] Requiring companies to take organizational and technical measures to embed privacy into data processing, technology products, and corporate structures is known as "privacy by design." It is part of the General Data Protection Regulation, taking effect in under a month in Europe,[27] and it has been endorsed by the Federal Trade Commission[28] and the Office of the Attorney-General of California.[29] When this Congress acts to protect the privacy of all Americans, it should include "privacy by design" mandates, as well.

2. *Regulate artificial intelligence*. We cannot hope that Facebook will adequately regulate itself. We have hoped that since 2007, and Facebook has continued to misuse our data, apologize, and go right back to its bad behavior. At a minimum, the artificial intelligence tools Facebook plans to use to right its ship must be regulated in many of the ways that University of Maryland law professor Frank Pasquale has recommended in his book, *Black Box Society*. Facebook must be transparent about how its AI works, about how and where AI is being deployed, and about the values embedded in those algorithms. If Facebook had been more transparent about its technical processes and its values, users and businesses would be less surprised when posts get taken down or when audiences narrow. But transparency is only the beginning. Facebook users deserve to understand and evaluate the values embedded in how those algorithms are designed.

3. *Recognize that we entrust our data to companies like Facebook*. Facebook is not a passive conduit of information. We recognize that we give over information to Facebook in exchange for a free service that "scratches its users social itches," to use James Grimmelmann's phrase. But that does not mean we relinquish that information with the expectation that it will be seen and used by everyone under the sun. Rather, we entrust our information to Facebook under specific circumstances and for specific purposes. Facebook should be understood as a trustee of that information, as Yale Law School Professor Jack Balkin has argued, with many of the legal implications that come with it.

---

[25] WOODROW HARTZOG, PRIVACY'S BLUEPRINT: THE BATTLE TO CONTROL THE DESIGN OF NEW TECHNOLOGIES (2018).

[26] Ari Ezra Waldman, *Designing Without Privacy*, 55 HOUSTON L. REV. 659 (2018).

[27] *See* General Data Protection Regulation, Art. 25, "Data protection by design and by default", http://ec.europa.eu/justice/data-protection/reform/files/regulation_oj_en.pdf (last visited Jan. 17, 2018).

[28] Fed. Trade Comm'n, Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers, at 22 (2010), *available at* https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf.

[29] CAL. DEP'T OF JUSTICE, OFFICE OF THE ATTORNEY GENERAL, PRIVACY ON THE GO: RECOMMENDATIONS FOR THE MOBILE ECOSYSTEM, at 1, 4 (Jan. 2013), *available at* https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/privacy_on_the_go.pdf ("Our recommendations, which in many places offer greater protection than afforded by existing law, are intended to encourage all players in the mobile marketplace to consider privacy implications *at the outset* of the design process.") (emphasis in original).