PREPARED TESTIMONY AND STATEMENT FOR THE RECORD
OF

WOODROW HARTZOG
PROFESSOR OF LAW AND COMPUTER SCIENCE
NORTHEASTERN UNIVERSITY
SCHOOL OF LAW & KHOURY COLLEGE OF COMPUTER SCIENCES

HEARING ON

"POLICY PRINCIPLES FOR A FEDERAL DATA PRIVACY FRAMEWORK IN THE UNITED STATES"

BEFORE THE

COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION
U.S. SENATE

February 27, 2019
Hart Senate Office Building 216
Washington, DC

I. INTRODUCTION

Chairman Wicker, Ranking Member Cantwell, and Members of the Committee, thank you for inviting me to appear before you and provide testimony on this important issue. My name is Woodrow Hartzog and I am a Professor of Law and Computer Science at Northeastern University's School of Law and Khoury College of Computer Sciences. I am also a Non-resident Fellow at The Cordell Institute for Policy in Medicine & Law at Washington University, a Faculty Associate at the Berkman Klein Center for Internet & Society at Harvard University, and an Affiliate Scholar at the Center for Internet and Society at Stanford Law School. I have written extensively about privacy and data protection issues, including over thirty scholarly articles, essays, and book chapters. I have specifically addressed policy principles for data privacy frameworks in a number of academic articles.[1] My recent book explores possible privacy principles for the design of data technologies.[2] My comments today will address what I've learned from this research. I make these comments in my personal, academic capacity. I am not serving as an advocate for any particular organization.

The effort to identify policy principles for a federal data privacy framework is necessary, urgent, and expansive. There are so many different issues to consider, including questions about preemption, enforcement mechanisms, regulatory structure, civil rights implications, law enforcement access, algorithmic accountability, and more. Policymakers should consider many different perspectives, including but not limited to people of color, people in the LGBTQ+ community, people with disabilities, women, and all communities that privacy law affects in different ways.

Because my time is limited, I focus my remarks on the topic I have spent most of my efforts researching over the past few years—the way our current privacy regime asks too much of people and too little of those entrusted with our data. I make two recommendations for the Committee.

First, I recommend that lawmakers should resist the traditional approach to data protection, which emphasizes transparency through notice to users and choice through user consent. It passes the risk of online interaction from data collectors onto people under an illusion of protection. This "notice and choice" approach has failed.

---

[1] Woodrow Hartzog, *The Inadequate, Invaluable Fair Information Practices*, 76 MD. L. REV. 952 (2017); Woodrow Hartzog, *The Case Against Idealising Control*, 4 EUROPEAN DATA PROTECTION L. REV. 423 (2018); Neil Richards & Woodrow Hartzog, *The Pathologies of Consent*, 96 WASH. U. L. REV. (forthcoming 2019); Neil Richards & Woodrow Hartzog, *Privacy's Trust Gap*, 126 YALE L. J. 1180 (2017); Neil Richards & Woodrow Hartzog, *Taking Trust Seriously in Privacy Law*, 19 STAN. TECH. L. REV. 431 (2016); Woodrow Hartzog & Neil Richards, *It's Time to Try Something Different on Internet Privacy*, THE WASHINGTON POST (Dec. 20, 2018). Parts of this testimony are adapted from some of this research.

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

Second, I will argue that the best path forward is to move beyond traditional procedural regimes towards substantive and robust rules that protect people's trust in entities and establish firm data boundaries that companies cannot cross without consequences.

Meaningful data privacy reform must do more than merely turn up the volume of fraught concepts like transparency, consent, and control. Second helpings of "I agree" buttons and turgid, unreadable terms of use would not have prevented the Cambridge Analytica debacle, the epidemic of data breaches, or the harmful decisions and predictions made by wrongfully biased algorithms powered by our data. Nor will they prevent the problems of manipulation, discrimination, and oppressive surveillance that we face in a future of automation. Lawmakers should instead create non-waivable robust and substantive duties and data mandates for companies.

II. NOTICE AND CHOICE IS IRREPARABLY BROKEN

The state of privacy in the United States is bad and getting worse. For years, privacy failures and data breaches have trickled in. But recently the fragile wall that policymakers constructed forty years ago to mitigate the risks of databases is cracking. To many, the answer to our modern privacy dilemma is simple: give users more control. From social media to big data to biometrics, proposed solutions include some combination of "privacy self-management" concepts like control, informed consent, transparency, notice, and choice.[3] These concepts are attractive because they seem empowering. They promise to put people in charge of their own data destiny. While notice and choice regimes facilitate data exchange, they have left people exposed and vulnerable. Meaningful progress requires more.

In basing policy principles for data protection on notice and choice, privacy frameworks are asking far too much from a concept that works best when preserved, optimized, and deployed in remarkably limited doses. Our personal agency is required for self-management concepts to work and, after all, we are only human. Even under ideal circumstances, our consent is far too precious and finite to meaningfully scale.

A. Notice and Choice Models Will Never Work for Data at Scale

The problem with notice and choice models is that they create incentives for companies to both hide the risks in their data practices though manipulative design, vague abstractions, and excessive and complex words while at the same time shifting risk by engineering a system meant to expedite the transfer of rights and relinquishment of protections.

But even the idealized, perfected transparency and control models contemplated by these frameworks is impossible to achieve in mediated environments. There are several reasons why.

---

[3] *See* Daniel J. Solove, *Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880 (2013).

First, the control companies promise people is an illusion. Entities inescapably engineer their technologies to produce particular results. People's choices are constrained because others design the tools they use. People get boxes to check, buttons to press, switches to activate and deactivate, and other settings to fiddle with.

Data collectors have incentives to make users believe they have more control than they are actually given. People can only click on the options that companies create for them. Think of how parents create this illusion of control for their children, such as when I give my kids a choice between going to the park or the movies. They feel empowered and I avoid a trip to the pet store so I can stave off a conversation about a new puppy for one more week.

Data collectors also have incentives to leverage design to extract our consent. Companies deploy "dark patterns" and exploit our built-in tendencies to prefer shiny, colorful buttons and ignore dull, grey ones. They may also shame us into feeling bad about withholding data or declining options. Many times, companies make the ability to exercise control possible but costly through forced work, subtle misdirection, and incentive tethering.[4] Sometimes platforms design online services to wheedle people into oversharing, such as keeping a "streak" going or nudging people to share old posts or congratulate others on Facebook. Companies know how impulsive sharing can be and therefore implement an entire system is set up to make it so easy.

Second, notice and choice regimes are overwhelming. They simply do not scale because they conceive of control and transparency as something people can never get enough of. People are gifted with a dizzying array of switches, delete buttons, and privacy settings. We are told that all is revealed in a company's privacy policy, if only we would read it. After privacy harms, companies promise more and better controls. And if they happen again, the diagnosis is often that companies simply must have not added enough or improved dials and check boxes.

Control over personal information is attractive in isolation. But often it's a practical and overwhelming obligation. If people do not exercise that control, they are at risk. Companies can take your inaction as acquiescence. While you might remember to adjust your privacy settings on Facebook, what about Instagram, Twitter, Google, Amazon, Netflix, Snapchat, Siri, Cortana, Fitbit, Candy Crush, your smart TV, your robot vacuum cleaner, your WiFi-connected car, and your child's Hello Barbie? Mobile apps can ask users for over two hundred permissions and even the average app asks for about five.[5] The problem with thinking of privacy as control is that if we are given our wish for more privacy, it means we are given so much control that we choke on it.

One remedy policymakers have proposed is to make all choices privacy protective by default. However, even if the default works, ceaseless demands are still making us

---

[4] For more information on the concept of dark patterns, see Harry Brignull's http://www.darkpatterns.org.

[5] Kenneth Olmstead and Michelle Atkinson, *Apps Permissions in the Google Play Store*, PEW RESEARCH CENTER, (Nov. 10, 2015), http://www.pewinternet.org/2015/11/10/apps-permissions-in-the-google-play-store/.

relent.[6] Anyone that has turned off notifications on their mobile apps can attest to the persistent, grinding requests for the user to turn them back on almost every time they open the app. Many can relate to the experience of a child asking for candy, over and over, until the requests become too much to ignore and we give in, simply to quiet them. Willpower can feel like a finite, vulnerable, and subjective resource, and companies design systems to deplete and erode it. Once our willpower and ability to make choices has been compromised, the control we have been given is meaningless.

Even if a company achieves the platonic ideal of how to give data subjects' notice and choice, it wouldn't solve people's limited bandwidth dilemma. People only have twenty four hours in a day and every service they use wants them to make choices about how they can handle their data. Meaningful individual control over one data flow between a person and a data collector won't change the fact that the data ecosystem is vast. And it should be if the market is to be competitive. The modern data ecosystem is mind-bogglingly complex, with many different kinds of information collected in many different ways, stored in many different places, processed for many different functions and shared with many other parties. And even if every tech company merged together until only one giant tech company existed, the tension between simplicity and nuance in privacy policies would seem irresolvable. This is because when companies try to simplify and shorten information nuance is lost. Risk is either hidden in terms of use through abstraction or made so explicit and voluminous we don't even know where to begin.

The collective result of atomized online decisions is not the best guide for privacy policy. Research shows that peoples' privacy preferences are uncertain, contextually dependent, and malleable.[7] The availability of knowledge doesn't necessarily translate into meaningfully informed decisions. People will always know less than companies regarding the wisdom of a decision. However, notice and choice regimes ask them to consider the privacy implications of each post they create and every action they take online—an impossibly complex calculation to make about future risks and consequences. In a world of predictions and group privacy, sometimes a person's consent is beside the point. For example, when members of my family consent to the practices of genetic testing companies by sending their DNA off for analysis, the DNA overlap implicates my privacy, but my consent is irrelevant.

Defending notice and choice regimes requires so much justification, so much stretching, bending, and tying ourselves in knots, that it feels like it's merely serving as a proxy for some other protection goal that is just out of reach. But it is not clear what the result that policymakers, industry, advocates, and data subjects are in truth hoping for. Control ostensibly serves autonomy, but in mediated environments involving personal data, idealizing control actually seems corrosive to autonomy. Is control valuable because people have such different privacy preferences? Or does it just appear that way because personal data risks are virtually impossible for people to consistently assess?

---

[6] Article 25 of Regulation (EU) 2016/679 on data protection by design and by default. [2016] OJ L119/1. *See also* Recital 78 of Regulation (EU) 2016/679.

[7] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein, *Privacy and Human Behavior in the Age of Information*, 347 SCIENCE 509 (2015).

If data processing is so dangerous that it requires formal permission, and meaningful choices can only be made in elusive, demanding, and bounded environments with preconditions such as "freely given, specific, informed, retractable, and unambiguous," then it is worth asking why companies are allowed to engage in what feels like a fiction, even under optimal conditions. Is notice and choice just a contorted and indirect way to pressure companies to lay off risky data practices? If so, lawmakers should dispense with the pretense of demanding a form of notice and choice that seems destined to misdirect industry efforts towards formalistic compliance without a meaningful change in processor behavior.

B. The Fair Information Practices are Necessary But Not Sufficient

The push for consent and control partially springs from the original principles used to ensure fair data processing, referred to as the "Fair Information Practices" or the FIPs. These aspirational principles developed over the past fifty years are used to model rules for responsible data practices. They are the bedrock of modern data protection regimes around the world: Transparency of business practices; Access and correction rights; Data collection and use limitations; Accountability; Data minimization and deletion; Data accuracy; Purpose specification; and Confidentiality/security.

The FIPs provide a common set of values, which is necessary as data flows from one country to another at the speed of light. The FIPs provide a benchmark for industry, advocates, and policymakers to analyze new technologies. Privacy as a general concept is vague and messy. But the FIPs are a little more concrete. This clarity gives everyone a more useful litmus test for determining whether companies are being responsible with people's data. In short, the FIPs are invaluable for the modern world.

The FIPs have also painted lawmakers into a corner. A sea change is afoot in the relationship between privacy and technology. FIPs-based regimes were relatively well-equipped for the first wave of personal computing. But automated technologies and exponentially greater amounts of data have pushed FIPs principles like data minimization, transparency, choice and access to the limit. Advances in robotics, genetics, biometrics and algorithmic decision making are challenging rules meant to ensure fair aggregation of personal information in databases.

While the FIPs are a necessary as a component of a Federal Data Privacy Framework, they are not sufficient for several reasons. First, the FIPs have several blind spots. They are largely focused on data aggregation by industry. They do not directly contemplate peoples' vulnerabilities to each other on platforms like social media, peoples' susceptibility to manipulation, and issues of platform power. Anthropomorphized robots, fMRIs that measure brain activity, and advances in genetics raise problems like people's susceptibility to things that look and act human, their inability to hide thoughts, and discrimination based on speculative predictions and forecasting of things that haven't even happened yet. These problems are generally beyond the scope of the FIPs.

Often, traditional data protection frameworks are so focused on the individual that they overlook important social and civil rights implications of collecting and processing personal data. Marginalized communities, particularly communities of color, shoulder a disproportionate burden from privacy abuses.[8] I would recommend frameworks that go beyond narrow and individualized conceptions of privacy to incorporate more communal and civil rights-based protections.

We are only just beginning to see the human and societal costs of massive scale of data processing and platform dominance. In addition to core privacy related harms associated with data collection and data use, companies' insatiable hunger for personal information is negatively affecting our attention and how we spend our time, how we become informed citizens, and how we relate to each other. Phenomena like "fake news," "deep fakes," non-consensual pornography and harassment, oversharing on social media, addiction by design, and lives spent staring into our phones are at least partially attributable to or made worse by the personal data industrial complex. We need broader frameworks for personal data not just because information is personal to us, but because the incentive to exploit it creeps into nearly every aspect our technologically-mediated lives.

The upshot is that existing data protection frameworks are important to build on, but they are still incomplete. That's why states play such a crucial role in the development of privacy policy in the U.S. Not only have states become quite involved in creating innovative privacy rules and frameworks, but they help carry the load of enforcement. Legislation that preempts the states privacy regulatory and enforcement efforts could have net negative effects for privacy as well as jeopardize the international flow of data if U.S. privacy law appears weaker as a result. Diluted preemptive federal law risks diminishing currently strong state rights and rules. State Attorneys General have played a key role in crafting and enforcing privacy rules.[9] States have also shown a willingness to exercise new and innovative approaches to privacy law, including pioneering breach notification statutes and biometric privacy protections. Baseline legislation would follow the tradition of having the federal government create a floor, not a ceiling, for privacy rules.

One temptation might be for lawmakers to seek a singular set of data protection rules to ease the cost of compliance. While harmonization of data protection rules into a national or even global standard would have benefits, an unwavering commitment to harmony with other regimes makes future progress difficult. Ossification of this sort would be fine if notice and choice were all the world needed to prepare for our future of algorithms and data. But U.S. privacy law needs more.

---

[8] The Leadership Conference on Civil & Human Rights, Letter to Senate and House Chairs Wicker, Graham, Pallone, and Nadler, and Ranking Members Cantwell, Feinstein, Walden, and Collins (Feb. 13, 2019), https://civilrights.org/resource/address-data-driven-discrimination-protect-civil-rights/.

[9] *See* Danielle K. Citron, *The Privacy Policymaking of State Attorneys General*, 92 NOTRE DAME L. REV. 747 (2017).

III. PRIVACY RULES SHOULD MOVE BEYOND CONSENT AND TRANSPARENCY

Notice and choice frameworks and overly-procedural privacy laws have resulted in a sea of blindly-clicked "I Agree" buttons, unread fine print, and constant anxiety about what our home Internet-of-Things device is listening to. What seems sensible on a case-by-case basis will in the aggregate continue to overwhelm people who simply want to be safe when they go online.

What the United States needs from federal legislation is a set of rules that reallocates the obligations from people to companies who invite their trust. Data collectors and processors are in the better position than we are to foresee how their tools and practices might be used in ways adverse to us. They are also in the better position to correct course.

The U.S. needs a set of rules with a firmer moral floor that furthers many different privacy-related values and discourages treating privacy as a procedural and formalistic compliance exercise. We should not distill privacy protection into a rote checklist that allows data collectors to flash an insignia of compliance without making more substantive efforts to protect our vulnerabilities. To that end, I recommend rules structured to protect peoples' trust and rules that place robust, clear, and non-waivable boundaries around how information technologies are designed, what data companies may collect, and how that data may be used.

A. Trust Rules Can Help Create Safe and Sustainable Information Relationships

There are ways to balance utilizing data and protecting people, but it requires a framework that reimagines the relationship between people and the companies they interact with in a way that places trust at the center. Being trustworthy in the digital age means companies must be discreet with our data, honest about the risk of data practices, protective of our personal information and, above all, loyal to us — the data subjects.[10] Our privacy frameworks should be built to encourage and ensure this kind of trustworthy conduct.

The United States needs to improve its poor international reputation regarding privacy. The world is watching, and the economic stakes are enormous. International data flows are essential for the global economy to function without fundamentally — and expensively — restructuring the Internet. American tech companies depend on being able to smoothly bring Europeans' data here for processing. But our current data-sharing agreement with Europe, the E.U./U.S. Privacy Shield, seems to be on thin ice. If it fails,

---

[10] For more information on taking trust seriously in privacy law *SEE* ARI EZRA WALDMAN, PRIVACY AS TRUST: INFORMATION PRIVACY FOR AN INFORMATION AGE (2018); DANIEL SOLOVE, THE DIGITAL PERSON: TECHNOLOGY AND PRIVACY IN THE INFORMATION AGE 102-104 (2004); Ian Kerr, "Personal relationships in the Year 2000: Me and My ISP" in NO PERSON IS AN ISLAND: PERSONAL RELATIONSHIPS OF DEPENDENCE AND INDEPENDENCE (2002); Neil Richards & Woodrow Hartzog, *Privacy's Trust Gap*, 126 YALE L. J. 1180 (2017); Neil Richards & Woodrow Hartzog, *Taking Trust Seriously in Privacy Law* 19 STAN. TECH. L. REV. 431 (2016); Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. DAVIS L. REV. 1183 (2016).

we will need a good replacement grounded strong enforcement and effective, protective, and substantive rules.

Trust rules could not only help America establish its own robust privacy identity, but it can also serve to mutually benefit industry, people as individuals, and society at large by nourishing safe and sustainable information relationships. Concepts like "big data" and machine learning seem exciting to many who work in technology. Data promises to revolutionize our work and finances, improve our health, and make our lives easier and better. But the scale and complexity of these concepts can also be scary and intimidating. The public's paranoia should be understandable. We are perpetually unsure about organizations' motives and methods. What do companies know about us? Are they keeping their insights safe from hackers? Are they selling their insights to unscrupulous parties? Most importantly, do organizations use our data against us? Like so many things in life, these relationships are a matter of trust.

In my research with Neil Richards, we have argued that trustworthy data stewards have four characteristics that promote trust: they are honest, discreet, protective, and loyal. Trustworthy stewards are honest because they forthcoming about the most important information for our well-being being, even if it might discourage use. Honesty rules place the obligation of being understood on the steward, rather than on peoples' ability to scrutinize the dense, vague, and protean language of privacy policies and terms of service. Stewards are also discreet because they treat our data as presumptively confidential and sensitive and do not disclose it in ways contrary to our interests or expectations. Discretion could involve robust de-identification efforts as well as nondisclosure. Stewards are protective because they hold the data securely against third parties, doing everything within reason to protect us from hacks and data breaches. Most fundamentally, keeping a trust requires loyalty. This involves data collectors putting the interests of those who trusted them with their data ahead of their own short-term potential for gain. Loyalty obligations would prohibit data collectors from, among other things, leveraging peoples' own resource and cognitive limitations against them or engaging in unreasonable self-dealing when collecting and processing data.

To be effective, trust frameworks should also give regulators the resources they need to enforce privacy protections and prohibit companies from using dense terms-of-use agreements to get us to waive those obligations. Companies should be trustworthy regardless of what we "agree" to online.

B. Data Boundaries Can Restore Balance

The modern data ecosystem is a runaway train. Trust rules can help, but they will not be enough. Some practices might be so dangerous that they should be taken off the table entirely. Others might be harmful to society in ways that don't implicate a violation of any trust. To be fully responsive to modern data problems, a meaningful federal privacy framework needs to embrace substantive boundaries for data collection and use.

In some contexts, this might mean rules that simply prohibit certain practices. For example, lawmakers could outright prohibit collection or aggregation of certain kinds of data, such as biometric and genomic data. Lawmakers could mandate deletion. They could get serious about purpose limitations and requiring companies to have a "legitimate interest" in processing data. And in cases where technologies represent a grave danger to civil liberties, they should not rule out a moratorium or ban. Strong rules limiting collection and storage on the front end can mitigate concern about the privacy problems raised through data analytics, sharing, and exploitation.

Finally, without structural support, resources, and a strong political mandate for enforcement, any data protection framework will merely be a pretext for exploitation. Whether legislation creates a new data privacy agency or emboldens existing federal agencies, regulators must have broad grants of authority, including rulemaking provisions where necessary, robust civil penalty authority, and the ability to seek injunctions quickly to stop illegal practices. Individuals should have private causes of action and rights as data subjects.

IV. CONCLUSION

Lawmakers must leave notice and choice frameworks behind in order to meaningfully address privacy in the United States. Companies have consistently hailed strategies for increasing transparency and control as solutions to our privacy woes, but time and time again doing further exacerbated the problem. People need to be able to trust the entities they deal with online and feel safe when they share information. Data protection rules should enforce that trust and crate substantive boundaries in service privacy as well as more diverse values like civil rights, due process, and psychological well-being. If people and groups are protected instead of saddled with the risk of online interaction, our digital ecosystem can become sustainable.

BIOGRAPHY

Woodrow Hartzog is a Professor of Law and Computer Science at Northeastern University School of Law and the Khoury College of Computer Sciences. He is also a Non-resident Fellow at The Cordell Institute for Policy in Medicine & Law at Washington University, a Faculty Associate at the Berkman Klein Center for Internet & Society at Harvard University, and an Affiliate Scholar at the Center for Internet and Society at Stanford Law School.

His research focuses on privacy, media, and technology. His work has been published in scholarly publications such as the *Yale Law Journal, Columbia Law Review, California Law Review*, and *Michigan Law Review* and popular publications such as *The Washington Post, BBC, CNN, The Guardian, Wired, The Atlantic* and *The Nation*. He has been quoted or referenced in numerous articles and broadcasts, including *NPR, The New York Times, The Los Angeles Times,* and *The Wall Street Journal*.

He holds a Ph.D. in mass communication from the University of North Carolina at Chapel Hill, an LL.M. in intellectual property from the George Washington University Law School and a J.D. from Samford University's Cumberland School of Law.

He is the author of *Privacy's Blueprint: The Battle to Control the Design of New Technologies,* published in 2018 by Harvard University Press.