UNITED STATES OF AMERICA
Federal Trade Commission
WASHINGTON, D.C. 20580

Office of Commissioner
Rebecca Kelly Slaughter

THE NEAR FUTURE OF U.S. PRIVACY LAW

Remarks of Commissioner Rebecca Kelly Slaughter[1]

Silicon Flatirons—University of Colorado Law School

September 6, 2019

Good morning! I am Rebecca Kelly Slaughter, and I have the honor of serving as a Commissioner on the United States Federal Trade Commission. I want to thank Silicon Flatirons and the University of Colorado Law School for hosting today's important event. It is an honor to be here and I welcome the opportunity to talk about the near future of U.S. privacy law.

Along with all four of my fellow Commissioners, I was sworn in to my job about a year ago. This is quite the time to be at the Federal Trade Commission. The same summer we began our jobs, GDPR went into effect, and the CCPA was signed into law. Hardly a day passes without headlines about some newly revealed data breach, a tech company practice that compromises consumer privacy, or a merger between companies that control enormous amounts of consumer data. The steady drumbeat of these stories shows what we at the FTC—the federal agency with primary responsibility over data protection issues—know to be true: This is a moment of weighty responsibility for the agency, but it is also one of opportunity.

I want to take my time today to share a little bit about how I believe the agency should meet this moment. I will start by laying out three observations about consumer data that inform how I think about both policymaking and enforcement. First, our concern needs to extend beyond a narrow concept of privacy to data abuse more broadly. Second, it is time for the reign of notice and consent to end. And third, as we consider what should replace notice and consent, we need to be especially careful to consider how data abuses affect vulnerable populations.

Then I will lay out the tools the FTC currently has at our disposal to protect against data misuse and abuse—as well as those that are critically missing from our toolbox. I will share a bit about how I hope Congress will fill in those gaps. But, even without Congressional action, we need to make the most of the authorities that we have today, both in prosecuting cases and in writing rules, so to wrap up I will share how I think we should approach those challenges.

Data Abuses Are Broader than Privacy

I know that we are here to discuss the near future of US *privacy* law, but my first observation is that I think we need to tweak the framing a little bit. Rather than simply thinking

[1] The views expressed in these remarks are my own and do not necessarily reflect the views of the Federal Trade Commission or any other commissioner.

narrowly about data privacy, I want us to be thinking in terms of data *abuses* more broadly. Privacy generally refers to limits on the collection or sharing of data that an individual would prefer to keep private. But we cannot and should not separate problems involving *collecting* data *about* individuals from problems involving the *targeting* of information *to* individuals or other decisions made *for* individuals (often based on the collected data).

Let me share an anecdote to illustrate this point: My seven-year old son is into jigsaw puzzles, but doing a traditional one is a high-risk proposition in my house with a "helpful" five-year old sister, as well as a roving toddler who will at best hide and at worst eat the pieces. So I wanted to get him a digital puzzle app. I found two options: one free app that was ad-supported, and one for which I had to pay. I will confess to being relatively cheap, so I downloaded the free one and set him off to solve.

A little while later, my husband came over asking what on earth I had put on our son's device. My husband had overheard my son listening to some pretty aggressive propaganda decrying the perils of women working outside the home (not my usual messaging, you will be shocked to know). When he asked Teddy what was happening, Teddy explained that to get more virtual coins to buy new puzzles in his app he just needed to watch a few short videos. "No big deal, Dad!" Needless to say, we deleted that app and replaced it with the paid, ad-free version. It was easy enough for us, but not everyone has the resources to do so.

This is just one example of an abusive data practice that does not fall squarely in the traditional orbit of "privacy" but is closely related and must be considered in tandem. The targeting of manipulative content to individuals—whether it is political or commercial—is a problem that disproportionately harms vulnerable populations (in this case, children and potentially lower-income individuals), but it is not the only one. It is also one that notice and consent does nothing to address.

Limitations of Notice and Consent

This brings me to the second observation I want to discuss: It is time for us to move past notice and consent as a panacea for data abuses. Much of our FTC Act authority and some of our privacy rules have, up to this point, been grounded in the principles of notice and consent. The notice-and-consent framework began as a sensible application of basic consumer protection principles to privacy—tell consumers what you are doing with their data, secure consent, and keep your promises.

But in order for a notice-and-consent regime to be effective, each element must be meaningful—notice must give consumers information they need and can understand, and consumers must have a choice about whether to consent. When it comes to our digital lives today, neither notice nor consent feels particularly meaningful.

As every consumer knows, notice is mostly in the form of lengthy click-through contracts. Few consumers can dedicate the time and legal parsing required to understand them.[2] And choice is illusory at best. Consumers do not actually have bargaining power—even if they

[2] A widely-cited article calculated that it would take a consumer 76 work days to read all the privacy policies she encounters each year. Alexis C. Madrigal, *Reading the Privacy Policies You Encounter in a Year Would Take 76 Work Days*, The Atlantic (Mar. 1, 2012), https://www.theatlantic.com/technology/archive/2012/03/reading-the-privacy-policies-you-encounter-in-a-year-would-take-76-work-days/253851/.

could read and understand the lengthy terms of contracts they must sign, their options are only to agree and access the service or to refuse and be denied access.

This means that consumers often must cede all control over their data to participate in or use certain services that are critical to their everyday lives. They do not have the ability to bargain, nor can they turn to a competing, more privacy-protective service; in too many cases, there is *no* viable competing service (an important reminder that we have serious competition problems to tackle in this space as well).

Choice is illusory in other ways as well: Many sites are designed to optimize the number of "opt-ins," including through "dark patterns," where tricks are employed by designers or developers to make users do something they otherwise would not want to do. In other words, what feels like "choice" may in fact be the product of manipulation.

We can consider a real-world example of increased opportunities for notice and consent when we look at what happened when GDPR was implemented last year. That law has the laudable goal of improving consumers' control over their data. In practice, the rollout resulted in a significant increase in opt-in consent requests any time a consumer opened a website. What was the result? People became numb to the questions; the "opt-in" too often became mere friction for consumers to ignore.

Finally, a data regime built entirely on notice and consent puts all of the burden on consumers to protect their privacy even though consumers have very little control over that data. Companies can and do track consumers across their devices and locations, and data about consumers is shared, sold, or used for targeting. Much of this happens between and among companies with which consumers never choose to interact. The companies that have control over data should have the burden of properly using and protecting it.

So I believe we must move beyond notice-and-consent as our governing principle in privacy. As a general rule, we should consider reasonable consumer expectations about data collection and data use as guideposts. Thoughtful purpose and use limitations will also be critical to protect consumers, especially when it comes to sensitive data. For example, biometrics and location information are among the categories of data that I believe should be subject to more robust protections than generic notice-and-consent. Of course there may be some places where clear, prominent, plain-English notice and true consumer consent can play a valuable role. But the onus to carefully protect data should be on the companies that collect, use, and share it, not to the consumers alone.

Data Abuses and Vulnerable Populations

The reason I feel strongly about consumers not bearing the full or exclusive burden to protect themselves from data abuses is closely tied to my third observation: Data abuses disproportionately affect vulnerable populations. We need to ensure that our laws and their enforcement reflect democratic values and principles of equality and thereby protect everyone.

One of the reasons I am particularly proud to serve as an FTC Commissioner is the agency's long history of making sure our consumer protection mission reaches all consumers. Our staff works hard to consider whether vulnerable groups might be getting hit harder by certain illegal practices, which groups might be underreporting complaints, and which groups we might be missing in our outreach. This is a continual process at the FTC—we know that we need

to be proactive to ensure that our efforts reach consumers who, for a variety of reasons, may be more vulnerable to bad practices or less visible to law enforcement. Striving to serve as a source of protection and empowerment for those left behind is not just an agency mission—it is a core value for me personally as well.

In the data protection context, this mission requires studying and acknowledging the ways certain harms fall disproportionately on disadvantaged or vulnerable populations—such as children, lower-income consumers, people of color, the LGBT community, immigrants, veterans, and our seniors. And, even more challenging, it requires us to consider how we can safeguard against a default system where the privileged are more protected from data abuses. A world where the privileged pay for access to services with their dollars and everyone else pays with their data—or worse, by suffering through manipulative content—is simply not acceptable.

Let me expand on the concept of disproportionate harm, because it goes well beyond the rogue ads I mentioned earlier. In 2016, the FTC published a report[3] that focused in part on the negative effects data collection can have on low-income and underserved populations. In the years that followed, these negative effects have only grown, including:

- Individuals being denied opportunities based on the actions of others.[4]

- Discriminatory algorithms and data practices foreclosing important life opportunities such as jobs and loans.[5]

- Fraud[6] or predatory payday lending targeting vulnerable consumers based on their personal data or demographic characteristics.

- Outsized impact of data breaches on lower-income individuals. Low-income victims of identity theft often face limited time and resources to access help and face harsher

[3] Fed. Trade Comm'n, Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues ("*FTC Big Data Report*") 9–12 (2016), https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf.

[4] *See FTC Big Data Report* at 9 (noting that big data could result "in more individuals mistakenly being denied opportunities based on the actions of others" and pointing to concerns raised by commenters that "some credit card companies have lowered a customer's credit limit, not based on the customer's payment history, but rather based on analysis of other customers with a poor repayment history that had shopped at the same establishments where the customer had shopped.").

[5] *See* Jeffrey Dastin, *Amazon scraps secret AI recruiting tool that showed bias against women*, Reuters (Oct. 9, 2018), https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G ("But by 2015, the company realized its new system was not rating candidates for software developer jobs and other technical posts in a gender-neutral way. That is because Amazon's computer models were trained to vet applicants by observing patterns in resumes submitted to the company over a 10-year period. Most came from men, a reflection of male dominance across the tech industry."); *see also FTC Big Data Report* at 9–10.

[6] *See FTC Big Data Report* at 10–11.

consequences for the types of harms that identity theft can cause: credit damage, collection efforts, and depletion of funds.[7]

- The collection of data about and targeting of information towards children.

All of these problems are compounded by the fact that it is very difficult for any of us to know what data has been amassed about us and by whom, and even harder to correct mistakes.

We must consider how to mitigate these disproportionate effects. For example, the right to access your data and seek correction is available to consumers in the U.S. on a limited basis right now, mostly cabined to credit reports. Dramatically improving the functionality of this process, applying it to personal data more broadly, and coupling it with strong enforcement could be one strategy to help protect against the spread of incorrect and harmful data.

We should also consider ways to require both visibility into and accountability for the decisions that are currently hidden behind the veil of "artificial intelligence." New draft legislation incorporates at least some of the goals of GDPR's right to an explanation for AI decisions that significantly impact individuals—though for now the discussions focus on auditing and justification rather than giving individuals specific rights.[8] We would benefit from serious consideration of the GDPR principles that protect against harms to vulnerable groups, even if we end up with different solutions.

And finally, as a society we need to take a hard look at whether the purported benefits of unlimited or unregulated behavioral advertising are worth the costs. Behavioral advertising is facilitated by the collection of vast amounts of data about individuals across platforms and devices. Its proponents tout the ability of behavioral advertising to tailor content to what an individual most wants to see, and they celebrate the increased revenue associated with behavioral advertising as facilitating high-quality content that might not otherwise come to market.

Both of these assertions merit skepticism. Delivering information to individuals that they want to see sounds fine, but targeting actually facilitates discrimination between individuals, and does so in a largely invisible way—I know what content I see, but if you are targeted with different content, I have no idea what you are seeing. Some of this targeting may be useful—for example, I personally find appealing the ad for the color-block swimsuit (ubiquitous in my social media feeds). But in the grand social scheme is that worth having teenage boys watching gamer

---

[7] *See* Madden, Mary and Gilman, Michele E. and Levy, Karen and Marwick, Alice E., *Privacy, Poverty and Big Data: A Matrix of Vulnerabilities for Poor Americans* (Mar. 9, 2017), 95 Washington University Law Review 53, 63 (2017), https://ssrn.com/abstract=2930247 ("Consider identity theft, a growing concern shared across social classes. This crime is particularly devastating for low-income individuals, who face not only financial losses that impact their ability to meet basic needs such as housing and utility services, but are also left coping with more severe consequences of someone else using their identity, such as wrongful arrests, improper child support garnishments, and harassment by collection.").

[8] *See* Adi Robertson, *A new bill would force companies to check their algorithms for bias*, The Verge (Apr. 10, 2019), https://www.theverge.com/2019/4/10/18304960/congress-algorithmic-accountability-act-wyden-clarke-booker-bill-introduced-house-senate.

videos targeted for white supremacist recruitment?[9] I hardly think so.

I also query whether behavioral advertising really enables the creation of valuable content that would be unavailable if incentivized only by traditional, contextual advertising. I want to call attention to a recent finding presented by Professor Alessandro Aqcuisti of Carnegie Melon University: The percentage of higher revenue generated from behavioral advertisements versus contextual advertisements may be quite small.[10] And that does not account for the increased costs associated with facilitating targeted behavioral ads—let alone the societal costs.

There are other studies that suggest a much stronger value correlated with behavioral advertising; I think we simply do not have enough information to know for sure. But I am confident that we would benefit from serious consideration of ways to capture whatever benefits of targeted advertising exist while limiting its substantial harms. Maybe this means banning it entirely in certain contexts; maybe there is a more targeted—as it were—way to regulate data collection and use. It is certainly worth substantial thought and debate, rather than just accepting the proliferation of widespread data targeting as inevitable.

I have laid out some observations I think need to inform how we think about data use and abuse, and in doing so I have highlighted some of the ways I think our citizens are particularly vulnerable today. This brings me to the logical question of what we ought to do about it.

Federal Trade Commission's Data-Privacy Authority

Let me begin by talking about the FTC's current data privacy authority and enforcement agenda. Today the FTC's privacy enforcement centers around the FTC Act's prohibition on unfair or deceptive acts or practices, as well as a handful of sector specific statutes—FCRA, COPPA, and the Safeguards Rule. These statutes allow us to protect children's information online and to help ensure that non-bank financial institutions and the CRAs are protecting consumer data. These statutes also give the FTC traditional rulemaking authority under the Administrative Procedure Act. In the case of COPPA and FCRA, the FTC also has the ability to seek money damages—"civil penalties"—from companies that violate the rules we promulgate.

These existing rules are important as far as they go, but they leave some gaping holes. Large categories of personal data are wholly uncovered by our rules: What we share on social media, what we share with many retailers, including our largest online retailers, and what we share with apps and devices, even when we share personal health or relationship information. And that is just the data that we intend to share. What about when our data are harvested and collected without our knowledge or expectation? In most cases, our rules do not cover these practices either.

To protect consumers' data and privacy beyond the narrow, sector-specific fields covered

---

[9] *See, e.g.*, Anya Kamenetz, *Right-Wing Hate Groups Are Recruiting Video Gamers*, NPR (Nov. 5, 2018), https://www.npr.org/2018/11/05/660642531/right-wing-hate-groups-are-recruiting-video-gamers.

[10] *See* Marotta et al., *Online Tracking and Publishers Revenues: An Empirical Analysis* at 13, Fed. Trade Comm'n: PrivacyCon Presentation (June 27, 2019), https://www.ftc.gov/system/files/documents/public_events/1415032/privacycon2019_acquisti_online_tracking_and_publishers_revenues.pdf ("After controlling for other factors, when [the] tracking cookie is available, revenue does increase[] – approximately by 4%, relative to when [the] cookie is not available.").

by our rules, we must rely on our century-old Section 5 unfairness and deception authority. We routinely use this authority to stop unfair practices that harm consumers, such as unreasonable data security practices or data tracking without consumer consent. We have brought cases to protect consumers against unauthorized and undisclosed surveillance by mobile devices, undisclosed tracking of content viewing, and numerous cases against companies that failed to secure consumer data.

Our remedies in these cases can be limited; we do not have the ability to seek civil penalties. Instead, we must make a case for consumer injury or disgorgement. Under our general FTC Act authority, we have the ability to seek civil penalties only if a company violates an existing FTC order; in other words, only a repeat-offender might pay a penalty. Even in these cases, we do not get to simply levy a fine; either we negotiate an appropriate penalty with the offender or we sue and ask a court to determine a violation occurred and weigh the violation within a range of statutory factors to assess a penalty.

The FTC staff have endeavored to be nimble and aggressive in their attempts to use this hundred-year-old statute to police today's technology-driven marketplace—with many successes. But we face real limitations proceeding under Section 5. Moreover, without specific statutes or rules defining practices in this area, both courts and companies have been left with questions about whether particular behavior is prohibited.

Because of these limitations, a majority of the FTC's commissioners has repeatedly urged Congress to pass federal privacy legislation. Specifically, we have asked for legislation that does three things in terms of FTC enforcement: (1) empowers the FTC to seek significant civil penalties for privacy violations in the first instance; (2) gives us APA rulemaking authority, to craft flexible rules that reflect stakeholder input and can be periodically updated; and (3) repeals the common carrier and nonprofit exemptions under the FTC Act to ensure that more of the entities entrusted with consumer data are held to a consistent standard.

And of course it is not just the FTC calling on Congress to act: Increased federal privacy protections enjoy widespread popular support.[11]

I know—from personal experience—that legislation takes time and that thoughtful, consensus-driven legislation takes lifetimes. The FTC will continue to use its current authorities while calling on Congress to empower us to do more. And I remain hopeful that the future holds comprehensive federal privacy legislation.

What Else Can We Do Now?

But you did not ask me to speak about the future; you asked me to speak about the "near" future. And, although I am optimistic about the prospects for federal privacy legislation, we cannot simply hold our breath and wait. So there are two things I think we need to do right away: The first is be as forward-looking and aggressive as we can be in our approach to case resolution under current law, and the second is to consider initiating a rulemaking proceeding now to address data abuses.

[11] *See* Felix Richter, *Infographic: Most Americans Support Consumer Data And Privacy Protection Law*, International Business Times (May 22, 2019), https://www.ibtimes.com/infographic-most-americans-support-consumer-data-privacy-protection-law-2794205 ("83 percent of registered voters in the U.S. agree that the country needs federal laws protecting consumer data and privacy.").

In terms of case prosecution, I hope that we will continue to prioritize cases where the harm falls disproportionately on vulnerable consumers. I also believe we should look for cases that illustrate the unfair burdens that the notice-and-consent regime imposes on consumers—and the fact that disclosures that masquerade as notice-and-consent often provide neither.

And when we evaluate how to resolve cases, we need to be considering how our actions in any given case create incentives for compliance not only by the company or individuals at issue but also by all other companies and individuals in the marketplace. That does not mean aiming for the biggest dollar amount we can efficiently extract in penalties or the speediest settlement. Instead, in each case we must carefully consider whether any particular settlement is likely to deter future wrongdoing.

I have been proud and happy to vote in favor of many settlements of data privacy cases in my limited time at the Commission.[12] But in a few cases, especially those concerning particularly large and profitable companies, I think we could and should have done more. In the two recent data privacy cases in which I dissented, the Commission had civil penalty authority but I do not believe the penalties we sought and the companies agreed to pay were nearly enough to deter future wrongdoing given the scope of the violations and the profitability of the companies. And, in both cases, I feared that the injunctive provisions did not meaningfully change the incentives for future abuse of data. I also worry that, when it comes to large companies, we have at times sacrificed a robust analysis of individual accountability for the efficiency of a settlement.

Individual liability is not about vindictiveness. It is about ensuring accountability and incentivizing a culture of compliance that starts from the very top of a company. Sarbanes-Oxley did this in the accounting sphere when it made officers personally certify the validity of a public company's bookkeeping; that is an example we should take seriously and consider applying in the data space.

Finally, I want to talk about rulemaking. As all of the admin-law wonks in the room know, most federal rules are promulgated under the Administrative Procedure Act, which provides a relatively efficient mechanism for rules to be proposed with a notice in the Federal Register, commented on by the public, and then finalized after consideration of the comments. In the 1970s, Congress removed the FTC's general ability to issue consumer protection rules under the APA; instead, it saddled us with the Magnuson-Moss Act.

Mag-Moss, as it is colloquially known, has a reputation for being like the cranky neighbor of the APA. The procedures required to issue a rule under Mag-Moss are substantially more detailed than under the APA. It requires the additional steps of a pre-rulemaking advance notice and comment period as well as a special heads-up to Congress, and public hearings among other logistical constraints. So the Commission has shied away from extensive Mag-Moss rulemaking as not worth the trouble.

But I am not convinced that this cranky reputation, however well-deserved, should scare us off the lawn. I believe the time has come to consider a Mag-Moss data-protection rule. The

[12] For example, I supported the Commission's resolutions in Unrollme, Equifax, Dealerbuilt, D-Link, ClixSense, iDressup and TikTok—everything from an email consolidation tool that misled consumers as to how their emails were handled to a popular music application that failed to take children's privacy seriously.

FTC has been incredibly innovative in its approach to consumer-data privacy from a law enforcement perspective. The agency has used a hundred-year-old statute and a handful of sector-specific laws to bring over 200 actions to protect consumers' data and privacy.[13] It has been an uphill battle, but one that has paid off on numerous fronts. The agency should be just as creative, just as dogged, in using its rulemaking tools. Congress *should* be the one to act here, but, unless and until it does, the FTC must use every existing tool—even the dull, rusty ones—to protect consumer privacy.

This path is not an easy one. This type of rulemaking initiative might take years and cost countless staff hours that would otherwise be spent on enforcement efforts. Of course, even as it continues to seek consensus on substantive privacy legislation, Congress could allocate significantly more resources to the FTC that we could use to increase enforcement while taking on this type of rulemaking initiative. The study, public commentary, and dialogue that a Mag-Moss rulemaking effort would generate would be valuable even if Congress eventually intervenes because much of the inquiry could help inform Congressional debate and any superseding rulemaking effort Congress might direct us to undertake.

The worst-case scenario here is not that a Mag-Moss rulemaking takes years to complete; the worst-case scenario is that years from now Congress has still not acted and the FTC has still not begun.

The threats to consumer privacy are growing; they impact our most vulnerable citizens more than most, and they demand new solutions. My hope is that the "near future" brings renewed action on this front across the board: from the FTC, Congress, advocates and industry and I feel both humbled and privileged to get to take part in this effort. Thank you.

[13] Fed. Trade Comm'n, Federal Trade Commission 2018 Privacy and Data Security Update 3 (2019), https://www.ftc.gov/system/files/documents/reports/privacy-data-security-update-2018/2018-privacy-data-security-report-508.pdf.