

November 21, 2022

Via Regulations.gov

Before the
Federal Trade Commission
Washington, D.C.

Commercial Surveillance ANPR, R111004

The Commission asked for comments regarding the prevalence of commercial surveillance and harmful data security practices. We, the undersigned, are interdisciplinary academic researchers from Boston University, Northeastern University, Princeton University, University of California Davis, University of California Irvine, and University of Washington with expertise spanning computer science, information science, security, and law. Contributing researchers span these multiple institutions and several multi- and cross-disciplinary projects, including: the ProperData Center (BU, IMDEA, NU, UCI, UC Davis, USC), the Center for Information Technology Policy (Princeton), Security and Privacy Lab (UW), and the Cybersecurity and Privacy Institute (NU).

In this Comment we provide suggestions for the Commission to implement new trade regulation rules (or alternatives) concerning unfair or deceptive data practices. These recommendations are based on the vast range of empirical and policy-focused research conducted by the undersigned researchers, with all cited scholarly papers written by the undersigned appended at the end of the Comment or in supplementary files for this Comment. Please find recommendations organized according to the categories of questions posed by the FTC, in numerical order.

Thank you for the opportunity to submit public comments through the ANPR; we additionally welcome any questions to the contact information below.

Signatories

Signed,**

David Choffnes*
Associate Professor, Khoury College of Computer Sciences, Northeastern University
PI, ProperData Center

Hao Cui*
PhD Student, University of California Irvine

Daniel Dubois
Research Scientist, Khoury College of Computer Sciences, Northeastern University

Alexander Gamero-Garrido

Ford Foundation Post-Doc Fellow, Khoury College of Computer Sciences, Northeastern University

Jeffrey Gleason*
PhD Student, Khoury College of Computer Sciences, Northeastern University

Johanna Gunawan*
PhD Candidate, Khoury College of Computer Sciences, Northeastern University

Woodrow Hartzog
Professor of Law, Boston University

Basileal Imana*
PhD Candidate, University of Southern California

Umar Iqbal*
Postdoctoral Researcher, Security and Privacy Lab, University of Washington

Levi Kaplan
PhD Student, Khoury College of Computer Sciences, Northeastern University

Aleksandra Korolova*
Assistant Professor of Computer Science and Public Affairs, Princeton University
PI, ProperData Center

Hieu Le*
PhD Candidate, Samueli School of Engineering, University of California Irvine

Athina Markopoulou*
Professor of Electrical Engineering and Computer Science, Samueli School of Engineering,
University of California Irvine
PI and Director, ProperData Center

Alan Mislove*
Professor of Computer Science and Senior Associate Dean for Academic Affairs, Khoury
College of Computer Sciences, Northeastern University

Hooman Mohajeri Moghaddam*
Postdoctoral Researcher, Center for Information Technology Policy, Princeton University

Shaoor Munir*
PhD Student, Department of Computer Science, University of California Davis

Muhammad Talha Paracha
PhD Candidate, Khoury College of Computer Sciences, Northeastern University

Piotr Sapiezynski*
Associate Research Scientist, Khoury College of Computer Sciences, Northeastern University

Zubair Shafiq*
Associate Professor, Department of Computer Science, University of California Davis
PI, ProperData Center

Rahmadi Trimananda*
Project Scientist, Samueli School of Engineering, University of California Irvine

Janus Varmarken*
PhD Candidate, University of California, Irvine

Christo Wilson*
Associate Professor, Khoury College of Computer Sciences, Northeastern University

(*) Denotes those who provided considerable drafting assistance

(**) Disclaimer: Titles and affiliations are provided for identification purposes only. The comments are submitted on behalf of the contributing researchers in their personal capacity, and *not* on behalf of the Cybersecurity and Privacy Institute, Northeastern University, University of California, University of Washington, the entire [ProperData](#) team, or the National Science Foundation.

Contact: gunawan.jo@northeastern.edu

Table of Contents

Signatories	1
Table of Contents	4
Comment	5
Harms to Consumers	5
Regulations	12
Rulemaking Generally	12
Data Security	13
Collection, Use, Retention, and Transfer of Consumer Data	14
Automated Systems	20
Discrimination	21
Consumer Consent	25
Notice, Transparency, and Disclosure	26
Remedies	32
Obsolescence	32

Comment

Addressed questions are bolded and organized in original order according to the ANPR categories. Unanswered questions have been removed.

Harms to Consumers

This ANPR has alluded to only a fraction of the potential consumer harms arising from lax data security or commercial surveillance practices, including those concerning physical security, economic injury, psychological harm, reputational injury, and unwanted intrusion.

- 1. Which practices do companies use to surveil consumers?**
- 2. Which measures do companies use to protect consumer data?**
- 3. Which of these measures or practices are prevalent? Are some practices more prevalent in some sectors than in others?**
- 4. How, if at all, do these commercial surveillance practices harm consumers or increase the risk of harm to consumers?**
- 5. Are there some harms that consumers may not easily discern or identify? Which are they?**
- 6. Are there some harms that consumers may not easily quantify or measure? Which are they?**

Regarding Questions 1–4 for mobile apps, IoT, and virtual reality (VR) devices:

For mobile apps, we find a long history of collecting data about who the consumer is, where they are located, and what they are doing as they interact with devices. For example, we found that apps pervasively collect data attributed to hardware identifiers as well as (resettable) uniquely identifiable advertising IDs and geolocations—allowing them to track individuals and where they are located over time¹. Further, they sometimes gather contact information about consumers and contacts in their address books. The set of data about consumers collected by apps often changes over time,² but overall we find that more data about consumers is exposed over time as apps evolve.³ In another line of work, we found no evidence of surreptitious recording of consumers via

¹ Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2016. *ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic*. In Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '16). Association for Computing Machinery, New York, NY, USA, 361–374. <https://doi.org/10.1145/2906388.2906392>; Jingjing Ren, Martina Lindorfer, Daniel J. Dubois, Ashwin Rao, David Choffnes and Narseo Vallina-Rodriguez. *Bug Fixes, Improvements, ... and Privacy Leaks: A Longitudinal Study of PII Leaks Across Android App Versions*, Network and Distributed Systems Security (NDSS) Symposium 2018 18-21 February 2018, San Diego, CA, USA, <http://dx.doi.org/10.14722/ndss.2018.23143>

² Ren, *Bug Fixes*.

³ Christophe Leung, Jingjing Ren, David Choffnes, and Christo Wilson. 2016. *Should You Use the App for That? Comparing the Privacy Implications of App- and Web-based Online Services*. In Proceedings of the 2016 Internet Measurement Conference (IMC '16). Association for Computing Machinery, New York, NY, USA, 365–372. <https://doi.org/10.1145/2987443.2987456>

mobile device microphones, but instead found cases where videos of user interactions with mobile apps were sent to a third party.⁴

In terms of protections, we find that most apps encrypt their network connections to protect against eavesdroppers (though we found several cases where credentials were exposed in plaintext, which we responsibly disclosed).⁵

For Internet-of-things (IoT) devices such as smart TVs and smart speakers, we find that commercial surveillance is harder to quantify due to extensive use of end-to-end encryption, which protects consumer data from eavesdroppers but frustrates attempts to audit what surveillance may take place via the data in those connections. We nonetheless find numerous cases where IoT devices contact third parties (some of which are known trackers),^{6,7} and further we find that most of them serve no observable purpose for the main functionality of the devices. In fact, when we block such traffic, the devices continue to operate normally.⁸

For smart TVs in particular, we find that first, third, and platform parties collect a multitude of different data types that can be used to identify the user.⁹ Notably, we find that some destinations collect both dynamic identifiers (advertising ID) and static identifiers (serial number) jointly. This practice effectively eliminates the user's ability to opt out of targeted advertising by resetting their dynamic IDs, as the tracker can simply link the old and the new dynamic IDs by joining on the static ID. When inspecting the domains that Roku and Fire TV apps contact, we find significant differences in apps' behaviors: 60% of apps (out of 1K total per platform) contact a handful or fewer ad/tracking domains, but 10% of Roku and Fire TV apps contact 20+ and 10+ ad/tracking domains, respectively. Due to the closed nature of many smart TV platforms, the privacy enhancing technologies available to consumers are in general limited to in-network tools such as Pi-Hole that block ads and tracking at domain-level granularity. Unfortunately, we observe that such coarse-grained approaches commonly cause app breakage, miss some ads, or miss blocking requests that expose PII.

⁴ Elleen Pan, Jingjing Ren, Martina Lindorfer, Christo Wilson, and David Choffnes. 2018. *Panoptispy: Characterizing Audio and Video Exfiltration from Android Applications*. In the Proceedings on Privacy Enhancing Technologies (PETS'18). <https://petsymposium.org/popets/2018/popets-2018-0030.php>.

⁵ Ren, *ReCon*, and Leung, *Bug Fixes*.

⁶ Jingjing Ren, Daniel J. Dubois, David Choffnes, Anna Maria Mandalari, Roman Kolcun, and Hamed Haddadi. 2019. *Information Exposure From Consumer IoT Devices: A Multidimensional, Network-Informed Measurement Approach*. In Proceedings of the Internet Measurement Conference (IMC '19). Association for Computing Machinery, New York, NY, USA, 267–279. <https://doi.org/10.1145/3355369.3355577>

⁷ M. Hammad Mazhar and Zubair Shafiq. *Characterizing Smart Home IoT Traffic in the Wild*. ACM/IEEE Conference on Internet of Things Design and Implementation (IoTDI), 2020.

⁸ Anna Maria Mandalari, Daniel J. Dubois, Roman Kolcun, Muhammad Talha Paracha, Hamed Haddadi, and David Choffnes. 2021. *Blocking Without Breaking: Identification and Mitigation of Non-Essential IoT Traffic*. In Proceedings on Privacy Enhancing Technologies (PETS'21). <https://www.petsymposium.org/2021/files/papers/issue4/popets-2021-0075.pdf>

⁹ Janus Varmarken, Hieu Le, Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. 2020. *The TV is Smart and Full of Trackers: Measuring Smart TV Advertising and Tracking*. In Proceedings on Privacy Enhancing Technologies (PETS'20).

<https://web.cs.ucdavis.edu/~zubair/files/smarttv-tracking-pets2020.pdf>

When focusing on smart speakers, we find that the always-on microphones generate misactivations due to incorrectly interpreting a wake word,¹⁰ leading to audio recording that was not intended or expected by the user. While most of these are just a few seconds, they sometimes can last tens of seconds, leading to an invasion of privacy. We further find that information about consumers can be exposed to third-party software running on voice assistants. Specifically, we found that Amazon and Alexa Skills (apps) obtain information about consumers when a skill is installed and used¹¹ and that this information is likely shared with third-parties. We also found that this information is processed and used to target personalized ads to users (contrary to previous statements by the platform operator), and that most skills have either missing or vague privacy policies.

For VR devices, such as the Oculus VR, we find that apps send personal identifiers (e.g., email, device ID, etc.), metadata that can be used to fingerprint the user (e.g., system version, app name, etc.), and even platform-specific data such as VR movement and play area to first, third, and platform parties.¹² When inspecting domains that are contacted, we observe that the Oculus VR's advertising and tracking ecosystem is still in its infancy. It is mostly dominated by platform party (i.e., Oculus and Facebook); with several tracking domains from Google, Facebook, and Unity; and without any ad-related domains since there were no on-device ads. This contrasts with the large and diverse ecosystems of other more mature platforms, such as mobile and smart TVs.¹³ Protection is also greatly limited for VR, as Pi-Hole do not have curated blocklists for VR. In addition, the Oculus OS is not open-sourced and does not provide any means for auditing, making it extremely challenging to audit the data collection practices for its apps.

Regarding Questions 1,3, 4-6 for web applications:

As third-party cookies are being phased out, websites are resorting to new techniques for tracking. Our recent research shows that many ad-tech organizations are now using first-party cookies for tracking when third-party cookies are blocked¹⁴. Specifically, we found that more than 90% of the top-10K websites employ first-party tracking cookies. Our research also shows that websites are increasingly employing highly invasive browser fingerprinting techniques for cross-site tracking. Our research has shown that browser fingerprinting is now used on over a quarter of the top-10K websites by known cross-site tracking companies¹⁵. It is much harder (almost impossible, perhaps)

¹⁰ Daniel J. Dubois, Roman Kolcun, Anna Maria Mandalari, Muhammad Talha Paracha, David Choffnes, and Hamed Haddadi. 2020. *When Speakers Are All Ears: Characterizing Misactivations of IoT Smart Speakers*. In Proceedings on Privacy Enhancing Technologies (PETS'20). <https://petsymposium.org/popets/2020/popets-2020-0072.pdf>

¹¹ Iqbal et. al. *Your Echos are Heard*.

¹² Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. 2022. *OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR*. In Proceedings of 31st USENIX Security Symposium (USENIX Security'22). https://www.usenix.org/system/files/sec22summer_trimananda.pdf

¹³ Varmarken et al. *The TV is Smart and Full of Trackers*.

¹⁴ Shaor Munir, Sandra Siby, Umar Iqbal, Steven Englehardt, Zubair Shafiq, and Carmela Troncoso. 2022. *CookieGraph: Measuring and Countering First-Party Tracking Cookies*. Preprint in arXiv. <https://doi.org/10.48550/arXiv.2208.12370>

¹⁵ Umar Iqbal, Steven Englehardt and Zubair Shafiq. 2021. *Fingerprinting the Fingerprinters: Learning to Detect Browser Fingerprinting Behaviors*. In Proceedings of IEEE Symposium on Security and Privacy (SP'21). 1143-1161, <https://web.cs.ucdavis.edu/~zubair/files/fpinspector-sp2021.pdf>

for users to detect and protect themselves against these tracking techniques. Specifically, in contrast to third-party cookies that users can observe and delete from their browsers, users cannot easily block first-party cookies and the browser APIs used to generate fingerprints without compromising the functionality of the website. Similarly, we find that companies employ specialized techniques, such as randomizing URL components (subdomains, paths), obfuscating JavaScript, and obfuscating the HTML structure of the ad itself, to circumvent privacy-enhancing technologies like adblockers.¹⁶ It is well-known that tracking happens alongside these obfuscated ad requests. We show that these circumvention approaches are aggressive, requiring human experts that maintain privacy-protective blocking tools on almost an hourly basis to combat these aggressive and dynamic circumvention techniques. Notably, some of these circumvention techniques can negatively impact website loading times. In summary, our research clearly shows that many websites are actively attempting to undermine or thwart privacy protections.

Tracking data is widely disseminated amongst participants in the targeted advertising ecosystem. Ad networks, ad exchanges, supply-side, and demand-side platforms routinely share the unique identifiers they have assigned to users with each other (a process known as cookie matching or cookie syncing), as well as information about users' browsing history.¹⁷ Simultaneously, ad exchanges (also known as Real Time Bidding [RTB] exchanges) forward impressions to hundreds of demand-side platforms so that they may bid upon them. Each bid request includes users' unique identifiers, IP address, geolocation, and browsing history, thus diffusing this information to hundreds of companies within the exchange. One study found that hundreds of advertising companies are able to observe the majority of users' web browsing history, not because they are embedded within websites by publishers, but because they are present within major ad exchanges and glean tracking data from the bidstream.¹⁸ Another study found that advertising companies receive bid requests, and personal information therein, even though they do not actually participate in the auction. Specifically, "null" or zero bids made up over 22% of all bids in one study.¹⁹

The diffusion of tracking data via ad exchanges is particularly pernicious because it is invisible to users. Users have no way of peering inside ad exchanges to know which companies are present and thus receiving their PII. This lack of transparency challenges the notion of "consent" with respect to online tracking: people cannot reasonably consent to data collection and sharing when the scope of the sharing and the parties involved are opaque.

Regarding Questions 5 and 6 for UX/dark patterns:

¹⁶ Hieu Le, Athina Markopoulou, and Zubair Shafiq. 2021. *CV-Inspector: Towards Automating Detection of Adblock Circumvention*. In Proceedings of Networks and Distributed Systems Security Symposium (NDSS'21). <https://web.cs.ucdavis.edu/~zubair/files/cvinspector-ndss2021.pdf>

¹⁷ Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. 2016. *Tracing Information Flows Between Ad Exchanges Using Retargeted Ads*. In Proceedings of the 25th USENIX Security Symposium. https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_bashir.pdf

¹⁸ Muhammad Ahmad Bashir and Christo Wilson. 2018. *Diffusion of User Tracking Data in the Online Advertising Ecosystem*. In Proceedings of Privacy Enhancing Technologies Symposium. <https://petsymposium.org/popets/2018/popets-2018-0033.php>

¹⁹ John Cook, Rishab Nithyanand, and Zubair Shafiq. Inferring Tracker-Advertiser Relationships in the Online Advertising Ecosystem using Header Bidding. Privacy Enhancing Technologies Symposium 2020.

Non-material harms in particular are challenging for consumers to easily discern or identify. For online designs and user interfaces, dark patterns present a unique set of difficulties for understanding potential harms. For example, missing privacy controls (in particular controls for exiting an online service and deleting accounts) present ‘Roach Motel’ dark patterns, which unfairly trap consumers in data relationships with companies and disempower consumers trying to manage these relationships efficiently.²⁰ Such designs may present non-material harms, or exacerbate other privacy or financial harms. User studies by other scholars explore consumer ability to identify or discern manipulation or harm and find continued challenges in articulately assessing interface harms.

Regarding Questions 5 and 6 for mobile/IoT:

The data collection that occurs is often completely invisible to the user. For example, in the Harvest documentary²¹, we found that the subject’s geolocation was exposed to third parties thousands of times over the course of the week—providing detailed insights into where she lived, worked, and spent time with her family. This chilling short film is not an outlier—as regular consumers, the flow of data from our devices to commercial surveillance systems is opaque and difficult to capture/present to consumers.

One way to assess that this surveillance is unknown to consumers is to provide them with the information we find through our research and see how they react. We did just this: we provided participants in our ReCon project with information about how their personal data was exposed to other parties, and we asked participants, “Have you changed your ways of using your smartphone and its applications based on the information provided by our system?” Of those who responded to the voluntary survey, a majority (20/26) indicated that they found the system useful and changed their habits related to privacy when using mobile devices.²²

7. How should the Commission identify and evaluate these commercial surveillance harms or potential harms?

8. Which areas or kinds of harm, if any, has the Commission failed to address through its enforcement actions?

Regarding Questions 7 and 8 for UX/dark patterns:

Other regulatory agencies internationally have assessed damages nonmaterial harms as pertaining to issuing damages or consumer redress in two main methods: one that considers infringement sufficient for consumer redress, and another that requires more stringent evaluations of injury. The Commission might pursue similar strategies for assessing potential harms; issues of *de minimis*

²⁰ Johanna Gunawan, Amogh Pradeep, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A *Comparative Study of Dark Patterns Across Web and Mobile Modalities*. Proc. ACM Hum.-Comput. Interact. 5, CSCW2, Article 377 (October 2021), 29 pages. <https://doi.org/10.1145/3479521>

²¹ *Harvest*. [Documentary]. <https://vimeo.com/189449163>. See also <https://www.indevu.com/harvest>.

²² Ren, *ReCon*.

scale should not be deprioritized when unfair and deceptive practices (particularly in user experience design, e.g. dark patterns) may trigger such non-material harms.²³

The Commission has already addressed dark patterns and similar designs in a few cases²⁴ and is reasonably effective in regulating against deception in these designs. For example, the Commission's recent complaint against CreditKarma²⁵ and subsequent consent order address deception and request CreditKarma to cease making misrepresented claims. State privacy laws like the CCPA and CPRA address deceptive designs for consent interactions in particular, but the Commission's future rules should evaluate such designs for unfairness as well, under its capacity to do so. Dark patterns are not exclusively deceptive and some designs may be better regulated under unfairness claims.

Regarding Questions 7 and 8 for mobile/IoT:

Researchers have developed many tools to automatically analyze mobile apps that collect personal data, potentially in violation of privacy policies. There remain important barriers to scaling these analyses, and doing so would enable much more effective enforcement. For example, many apps have vague or permissive privacy policies that could allow just about any data collection. Instead, we need more interpretable privacy policies that can be easily mapped to specific types of data collection that can be automatically audited. Another issue is obfuscation, for example, through encrypting data or otherwise hiding the fact that commercial surveillance is occurring. There need to be transparency rules that at least provide regulators—if not independent auditors—the ability to assess whether such apps are in fact compliant with their policies and Section 5 of the FTC Act. This is particularly a problem for IoT, where researchers have much less insight into the code and data being processed by devices due to their closed nature and the difficulty with reverse engineering them.

12. Lax data security measures and harmful commercial surveillance injure different kinds of consumers (e.g., young people, workers, franchisees, small businesses, women, victims of stalking or domestic violence, racial minorities, the elderly) in different sectors (e.g., health, finance, employment) or in different segments or “stacks” of the internet economy. For example, harms arising from data security breaches in finance or healthcare may be different from those concerning discriminatory advertising on social media which may be different from those involving education technology. How, if at all, should potential new trade regulation rules address harms to different consumers across different sectors? Which commercial surveillance practices, if any, are unlawful such that new trade regulation rules

²³ Johanna Gunawan, Cristiana Santos, and Irene Kamara. 2022. *Redress for Dark Patterns Privacy Harms? A Case Study on Consent Interactions*. In *Proceedings of the 2022 Symposium on Computer Science and Law (CSLAW'22)*, November 1–2, 2022, Washington, DC, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3511265.3550448>

²⁴ The Federal Trade Commission. *FTC Report Shows Rise in Sophisticated Dark Patterns Designed to Trick and Trap Consumers*. September 15, 2022. <https://www.ftc.gov/news-events/news/press-releases/2022/09/ftc-report-shows-rise-sophisticated-dark-patterns-designed-trick-trap-consumers>

²⁵ FTC v. CreditKarma, LLC. [Complaint]. https://www.ftc.gov/system/files/ftc_gov/pdf/CK%20Complaint%209-1-22%20%28Redacted%29.pdf

should set out clear limitations or prohibitions on them? To what extent, if any, is a comprehensive regulatory approach better than a sectoral one for any given harm?

Regarding question 12 generally:

History has taught us that it is difficult to identify categories of data collections that completely protect consumers against potential harms. One can, for example, prohibit collection of GPS coordinates to protect consumers against extensive location tracking and the inferences that can be made with them (e.g., whether a consumer has been located near an abortion provider). However, a data-collecting entity can still recover someone's geolocation (and changes over time) via IP-to-geolocation translation followed by gyroscope data to conduct dead reckoning.²⁶ One cannot simply ban gyroscope data usage, since it is also important for other applications (e.g., games).

Instead, we encourage the FTC to consider rules that protect consumers against current and future harms by devising rules that focus on the usage of data and how the data use incurs harms on consumers. One vision for this is fiduciary duties, where the entity collecting data must have the best interest of the consumer in mind when using data, not other competing interests – or a duty of loyalty.²⁷ By focusing on use instead of specific data types, the rules can potentially offer broader protections for consumers that are more durable over time.

Of course, there are certain types of data that should always be protected or restricted, and already are restricted under existing laws (e.g., COPPA, HIPAA). Thus, we recommend a layered approach to rule formulation, where data collection with known harms are explicitly prohibited, and other data.

Regarding question 12 for UX/dark patterns:

New trade regulation rules should consider differences in user experiences of surveillance across 'stacks' of the internet economy, in particular examining how online services may make certain activities more difficult for users with varying access to multiple devices. Dark patterns impacting user controls that are inconsistent across versions of a service (e.g. mobile app, mobile browser, desktop browser, and other modalities like VR or voice) risk the potential of users who only own certain devices from effectively exercising autonomy.²⁸ This entails taking a platform-wide evaluation of a web service and auditing the accessibility and availability of privacy controls in every served modality, in order to preserve fairness across 'slices' of the service's user base.

²⁶ Kenneth Block and Guevara Noubir. 2018. My Magnetometer Is Telling You Where I've Been? A Mobile Device Permissionless Location Attack. In Proceedings of the 11th ACM Conference on Security & Privacy in Wireless and Mobile Networks (WiSec '18). Association for Computing Machinery, New York, NY, USA, 260–270. <https://doi.org/10.1145/3212480.3212502>; Sashank Narain, Triet D. Vo-Huu, Kenneth Block and Guevara Noubir. 2017. *The Perils of User Tracking Using Zero-Permission Mobile Apps*. In IEEE Security & Privacy. <https://ieeexplore.ieee.org/abstract/document/7891528>.

²⁷ Neil M. Richards and Woodrow Hartzog. 2020. *A Duty of Loyalty for Privacy Law*. 99 Washington University Law Review 961 (2021), Available at SSRN: <https://ssrn.com/abstract=3642217> or <http://dx.doi.org/10.2139/ssrn.3642217>

²⁸ Gunawan et. al., *A Comparative Study of Dark Patterns*.

Similarly, the emergence of voice technologies, which are deployable on standalone devices or as apps on other devices, presents new opportunities for interface unfairness and deception. Our exploratory user survey on manipulative designs in voice interactions considers how voice affordances might be used towards consumer detriment.²⁹ While sight-abled consumers may be able to access controls in visual interfaces, blind or visually-impaired consumers may be more dependent on voice-assisted technologies. As such, holistic evaluation of a service' experience offerings should take different user segments into account.

Regulations

Rulemaking Generally

30. Should the Commission pursue a Section 18 rulemaking on commercial surveillance and data security? To what extent are existing legal authorities and extralegal measures, including self-regulation, sufficient? To what extent, if at all, are self-regulatory principles effective?

Regarding Question 30:

The Commissions should pursue new rulemaking on commercial surveillance and data security.

Existing laws in the United States meant to reign in commercial surveillance are not strong enough, nor are they uniform enough. For example, recent work that has examined compliance with the California Consumer Privacy Act (CCPA) found that of the 10,000 websites that appear to meet the CCPA's applicability criteria, only 18% were meeting the bare-minimum requirements of the law.³⁰ Given that other states are now adopting similarly structured online privacy laws, it is reasonable to assume that similar non-compliance problems will affect these other laws.

Industry self-regulation is also not the solution. Industry has consistently sabotaged efforts to improve online privacy that were unfavorable to its interests (e.g., the Do Not Track [DNT] standard), while promulgating their own "solutions" that are cumbersome, counterproductive, or fail to be meaningfully enforced. Examples include the following:

- The Network Advertising Initiative's browser-based opt-out tool is not honored by all advertising industry participants³¹ and—because it is cookie-based—necessitates that users

²⁹ Kentrell Owens, Johanna Gunawan, David Choffnes, Pardis Emami-Naeini, Tadayoshi Kohno, and Franziska Roesner. 2022. Exploring Deceptive Design Patterns in Voice Interfaces. In Proceedings of the 2022 European Symposium on Usable Security (EuroUSEC '22). Association for Computing Machinery, New York, NY, USA, 64–78. <https://doi.org/10.1145/3549015.3554213>

³⁰ Maggie Van Nortwick and Christo Wilson. 2022. *Setting the Bar Low: Are Websites Complying With the Minimum Requirements of the CCPA?* In Proceedings on Privacy Enhancing Technologies Symposium (PETS'22). <https://doi.org/10.2478/popets-2022-0030>.

³¹ Zengrui Liu, Umar Iqbal, and Nitesh Saxena. *Opted Out, Yet Tracked: Are Regulations Enough to Protect Your Privacy?* Preprint on arXiv. <https://arxiv.org/pdf/2202.00885.pdf>

opt-out on every device, in every browser they use, and opt-out again if they clear their cookies.

- Android’s advertising opt-out does not actually prevent apps from reading the devices’ advertising ID, sharing it with first- and third-parties, or using the ID to target ads (Google is finally changing this behavior in 2023).
- Google’s efforts to deprecate third-party cookies from Chrome have been pushed back several times, and their proposed replacement technologies in the “Privacy Sandbox” create new privacy problems and opportunities for tracking.³²
- The Interactive Advertising Bureau (IAB) defines a transparency standard for online advertising exchanges called *sellers.json* that requires ad exchanges to reveal the publishers whose impression inventory they are authorized to sell. Despite being party to the standard, the company with the largest exchange—Google—has consistently failed to comply. As of this writing, Google’s *sellers.json* still obfuscates over 90% of their publisher partners.
- Further, even when available and not obfuscated, *sellers.json* of most ad exchanges is riddled with mistakes and misrepresentations (e.g., invalid domain names, duplicate seller identifiers), making them useless for performing even basic validation.³³ In sum, compliance with self-regulation efforts (e.g., IAB’s *ads.txt* and *sellers.json* standards) is lacking.

Data Security

31. Should the Commission commence a Section 18 rulemaking on data security? The Commission specifically seeks comment on how potential new trade regulation rules could require or help incentivize reasonable data security.

32. Should, for example, new rules require businesses to implement administrative, technical, and physical data security measures, including encryption techniques, to protect against risks to the security, confidentiality, or integrity of covered data? If so, which measures? How granular should such measures be? Is there evidence of any impediments to implementing such measures?

33. Should new rules codify the prohibition on deceptive claims about consumer data security, accordingly authorizing the Commission to seek civil penalties for first-time violations?

³² Alex Berke and Dan Calacci. 2022. *Privacy Limitations of Interest-based Advertising on The Web: A Post-mortem Empirical Analysis of Google's FLoC*. In Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22). Association for Computing Machinery, New York, NY, USA, 337–349. <https://doi.org/10.1145/3548606.3560626>; Bennet Cyphers. 2019. “Don’t Play in Google’s Privacy Sandbox.” *Electronic Frontier Foundation*. <https://www.eff.org/deeplinks/2019/08/dont-play-googles-privacy-sandbox-1>; Bennett Cyphers. 2021. “Google’s FLoC is a Terrible Idea.” *Electronic Frontier Foundation*. <https://www.eff.org/deeplinks/2021/03/googles-floc-terrible-idea>; Bennet Cyphers, 2021. “Google is Testing its Controversial New Ad Targeting Tech in Millions of Browsers. Here’s What We Know.” *Electronic Frontier Foundation*. <https://www.eff.org/deeplinks/2021/03/google-testing-its-controversial-new-ad-targeting-tech-millions-browsers-heres>

³³ Yash Vekaria, Rishab Nithyanand, Zubair Shafiq. The Inventory is Dark and Full of Misinformation: Understanding the Abuse of Ad Inventory Pooling in the Ad-Tech Supply Chain, arXiv:2210.06654, 2022.

Regarding Questions 31-33:

The undersigned generally agree that the Commission should in fact commence a Section 18 rulemaking on commercial surveillance and data security to provide the opportunity for privacy and security scholars to contribute knowledge to data security rules. Requirements in the rules may follow from the scholarly findings cited and detailed in the papers cited and appended within this Comment in response to related questions posed by the Commission. The works cited provide evidence of frequent data security failures in industry, and provide recommendations for various measures to differing granularity. Thus this work broadly supports new rules and requirements regarding data security writ large.

36. To what extent, if at all, should the Commission require firms to certify that their data practices meet clear security standards? If so, who should set those standards, the FTC or a third-party entity?

Regarding Question 36:

While there are many examples of best practices and rules around data security (e.g., FIPS), data security standards evolve over time as new threats and countermeasures evolve. We recommend that the FTC establish a body of independent data security experts to set the standards to be enforced by the FTC, and leverage a combination of independent auditors and internal audits to enforce compliance. A diversity of auditing providers can address potential inconsistencies and adversarial behavior that might otherwise occur. Some audits of data security practices can be automated in a straightforward manner, e.g., ensuring correct and secure encryption for data in flight.³⁴

Collection, Use, Retention, and Transfer of Consumer Data

38. How do companies collect consumers' biometric information? What kinds of biometric information do companies collect? For what purposes do they collect and use it? Are consumers typically aware of that collection and use? What are the benefits and harms of these practices?

Regarding Question 37 for VR:

In the case of VR, we have looked into the Oculus VR ecosystem as currently, arguably, the most widely used VR platform. The apps, platform, and third parties in Oculus VR collect (including but

³⁴ Amogh Pradeep, Muhammad Talha Paracha, Protick Bhowmick, Ali Davanian, Abbas Razaghpanah, Taejoong Chung, Martina Lindorfer, Narseo Vallina-Rodriguez, Dave Levin, and David Choffnes. 2022. *A comparative analysis of certificate pinning in Android & iOS*. In Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22). Association for Computing Machinery, New York, NY, USA, 605–618. <https://doi.org/10.1145/3517745.3561439>; Muhammad Talha Paracha, Daniel J. Dubois, Narseo Vallina-Rodriguez, and David Choffnes. 2021. *IoTLS: understanding TLS usage in consumer IoT devices*. In Proceedings of the 21st ACM Internet Measurement Conference (IMC '21). Association for Computing Machinery, New York, NY, USA, 165–178. <https://doi.org/10.1145/3487552.3487830>

not limited to) play area, movement, field of view, and pupillary as reported in the OVRseen work.³⁵ They are collected by the sensors on the device and these data types were found in the (decrypted) network traffic. Furthermore, they also collect voice, and, in the near future, facial expressions.³⁶ The purposes of data collection, as described in the OVRseen paper, include core (functionality, security, etc.) and non-core (advertising, analytics, etc.) purposes. Depending on how the raw/processed data is collected and processed, there may be privacy harms to users, e.g., the Meta Quest Pro headset announcement claims that the facial expression data are going to be kept locally on the headset and only the processed data will be sent to the servers, but it is unclear whether the processed data can still be de-anonymized or not.

38. Should the Commission consider limiting commercial surveillance practices that use or facilitate the use of facial recognition, fingerprinting, or other biometric technologies? If so, how?

Much of the cited scholarship in this Comment supports limitations on commercial surveillance practices.

Regarding Question 38 for VR:

Given the findings discussed regarding virtual reality in response to Question 37 of this ANPR, it is best to limit the commercial surveillance practices w.r.t. facial recognition, fingerprinting, or similar biometric technologies; at least, in the sense that one can audit that although the data are collected they are only meant for the context of core purposes, e.g., functionality of the app.³⁷ Further, local processing should be enforced and audited, instead of cloud processing. Collection for other non-core purposes, e.g., advertising, should definitely be limited/restricted/regulated more carefully. Furthermore, it is also important for the companies to provide means for auditing these practices, namely to increase transparency beyond just stating the practices in their privacy policies.

Regarding Question 37 and 38 for Smart Speakers:

Our research has shown that smart speakers collect and use voice data (either directly or indirectly) to infer user interests and use it to serve behaviorally targeted ads³⁸. This is despite them making deceptive public promises that they do not use this information to target ads³⁹. Furthermore, smart speaker vendors have patented several privacy-infringing practices (e.g.,

³⁵Trimananda et al., "OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR." *31st USENIX security symposium (USENIX security 22)*. 2022.

³⁶ <https://www.oculus.com/blog/meta-quest-pro-privacy/>

³⁷ Trimananda et. al., *OVRSeen*.

³⁸ Umar Iqbal, Pouneh Nikkha Bahrami, Rahmadi Trimananda, Hao Cui, Alexander Gamero-Garrido, Daniel Dubois, David Choffnes, Athina Markopoulou, Franziska Roesner, and Zubair Shafiq. 2022. *Your Echos are Heard: Tracking, Profiling, and Ad Targeting in the Amazon Smart Speaker Ecosystem*. Preprint in arXiv. <https://doi.org/10.48550/arXiv.2204.10920>

³⁹ Sapna Maheshwari. 2018. "Hey, Alexa, What Can You Hear? And What Will You Do With It?" *New York Times*. <https://www.nytimes.com/2018/03/31/business/media/amazon-google-privacy-digital-assistants.html>

inference of physical and emotional characteristics of users' voices) to monetize voice input⁴⁰. Since voice data is biometric information about the user, and it is nearly impossible to change, the Commission should seriously consider limiting its usage for non-essential purposes (e.g., for advertising) that go beyond functional purposes.

39. To what extent, if at all, should the Commission limit companies that provide any specifically enumerated services (e.g., finance, healthcare, search, or social media) from owning or operating a business that engages in any specific commercial surveillance practices like personalized or targeted advertising? If so, how? What would the relative costs and benefits of such a rule be, given that consumers generally pay zero dollars for services that are financed through advertising?

Regarding Question 39:

During the peak of the COVID-19 pandemic, many healthcare-related services sprang to action towards the goal of containing the viral spread; this provides a recent and poignant illustration of the trust gaps that remain between consumers and technologies, and especially between consumers and sensitive services like healthcare. Customers unable to easily trust the types of technologies used in COVID responses (spanning facial recognition and temperature scanners, contact tracing apps, and others) due to the failures in secure data practices, honest and non-exploitative interfaces, and exploitation of PII for personalization, targeting, and increased engagement with the service.⁴¹ Trust is necessary across all online services, but especially so in the finance and healthcare sectors, or other critical services. Eroding trust within such sectors may leave consumers vulnerable to other harmful actors and without appealing or viable alternatives that will protect their sector-based interests (e.g. financial interests or well-being) as well as their data. Thus we recommend the Commission limits such services from engaging in commercial surveillance practices, for many of the other reasons articulated by other work cited in this Comment but additionally to maintain consumer trust in services available to them.

This concept is additionally noted in a 'duty of loyalty' to data subjects, which suggests that companies should behave in a manner that minimizes digital opportunism while maximizing the data subject's best interests for their purposes in using the service.⁴²

41. To what alternative advertising practices, if any, would companies turn in the event new rules somehow limit first- or third-party targeting?

Regarding Question 41:

⁴⁰ Huafeng Jin and Shuo Wang. *Voice-based determination of physical and emotional characteristics of users*, US Patent 10,096,319. <https://patents.google.com/patent/US10096319B1/en>

⁴¹ Johanna Gunawan, David Choffnes, Woodrow Hartzog, and Christo Wilson, 2021. *The COVID-19 Pandemic and the Technology Trust Gap*. 51 Seton Hall Law Review 1505 (2021), Northeastern University School of Law Research Paper No. 421, Available at SSRN: <https://ssrn.com/abstract=3874152>

⁴² Richards and Hartzog, *A Duty of Loyalty*.

One possibility is that companies will turn towards contextual advertising, which was long the cornerstone of the advertising industry before the advent of ubiquitous tracking made more invasive forms of targeting possible.

Another possibility is that companies will rely more heavily on custom audience-based advertising. By this, we refer to products where an advertiser uploads a list of PII to an ad network, and the ad network matches the PII to specific online users who may then be targeted with ads. It is our position that custom audience-based products are harmful and should be heavily regulated: for example, these products may be used to facilitate ad campaigns that discriminate against legally protected classes of people.⁴³ It is unclear whether the Commission views custom audience-like products as “first- or third-party targeting”, and consequently whether they would fall within the scope of proposed rulemaking.

Yet another possibility is that platform developers—by which we mean companies that develop operating systems and browsers—will embed advertiser-friendly tracking and targeting technologies into client-side software. A prime example of this are the Federated Learning of Cohorts (FLoC) and “Topics” APIs that Google is experimenting with in Chrome: both aggregate data about users’ behavior within Chrome itself and then make some of this information available to advertisers. While Google has positioned these technologies as being part of their “Privacy Sandbox” effort, these are not Privacy Enhancing Technologies (PETs). Further, it is unclear whether the Commission views ad targeting based on client-side data as falling within the scope of “first- or third-party targeting”.

42. How cost-effective is contextual advertising as compared to targeted advertising?

Regarding Question 42:

A 2019 study of the “interest profiles” that large ad networks (e.g., Google and Facebook) had inferred about study participants found that less than 30% of these interests were relevant to participants.⁴⁴ Furthermore, study participants overwhelmingly reported that when they were served ads targeted to incorrectly inferred interests, they found these ads to be irrelevant.

The headline result of this study is that the data used to target many forms of online ads are low-quality and inaccurate. Money spent targeting ads to people based on inaccurate data is likely to be wasted. An advertiser might assert that incorrectly targeted ads may still function to build awareness of their brand, but this argument fails to take costs into account: targeted ads tend to be more expensive than contextual ads, and thus they are an inefficient mechanism for deploying brand-awareness campaigns.

⁴³ Piotr Sapiezynski, Avijit Ghosh, Levi Kaplan, Aaron Rieke, and Alan Mislove. 2022. *Algorithms that "Don't See Color": Measuring Biases in Lookalike and Special Ad Audiences*. In Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES '22). Association for Computing Machinery, New York, NY, USA, 609–616. <https://doi.org/10.1145/3514094.3534135>.

⁴⁴ Muhammad Ahmad Bashir, Umar Farooq, Maryam Shahid, Muhammad Fareed Zaffar, and Christo Wilson. 2019. *Quantity vs. Quality: Evaluating User Interest Profiles Using Ad Preference Managers*. In Proceedings of Networks and Distributed Systems Security Symposium (NDSS'19). https://www.ndss-symposium.org/wp-content/uploads/2019/02/ndss2019_04B-5_Bashir_paper.pdf

To the best of our knowledge, there is no credible, independently conducted research demonstrating that targeted ads are more cost effective than contextually targeted ads.

43. To what extent, if at all, should new trade regulation rules impose limitations on companies' collection, use, and retention of consumer data? Should they, for example, institute data minimization requirements or purpose limitations, i.e., limit companies from collecting, retaining, using, or transferring consumer data beyond a certain predefined point? Or, similarly, should they require companies to collect, retain, use, or transfer consumer data only to the extent necessary to deliver the specific service that a given individual consumer explicitly seeks or those that are compatible with that specific service? If so, how? How should it determine or define which uses are compatible? How, moreover, could the Commission discern which data are relevant to achieving certain purposes and no more?

Regarding Question 43:

Data minimization and purpose limitations are good examples of principles that can impose limitations on collection, use, and retention of consumer data and thus reduce the harms that would otherwise occur.

One example of how to test for data minimization and purpose limitation is to block the data from being collected and observe whether the service continues to work as expected without the data. We implemented this approach in our Blocking without Breaking paper⁴⁵, demonstrating that evaluating purpose limitation is in fact feasible through technical means. There remains more work to be done to explore the set of tests that can be done to provide similar assessments for other environments.

45. Pursuant to a purpose limitation rule, how, if at all, should the Commission discern whether data that consumers give for one purpose has been only used for that specified purpose? To what extent, moreover, should the Commission permit use of consumer data that is compatible with, but distinct from, the purpose for which consumers explicitly give their data?

Regarding Question 45:

In the Internet of Things (IoT) ecosystem, there is a discrepancy between users' expectations and how companies use their data as demonstrated by research⁴⁶. The Commission should consider

⁴⁵ Mandalari et. al., *Blocking without Breaking*.

⁴⁶ Nathan Malkin*, Julia Bernd, Maritza Johnson, and Serge Egelman. 2018. *What Can't Data Be Used For? Privacy Expectations about Smart TVs in the U.S.* In Proceedings of European Workshop on Usable Security (EuroUSEC'18). <https://dx.doi.org/10.14722/eurosec.2018.23016>

identifying users' expectations, using existing methods such as contextual integrity⁴⁷ to find the right balance between user expectations and data collection and sharing practices by companies. Furthermore, IoT devices can now act as software platforms, offering the ability to third party developers to develop applications on their devices. IoT platform privacy policies allow use of users' data for certain purposes such as fraud detection and security, while prohibiting others such as targeted advertising when users opt out of interest-based advertising. However, often third-party app developers are delegated and trusted with the task of complying with these privacy policies and it is not clear what level of enforcement, if any, currently exists in order to distinguish between different use cases of users' data⁴⁸. The Commission should incentivize IoT platforms to develop methods for testing whether applications follow data privacy guidelines on their devices and identify applications that are not complying with platform policies.

52. To what extent, if at all, do firms that now, by default, enable consumers to block other firms' use of cookies and other persistent identifiers impede competition? To what extent do such measures protect consumer privacy, if at all? Should new trade regulation rules forbid the practice by, for example, requiring a form of interoperability or access to consumer data? Or should they permit or incentivize companies to limit other firms' access to their consumers' data? How would such rules interact with general concerns and potential remedies discussed elsewhere in this ANPR?

Regarding Question 52:

While restrictions around third-party cookies are expected to enhance and protect consumer privacy, research⁴⁹ has shown that blocking third-party cookies has motivated the trackers and advertisers to look towards first-party cookies for identifying and tracking users. Even with complete blockage of third-party cookies at least a full year away (Google Chrome aims to block third-party cookies by the end of 2024), a majority of trackers and advertisers have been shown to rely on first-party cookies for both inter and intra-site tracking. Trackers and Advertisers like Google, Microsoft, and Criteo were found using first-party cookies on more than two-thirds of the sites tested. These cookies were found to be not only in use by these trackers themselves, but they are also being shared in abundance with other trackers and advertisers. With this network of information sharing which exists in spite of restrictions around cookies, it's unclear if blocking third-party cookies is a silver bullet for consumer privacy. Countermeasures against use of first-party cookies and other potential replacements for third-party cookies are needed to ensure consumer privacy remains intact.

In the context of Internet of Things (IoT), IoT devices can now act as software platforms, offering the ability to third party developers to develop applications on their devices. However, besides an

⁴⁷ Noura Abdi, Xiao Zhan, Kopo M. Ramokapane, and Jose Such. 2021. Privacy Norms for Smart Home Personal Assistants. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 558, 1–14. <https://doi.org/10.1145/3411764.3445122>

⁴⁸ Hooman Mohajeri Moghaddam. 2022. *Tracking and Behavioral Targeting on Connected TV Platforms*. Dissertation, Princeton University. <https://www.proquest.com/docview/2721346426/>

⁴⁹ Munir et al., *CookieGraph*.

agreement at the time of the registration that asks the third-party developers to acknowledge and abide by the platform's privacy policies, third-party applications are often trusted with the enforcement of the platform privacy policies⁵⁰⁵¹. While this shields the platforms from legal liability, it does not address the underlying problem. Moreover, IoT platform developers themselves engage in data collection practices that are independent of other firms and do not require data sharing methods and these practices are often not well understood, because these platforms often use proprietary software and hardware. In fact, we have already seen instances of failure to comply with privacy policies⁵². The Commission should incentivize IoT platforms to disclose more about their data collection practices and invest more on enforcement of their privacy policies by application developers.

Automated Systems

56. To what extent, if at all, should new rules require companies to take specific steps to prevent algorithmic errors? If so, which steps? To what extent, if at all, should the Commission require firms to evaluate and certify that their reliance on automated decision-making meets clear standards concerning accuracy, validity, reliability, or error? If so, how? Who should set those standards, the FTC or a third-party entity? Or should new rules require businesses to evaluate and certify that the accuracy, validity, or reliability of their commercial surveillance practices are in accordance with their own published business policies?

Regarding Question 56:

We propose an approach similar to the Digital Services Act in the EU, where companies would be legally mandated to share data with academic researchers for the specific purpose of auditing their automated systems. Under such a framework, FTC would act as a facilitator between companies and researchers to approve audit studies, settle disputes and more. But it is ultimately researchers that would identify questions to be studied, get access to the data, conduct audits and certify compliance with legally and ethically expected behaviors.

We have also proposed a framework for how such auditing can be performed for a specific type of platform and specific types of harms -- newsfeed-driven advertisement platforms and their potential for discriminatory advertisement delivery. In a paper to appear at CSCW '23, we outline

⁵⁰ Long Cheng, Christin Wilson, Song Liao, Jeffrey Young, Daniel Dong, and Hongxin Hu. 2020. *Dangerous Skills Got Certified: Measuring the Trustworthiness of Skill Certification in Voice Personal Assistant Platforms*. In Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS '20). Association for Computing Machinery, New York, NY, USA, 1699–1716. <https://doi.org/10.1145/3372297.3423339>

⁵¹ Moghaddam. *Tracking and Behavioral Targeting on Connected TV Platforms*.

⁵² Federal Trade Commission. 2017. "VIZIO to Pay \$2.2 Million to FTC, State of New Jersey to Settle Charges It Collected Viewing Histories on 11 Million Smart Televisions without Users' Consent." Press Releases of the Federal Trade Commission.

<https://www.ftc.gov/news-events/news/press-releases/2017/02/vizio-pay-22-million-ftc-state-new-jersey-settle-charges-it-collected-viewing-histories-11-million>

what specific piece of data would enable effective auditing, and how this data can be shared to researchers without compromising user privacy and while protecting companies' business interests.⁵³

Discrimination

65. How prevalent is algorithmic discrimination based on protected categories such as race, sex, and age? Is such discrimination more pronounced in some sectors than others? If so, which ones?

Regarding Question 65:

Our research has found a number of instances of algorithmic discrimination based on protected classes including race, sex, and age. We have focused primarily on the advertising industry, due to the ability to use advertising systems to study the underlying online platforms.

On Facebook, our group has studied algorithmic discrimination in the context of *ad delivery*.⁵⁴ In brief, ad delivery is the process by which the platform decides which advertiser wins an ad slot—an opportunity to show an ad to a user. Today, platforms often will subsidize ads they believe are relevant to users; these estimates of relevance are often determined using machine learning models. We studied Facebook's ad delivery algorithm by becoming a Facebook advertiser and running carefully controlled ads.

We found that biased ad delivery can occur due to the content of the ad itself (i.e., the ad headline, text, and image, collectively called the ad creative) or the content of the webpage that the ad leads to. For example, ads targeting the same audience but that include a creative that would stereotypically be of the most interest to men (e.g., bodybuilding) can deliver to over 80% men, and those that include a creative that would stereotypically be of the most interest to women (e.g., cosmetics) can deliver to over 90% women. Similarly, ads referring to cultural content stereotypically of most interest to Black users (e.g., hip-hop) can deliver to over 85% Black users, and those referring to content stereotypically of interest to white users (e.g., country music) can deliver to over 80% white users, even when targeted identically by the advertiser. Thus, despite placing the same bid on the same audience, the advertiser's ad delivery can be heavily skewed by the platform, often unbeknownst to the advertiser, based on the ad creative alone.

Worse, we found that these effects persist along gender and racial lines for ads for employment opportunities, and they are present despite Facebook having a distinct ad creation flow for housing, credit, and employment ads. In the most extreme cases, our ads for jobs in the lumber industry reach an audience that is 72% white and 90% male, our ads for cashier positions in supermarkets reach an 85% female audience, and our ads for positions in taxi companies reach a 75% Black

⁵³ Basileal Imana, Aleksandra Korolova, John Heidemann. *Having your Privacy Cake and Eating it Too: Platform-supported Auditing of Social Media Algorithms for Public Interest*, arXiv:2207.08773, 2022. To appear in Proceedings of the 26th ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW) 2023. <https://www.isi.edu/~johnh/PAPERS/Imana22a.pdf>.

⁵⁴ Ali et. al., *Discrimination through Optimization*.

audience, even though the targeted audience specified by us as an advertiser is identical and balanced across gender and racial lines for all three.

Further, we found that these biased delivery effects cannot be explained simply by population-level differences between demographic groups in professional qualifications. We found that when two jobs that require the same set of skills are advertised by different companies, the ad delivery is skewed based on existing employee demographics of those companies: our ads for software engineers at Nvidia delivered significantly more to men than our ads for software engineers at Netflix, despite identical targeting parameters and expected qualifications.⁵⁵ We have also shown that Facebook’s delivery algorithm can skew the reached audience based on the demographics of the person depicted in the job⁵⁶ and that advertisers may be deliberately choosing the advertising images in response to this⁵⁷.

Unfortunately, platforms’ efforts to combat discrimination are not always successful. Following a civil rights settlement Facebook removed demographic attributes as input from an algorithm for finding users similar to those provided by the advertiser. Our measurements showed that despite this change, the new algorithm produced results often not less biased than the original algorithm, i.e. if the advertiser provided a list of mostly male users, the algorithm would identify other male users as similar; the same effect held for age, race, and political leaning.⁵⁸ Following a settlement with the Department of Justice, Facebook removed advertisers’ access to either algorithm when running ads for housing, employment, and credit.

Furthermore, our group has studied algorithmic discrimination based on race and gender in the context of online labor markets and freelance marketplaces. In the context of online labor markets, we collected search results from Indeed, Monster, and CareerBuilder based on 35 job titles in 20 U.S. cities.⁵⁹ On the positive side, we found that the ranking algorithms used by all three hiring sites did not use candidates’ inferred gender as a feature. However, we did observe significant and consistent group unfairness against feminine candidates in roughly 1/3 of the job titles examined. Specifically, the ranking algorithms reflected structural gender inequalities that were embedded in the raw data. In the context of online freelance marketplaces, we collected worker profiles from TaskRabbit and

⁵⁵ Basileal Imana, Aleksandra Korolova, and John Heidemann. 2021. *Auditing for Discrimination in Algorithms Delivering Job Ads*. In Proceedings of the Web Conference 2021 (WWW '21). Association for Computing Machinery, New York, NY, USA, 3767–3778. <https://doi.org/10.1145/3442381.3450077>

⁵⁶ Levi Kaplan, Nicole Gerzon, Alan Mislove, and Piotr Sapiezynski. 2022. *Measurement and analysis of implied identity in ad delivery optimization*. In Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22). Association for Computing Machinery, New York, NY, USA, 195–209. <https://doi.org/10.1145/3517745.3561450>

⁵⁷ Varun Nagaraj Rao and Aleksandra Korolova. 2022. *An Audit of Images Used by Job Advertisers on Facebook*. Poster in 2nd ACM conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO'22).

⁵⁸ Sapiezynski et. al., *Algorithms that “Don’t See Color.”*

⁵⁹ Le Chen, Ruijun Ma, Anikó Hannák, and Christo Wilson. 2018. *Investigating the Impact of Gender on Rank in Resume Search Engines*. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18). Association for Computing Machinery, New York, NY, USA, Paper 651, 1–14. <https://doi.org/10.1145/3173574.3174225>

Fiverr.⁶⁰ In both marketplaces, we found that perceived gender and race were significantly correlated with worker evaluations (e.g. ratings, reviews). We also found that TaskRabbit’s algorithm generated results that were significantly correlated with perceived race and gender. Although we couldn’t determine the cause of this bias, one plausible explanation was that the algorithm was designed to take customer behavior (e.g. ratings, reviews, clicks on profiles) into account.

66. How should the Commission evaluate or measure algorithmic discrimination? How does algorithmic discrimination affect consumers, directly and indirectly? To what extent, if at all, does algorithmic discrimination stifle innovation or competition?

Regarding Question 66:

Our group has had significant success developing methodologies to perform external *algorithm audits*.⁶¹ In brief, these audits allow third parties who have no privileged access to the system of concern to study the algorithms that power them and reach scientific conclusions about how decisions are made. Algorithm audits are often highly customized to the particular system under study, as the levels of access, input data, output results, and other constraints often necessitate

⁶⁰ Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. *Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr*. In Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17). Association for Computing Machinery, New York, NY, USA, 1914–1933. <https://doi.org/10.1145/2998181.2998327>

⁶¹ Ali et. al, *Discrimination through Optimization.*; Kaplan et. al., *Measurement and analysis of implied identity.*; Basileal Imana, Aleksandra Korolova, and John Heidemann. 2021. *Auditing for Discrimination in Algorithms Delivering Job Ads*. In Proceedings of the Web Conference 2021 (WWW '21). Association for Computing Machinery, New York, NY, USA, 3767–3778. <https://doi.org/10.1145/3442381.3450077>; Imana et. al., *Having your Privacy Cake and Eating It Too.*; Mark Juarez and Aleksandra Korolova. 2022. *"You Can't Fix What You Can't Measure": Privately Measuring Demographic Performance Disparities in Federated Learning*. To appear in Proceedings of Machine Learning Research (PMLR'22). <https://arxiv.org/abs/2206.12183>.

different approaches. But we have successfully conducted audits of web search engines,⁶² hiring systems,⁶³ transportation network companies,⁶⁴ and advertising platforms.⁶⁵

Our group has also proposed a new framework for adding explicit platform support for auditing advertising platforms.⁶⁶ At a high-level, the framework would allow external researchers to directly measure various forms of algorithmic discrimination without the need to come up with the highly customized approaches and data collection methods that our prior external audits have had to develop. Giving researchers auditor-specific and privileged access to conduct audits helps avoid the constraints of fully external auditing methods which rely on the limited features that platforms make available to any user or advertiser. Our proposal shows how platforms can provide such access while safeguarding privacy interests of users and business interests of platforms.

67. How should the Commission address such algorithmic discrimination? Should it consider new trade regulation rules that bar or somehow limit the deployment of any system that produces discrimination, irrespective of the data or processes on which those outcomes are based? If so, which standards should the Commission use to measure or evaluate disparate outcomes? How should the Commission analyze discrimination based on proxies for protected categories? How should the Commission analyze discrimination when more than one protected category is implicated (e.g., pregnant veteran or Black woman)?

Regarding Question 67:

In addressing discrimination, the Commission should look end-to-end and not focus on individual components of a larger system, or rely solely on limiting the kinds of inputs provided to systems. Let us provide an instructive example, in which we have shown that such approaches are doomed to failure.⁶⁷ Facebook provides a tool to advertisers to find users who “look like” their current customers (this is literally called Lookalike Audiences on Facebook). Based on a machine learning

⁶² Aniko Hannak, Piotr Sapiezynski, Arash Molavi Kakhki, Balachander Krishnamurthy, David Lazer, Alan Mislove, and Christo Wilson. 2013. *Measuring personalization of web search*. In Proceedings of the 22nd international conference on World Wide Web (WWW '13). Association for Computing Machinery, New York, NY, USA, 527–538. <https://doi.org/10.1145/2488388.2488435>; Ronald E. Robertson, Shan Jiang, Kenneth Joseph, Lisa Friedland, David Lazer, and Christo Wilson. 2018. *Auditing Partisan Audience Bias within Google Search*. Proc. ACM Hum.-Comput. Interact. 2, CSCW, Article 148 (November 2018), 22 pages. <https://doi.org/10.1145/3274417>.

⁶³ Chen et al., *Investigating the Impact of Gender on Rank in Resume Search Engines*; Hannák et al., *Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr*

⁶⁴ Le Chen, Alan Mislove, and Christo Wilson. 2015. *Peeking Beneath the Hood of Uber*. In Proceedings of the 2015 Internet Measurement Conference (IMC '15). Association for Computing Machinery, New York, NY, USA, 495–508. <https://doi.org/10.1145/2815675.2815681>; Shan Jiang, Le Chen, Alan Mislove, and Christo Wilson. 2018. *On Ridesharing Competition and Accessibility: Evidence from Uber, Lyft, and Taxi*. In Proceedings of the 2018 World Wide Web Conference (WWW '18). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 863–872. <https://doi.org/10.1145/3178876.3186134>

⁶⁵ Ali et. al, *Discrimination through Optimization.*; Kaplan et. al., *Measurement and analysis of implied identity*; Imana et. al. *Auditing for Discrimination*.

⁶⁶ Imana et. al., *Having your Privacy Cake and Eating it Too*.

⁶⁷ Sapiezynski et. al., *Algorithms that “Don’t See Color.”*

algorithm, this product takes in a list of users provided by the advertiser, and constructs a larger list of users who are similar. In 2018, Facebook settled a lawsuit with the National Fair Housing Alliance concerning how this tool could lead to discrimination. Facebook agreed to modify the functionality of Lookalike Audiences when used to target housing, credit, and employment ads. In brief, Facebook created the Special Ad Audiences tool, which works like Lookalike Audiences, except its algorithm does *not* consider users' age, gender, relationship status, religious views, school, political views, interests, or zip code when detecting common qualities.

Facebook left both tools accessible, allowing us to compare whether removing these features actually reduces the level of disparity in the resulting audiences. We demonstrated that our Special Ad audiences are skewed to almost the same degree as Lookalike audiences, with many of the results being statistically indistinguishable. For example, when using a source audience that is all women, our Lookalike audience-targeted ad delivered to 96.1% women, while Special Ad audience-targeted ad delivered to 91.2% women. We also provide evidence indicating that both Lookalike and Special Ad audiences carry—to a certain extent—the biases of the source audience in terms of race and political affiliation.

68. Should the Commission focus on harms based on protected classes? Should the Commission consider harms to other underserved groups that current law does not recognize as protected from discrimination (e.g., unhoused people or residents of rural communities)?

Regarding Question 68:

Yes, the Commission should focus on both legally protected classes as well as other portions of the population that may be vulnerable but do not enjoy the same legal protections. In the context of advertising, most online platforms are built to identify which users are most “receptive” to ads – essentially, they are built to discriminate (in a statistical, not legal, sense). In many cases, being in a vulnerable population *is* predictive of whether you are likely to engage with a particular ad or opportunity. As a result, these systems will often pick up on and exploit such divisions. We make this suggestion following cites and responses to other questions regarding discrimination in this Comment.

Consumer Consent

80. Have opt-out choices proved effective in protecting against commercial surveillance? If so, how and in what contexts?

Regarding question 80:

Websites often utilize Consent Management Platforms (CMPs), such as OneTrust⁶⁸, to solicit user consent, with the expectation that the consent will be respected and the users data will not be

⁶⁸ OneTrust:<https://www.onetrust.com>

processed, shared, and sold. Our research⁶⁹ evaluated several CMPs, and has shown that user data is unfortunately still being collected, processed, and shared even when users opt-out. Specifically, even after opting-out third-parties (embedded on websites), including advertisers, still collect user cookies, share it with other third-parties (e.g., through HTTP redirects), and use it to serve targeted advertisements. Overall, our measurements cast a doubt if consent conveyed through CMPs is effective at protecting users' online privacy.

82. How, if at all, should the Commission require companies to recognize or abide by each consumer's respective choice about opting out of commercial surveillance practices—whether it be for all commercial surveillance practices or just some? How would any such rule affect consumers, given that they do not all have the same preference for the amount or kinds of personal information that they share?

Regarding question 82:

On Consent, Opt-Out, and Dark Patterns: If services can be used across modalities but opt-out choices are not equivalently provided, then they are not effective in these cases. Similarly, options should not asymmetrically introduce friction for options the company prefers (to consumers' potential detriment). Deploying such designs for opt-out regimes risks unequal experiences for different consumer demographics, with potentially worse outcomes for those with less access to multiple modalities (in particular, we find that apps tend towards having more dark patterns than websites, which increases risk for consumers who primarily or solely use app version of services out of necessity).⁷⁰ We additionally suggest that the FTC recommend that companies provide equivalent access to relevant controls for opt-out or similar controls in settings interfaces, and that the FTC especially hold companies to this standard if they provide services across multiple modalities. Choices should be controllable from all modalities and applied equivalently; e.g. turning off cookies on a web browser site should prevent similar information from being used in an app version, and consumers should be able to trust that certain settings are universally honored beyond the local modality used.

Notice, Transparency, and Disclosure

83. What kinds of information should new trade regulation rules require companies to make available and in what form?

Regarding Q83:

Work cited in this Comment affirms our stance that the Commission should consider rules requiring companies to give notice of their commercial surveillance practices. Current forms of notice include privacy policies, press releases, and interaction with the user during the service.

⁶⁹ Liu et al., *Opted Out, Yet Tracked: Are Regulations Enough to Protect Your Privacy?*

⁷⁰ Gunawan et. al., *A Comparative Study of Dark Patterns Across Web and Mobile Modalities*.

On privacy policies:

An entity's written privacy policy is the interface between privacy laws, and system design and implementation, as well as communication with the user. While providing a privacy policy is necessary, it is currently not sufficient to provide transparency due to the following challenges:

- A privacy policy must be written so that it meets the minimum requirements in order to be compliant with privacy laws.
- A privacy policy must accurately and comprehensively describe the data collection practices implemented by the system and any third-party libraries the system uses. There is currently a disconnect between the text of privacy policies and the design and implementation of the corresponding systems. There is currently no technical or automated way to audit or enforce the privacy policies.
- A privacy policy should be (i) comprehensive in its coverage of the system functionality and law requirements, (ii) self-consistent, i.e., without contradicting statements, and (iii) present information to the user in a way that is easily understood. Today's privacy policies are typically long, and written in a vague or misleading way. This creates a challenge for users to understand.

We believe that privacy policies should continue to be required for companies, so that they disclose their data collection (or "commercial surveillance") practices in legally-binding ways. However, we also believe that these policies should be written more clearly⁷¹, in a simple and standard way that makes it possible to understand, audit and enforce them on the corresponding systems.⁷²

(1) **Precise definition:** More specifically, the exact information of what must be disclosed in privacy policies should be guided by and aligned with the disclosure requirements in privacy laws. For example, privacy laws, such as GDPR and CCPA, currently require at least the disclosure of the following information regarding data collection, use, and sharing practices:

- Categories of personal information collected, used, or shared
- Source (GDPR) or categories of sources (CCPA) of the personal information
- Purposes for the collection, use, and sharing of personal information
- Categories of third parties with whom personal information is shared

This has led to privacy policies currently being written in sections with lists of items, e.g., "We collect the following categories of personal information....", "We use your personal information for....", "We disclose your information with". Although compliant with the disclosure requirements of privacy laws, this way of writing privacy policies gives only a vague description of what exact data types are collected, by which entity, and for which exact purpose. New rules should require that privacy policies have precise statements, such as: " *we collect your email address for the purpose of providing you service X ...; we share your email address with Google for monetization purposes... etc.*"

⁷¹ Hao Cui, Rahmadi Trimananda, Athina Markopoulou, and Scott Jordan. 2022. *PoliGraph: Automated Privacy Policy Analysis using Knowledge Graphs*. Preprint on arXiv. <https://arxiv.org/abs/2210.06746>

⁷² Athina Markopoulou, Rahmadi Trimananda, and Hao Cui. 2022. *A Contextual-Integrity Based Auditing Framework for Data Collection Practices*. Position Paper in Annual Symposium on Applications of Contextual Integrity (PrivaCI'22).

<https://cpb-us-e2.wpmucdn.com/faculty.sites.uci.edu/dist/5/718/files/2022/09/ACI-basedAuditingFrameworkForDataCollectionPractices.pdf>

Furthermore, many policies are also vague about the exact data types (e.g., using generic terms such as device information, personal information, etc.) and about the entities that collect or receive the data (e.g. *third-party vendor* instead of the exact company name). New rules should require that a policy clearly defines these terms.

(2) **Clear presentation.** To facilitate users' understanding, new rules should require a standardized concise representation of privacy policy that discloses only important aspects, e.g., the type of personal information collected, the categories of third parties with whom the company shares it, and the purposes for which the company uses it. Early examples to the right direction include Apple and Google's app stores' "privacy labels", which summarize privacy notices.

(3) **Auditing.** Most importantly, it is not enough to specify a privacy policy, if there is no technical way to check the actual behavior and consistency of the corresponding system. New rules should give users ways to inspect their personal information that companies collect (e.g., see GDPR Article 20 "Right to data portability"). There should be technical ways, e.g., interfaces provided by the companies for auditing, that allow an external auditor to inspect the system's data commercial surveillance practices, and declare if they are consistent with the privacy policy and the law. We have performed one such auditing of data collection practices and their consistency with their privacy policies, specifically for Oculus VR⁷³, which can be used as a template going forward.⁷⁴

84. In which contexts are transparency or disclosure requirements effective? In which contexts are they less effective?

Regarding Question 84:

Privacy policies and similar disclosures are commonly left unread by consumers with little time to read them, which is troubling as these documents may contain important information and even avenues for exercising data rights or controls. Compliance measured only by disclosures is ineffective; compliance must be accompanied with auditable behavior within the service that corresponds to the privacy policies or other documents.

Beyond improved auditing enforcement, transparency or disclosure requirements are ineffective if users are not able to easily access controls referred to in disclosures. One such example is lack of compliance to the CCPA's DNSMPI link requirements;⁷⁵ if these are not clearly, conspicuously, and transparently provided to users, or if other controls or instructions are buried within lengthy disclosures and not provided in relevant interactions, then disclosure alone does not effectively provide consumers with options. In such cases, the consumer option itself is not disclosed appropriately – this is a separate, but related issue to a failure to disclose practices writ large.

⁷³ Trimananda et al, *OVRSeen*.

⁷⁴ Markopoulou et. al., *A CI-Based Auditing Framework*.

⁷⁵ Van Nortwick and Wilson, *Setting the Bar Low*.

Please also see response to Q83 above, and citations therein.

85. Which, if any, mechanisms should the Commission use to require or incentivize companies to be forthcoming? Which, if any, mechanisms should the Commission use to verify the sufficiency, accuracy, or authenticity of the information that companies provide?

Regarding Question 85:

Current Internet and Web technologies on personal computers and mobile devices often have well-supported ecosystems for monitoring and intercepting network traffic that enable researchers to study the data sharing practices on these platforms, such as tools for modern web browsers⁷⁶ and mobile operating systems⁷⁷. However, the emerging Internet of things (IoT) devices often run on proprietary software and hardware and do not currently provide the same level of access for analyzing such devices, as already pointed out by researchers⁷⁸. This makes the task of understanding data sharing patterns and privacy policy enforcement more difficult on these platforms, because monitoring network traffic is often only possible by device manufacturers and developers. The Commission should incentivize companies to open up their devices to researchers, while being cognizant of the concerns around intellectual property. The Commission can lower the cost of monitoring and studying these platforms by providing incentives to companies to partner with researchers and providing them with special hardware prototypes or debug software licenses. The Commission can also engage with standardization efforts that include transparency and user privacy as their goals, such as the Digital Standard by the Consumer Reports.⁷⁹

Please also see response to Q83 above, regarding auditing, and cites therein.

86. The Commission invites comment on the nature of the opacity of different forms of commercial surveillance practices. On which technological or legal mechanisms do companies rely to shield their commercial surveillance practices from public scrutiny? Intellectual property protections, including trade secrets, for example, limit the involuntary public disclosure of the assets on which companies rely to deliver products, services, content, or advertisements. How should the Commission address, if at all, these potential limitations?

Regarding Question 86:

Please see answer to 83 above and cites therein, regarding vague, misleading privacy policies that are not necessarily consistent with the actual commercial surveillance practices of the system.

⁷⁶ openWPM Web Privacy Measurement Framework. <https://github.com/openwpm/OpenWPM>

⁷⁷ [The Haystack Project](https://haystack.mobi/) . <https://haystack.mobi/>

⁷⁸ Hooman Mohajeri Moghaddam, Gunes Acar, Ben Burgess, Arunesh Mathur, Danny Yuxing Huang, Nick Feamster, Edward W. Felten, Prateek Mittal, and Arvind Narayanan. 2019. *Watching You Watch: The Tracking Ecosystem of Over-the-Top TV Streaming Devices*. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19). Association for Computing Machinery, New York, NY, USA, 131–147. <https://doi.org/10.1145/3319535.3354198>

⁷⁹ <https://thedigitalstandard.org/>

89. To what extent should trade regulation rules, if at all, require companies to explain (1) the data they use, (2) how they collect, retain, disclose, or transfer that data, (3) how they choose to implement any given automated decision-making system or process to analyze or process the data, including the consideration of alternative methods, (4) how they process or use that data to reach a decision, (5) whether they rely on a third-party vendor to make such decisions, (6) the impacts of their commercial surveillance practices, including disparities or other distributional outcomes among consumers, and (7) risk mitigation measures to address potential consumer harms?

Regarding Question 89:

One interesting finding from the OVRseen paper⁸⁰ is that a lot of companies (app developers) in VR, at least, tend to neglect referencing third-party privacy policies in their privacy policies, whereas their app actually uses a lot of third-party libraries that collect data without their and, moreover, user's knowledge.

Please see answers to Questions 83: they apply here as well.

92. To what extent should the Commission, if at all, make regular self-reporting, third-party audits or assessments, or self-administered impact assessments about commercial surveillance practices a standing obligation? How frequently, if at all, should the Commission require companies to disclose such materials publicly? If it is not a standing obligation, what should trigger the publication of such materials?

Regarding Q92:

See also our response to questions 83 and 89 above, and citations therein.

At the very minimum, the Commission should require companies to make information available to the public about what data it collects and for what purpose, similar in spirit to Apple's App Tracking Transparency. Companies should be obligated to both list data collected for first-party use, as well as data that is collected and shared with third parties. The data collection information should be required to be comprehensive: it should list specific data types, collection frequency for each data type, and it should restrict use of vague, aggregate terms (e.g., "device serial number, device MAC address, advertising ID" instead of "various unique identifiers" and "first name, last name, phone number" instead of "personal information" or "contact details"). The format in which the data collection practices are presented should be standardized by the Commission. It is (a) comprehensible for the average consumer and (b) machine readable, to enable automated analysis of data collection practices. Inspiration may be drawn from the standardized privacy notices used by financial institutions.⁸¹ As third-party data collection often stems from integration of third party code, providers of such code should be obligated to make their data collection information in their privacy notice available s.t. other developers can easily include it in products that integrate said

⁸⁰Trimananda et al., *OVRSeen*.

⁸¹ Lorrie Faith Cranor et al. *A Large-Scale Evaluation of U.S. Financial Institutions' Standardized Privacy Notices*.

code. For instance, it is found in the Oculus VR ecosystem that the negligence of including references to these third-party privacy notices by developers is prevalent.⁸²

However, self-reporting cannot be relied on exclusively as it provides no means to detect inaccurate reporting. To minimize inaccuracies resulting from negligence, companies should be required to self-audit their products (i.e., actual testing of their products) to ensure that self-reporting stays consistent with the actual behavior of their products. The legislative requirements to the frequency of such self-audits should consider to what extent auditing can be automated (e.g., as has been shown to be possible for smart TVs⁸³, voice assistants⁸⁴, and VR⁸⁵), crowd-sourced, or has to be performed manually.

To hold companies accountable and truthful, the Commission should require companies to build mechanisms into their products that enable an independent (trusted) party (e.g., researchers, non-profit consumer advocates) to audit the data collection practices. There should be an interface provided by companies for auditing of their (commercial surveillance) practices and for checking their consistency with the privacy policies and compliance with privacy laws.

93. To what extent do companies have the capacity to provide any of the above information?

Given the potential cost of such disclosure requirements, should trade regulation rules exempt certain companies due to their size or the nature of the consumer data at issue?

Regarding Question 93:

We demonstrate through a comparative study of dark pattern experiences across three modalities that companies often do have the technical capacity to serve certain consumer-forward features (e.g. bulk settings toggles, simpler navigation, provision of options, among other designs) in at least one modality of that service.⁸⁶ However, ‘capacity’ at a *platform* level – that a service is technically capable of offering certain settings at all – should result in important options being provided to consumers regardless of modality; a service that offers certain user account settings in a website should offer those same settings via app if those are useful privacy or consumer controls and there are no legitimate design constraints on offering those controls everywhere the service is accessed in a similar fashion. Disclosures are typically less complicated to serve than highly interactive features, so companies that are capable of providing disclosures in one modality while providing more technologically complicated parts of their service across several modalities should be required to provide disclosures equivalently across all modalities.

To ensure that consumer experiences do not disparately impact those who have less access to multiple devices, trade regulation rules should institute disclosure requirements across modalities with minimal exemptions; the cost of implementing such disclosures should be outweighed by an

⁸² Trimananda et al., *OVRseen*.

⁸³ Varmarken et al. *The TV is Smart and Full of Trackers*.

⁸⁴ Iqbal et al. *Your Echos are Heard*.

⁸⁵ Trimananda et. al., *OVRseen*.

⁸⁶ Gunawan et. al. *A Comparative Study of Dark Patterns*.

interest in protecting all consumers of a given service. (Note that these recommendations are particular to multimodal services, not unimodal services like app-only or website-only tools).

Remedies

94. How should the FTC's authority to implement remedies under the Act determine the form or substance of any potential new trade regulation rules on commercial surveillance? Should new rules enumerate specific forms of relief or damages that are not explicit in the FTC Act but that are within the Commission's authority? For example, should a potential new trade regulation rule on commercial surveillance explicitly identify algorithmic disgorgement, a remedy that forbids companies from profiting from unlawful practices related to their use of automated systems, as a potential remedy? Which, if any, other remedial tools should new trade regulation rules on commercial surveillance explicitly identify? Is there a limit to the Commission's authority to implement remedies by regulation?

Regarding Question 94:

With regards to issuing damages for nonmaterial harms of deceptive designs and dark patterns, including privacy loss and associated nonmaterial harms, the Commission's authority is limited by extant by a lack of effective measurements for evaluating such harms towards a satisfiable legal threshold, much like the GDPR and EU examples discussed in our case study of consent dark patterns and opportunities for redress.⁸⁷ (We do acknowledge recent successes with dark patterns enforcement actions, which do compensate users for some deceptive behaviors). New rules should enumerate forms of damages for nonmaterial harms to users that are not explicit in the FTC Act, defining these further and drawing upon prior FTC dark patterns enforcement actions to build these definitions. Similar remedial tools must look beyond consent to protect consumers against commercial surveillance, and should explicitly discuss the role of deceptive designs in over-encouraging consumer data sharing generally, then prohibit the use of such designs without effective disclosure in high-risk consumer interactions.

Obsolescence

95. The Commission is alert to the potential obsolescence of any rulemaking. As important as targeted advertising is to today's internet economy, for example, it is possible that its role may wane. Companies and other stakeholders are exploring new business models. Such changes would have notable collateral consequences for companies that have come to rely on the third-party advertising model, including and especially news publishing. These developments in online advertising marketplace are just one example. How should the Commission account for changes in business models in advertising as well as other commercial surveillance practices?

Regarding Question 95:

⁸⁷ Gunawan et. al. *Redress for Dark Patterns*.

The aforementioned 'duty of loyalty' is one way to account for changes in business models and related practices because this duty is resilient to temporary trends or future innovation.⁸⁸ Loyalty as a principle for improving consumer privacy naturally requires companies to adapt to changing technologies and business models while keeping their duty to consumers constant. Rules that promote loyalty and trust, particularly rules that speak to the nature of the relationship between company and consumer, provide flexibility and some preventative measures against future abuses of advertising and commercial surveillance.

⁸⁸ Richards and Hartzog. *A Duty of Loyalty*.