MOZILLA'S RESPONSE TO THE FEDERAL TRADE COMMISSION'S ADVANCE NOTICE OF PROPOSED RULEMAKING ON COMMERCIAL SURVEILLANCE AND DATA SECURITY

(Commercial Surveillance ANPR, R111004)

November 4, 2022

Table of Contents

I. Mozilla and its Role in the Web Ecosystem

A. Introduction

Mozilla is the maker of the open-source Firefox web browser, the Pocket "read-it-later" application and other products and services that collectively are used by hundreds of millions of individuals around the world. Mozilla is also a global community of contributors and developers who work together to keep the internet open and accessible for all. As a mission-driven technology company owned by a not-for-profit foundation, we are dedicated to putting people in control of their online experience, and creating an internet that is open and accessible to all. To fulfill this mission, we are constantly investing in the security of our products and the privacy of our users.

B. Mozilla, Firefox & Our Public Mission

For Mozilla, privacy is not optional. It is an integral aspect of our Manifesto, where Principal 4 states that Individuals' security and privacy on the internet are fundamental and must not be treated as optional. We put privacy first in our own products with features like Enhanced Tracking Protection (ETP)[1], Total Cookie Protection (TCP)[2], DNS over HTTPS,[3] and our end-to-end encrypted Firefox Sync service.[4] However, we know we cannot do it alone. Consequently, we also work to advance privacy protections in the industry more widely, working with other browser makers, ad networks, publishers, and advertisers to put forward proposals[5] that would make online advertising less privacy-invasive and to improve people's privacy. Technical privacy protections by companies, however, are insufficient - complementary privacy regulation is necessary.

[1] Deckelmann, Selena. "Latest Firefox rolls out Enhanced Tracking Protection 2.0." August 4, 2020. https://blog.mozilla.org/en/products/firefox/latest-firefox-rolls-out-enhanced-tracking-protection-2-0-blocking-redirect-trackers-by-default/

[2] Huang, T,Hofmann, J and Edelstein, A. "Firefox 86 Introduces Total Cookie Protection." February 23, 2021. https://blog.mozilla.org/security/2021/02/23/total-cookie-protection/,

[3] Deckelmann, Selena. "Firefox continues push to bring DNS over HTTPS by default for US users." February 25, 2020. https://blog.mozilla.org/en/products/firefox/firefox-continues-push-to-bring-dns-over-https-by-default-for-us-users/

[4] Ritter, Tom. "Privacy by Design: How we build Firefox Sync." November 13, 2018. https://hacks.mozilla.org/2018/11/firefox-sync-privacy/

[5] "Building a more privacy preserving ads-based ecosystem", Mozilla. May 28, 2021. https://blog.mozilla.org/en/mozilla/building-a-more-privacy-preserving-ads-based-ecosystem/

Mozilla promotes privacy in our public advocacy, calling for comprehensive privacy legislation, greater ad transparency, and for strong enforcement around the world.[6] Mozilla has long been a supporter of data privacy laws and regulations that empower people, including landmark state privacy laws like the California Consumer Privacy Act (CCPA)[7] and the California Privacy Rights Act (CPRA)[8], the European Union's General Data Protection Regulation (GDPR)[9], and US federal privacy proposals like the American Data Privacy and Protection Act (ADPPA)[10] - but there's so much more to do, particularly in the US, where we lag behind most of the world when it comes to recognizing consumer privacy and protecting people from indiscriminate data collection and use. We're encouraged by the important questions that the Federal Trade Commission (FTC) is asking today in the Commercial Surveillance and Data Security Advance Notice of Proposed Rulemaking (ANPR).

II. Mozilla's Thinking on Privacy Practices Online

The internet is powered by consumer data. While that data has brought remarkable innovation and services, it has also put internet consumers, and trust online, at substantial risk. We believe that everyone should have control over their personal data, understand how it's obtained and used, and be able to access, modify, or delete it. Additionally, we believe that consumers should be protected from particularly egregious practices (such as third party tracking) by default, both in the form of technical measures but also strong legal protections.

As raised in *Question 2* regarding measures companies use to protect consumer data, Mozilla has long advocated for companies to adopt better privacy practices through our Lean Data Practices methodology.[11] It includes being conscious to collect only data we need, clearly and concisely explaining what we collect, how we use it, how we mitigate risks and so on. This includes engaging consumers (making privacy policies more accessible and explaining data collection through "just-in-time" notifications), staying

[6] Mozilla Open Policy & Advocacy Blog. https://blog.mozilla.org/netpolicy/category/privacy/
[7] Davidson, Alan. "Bringing California's privacy law to all Firefox users in 2020." December 31, 2019. https://blog.mozilla.org/netpolicy/2019/12/31/bringing-californias-privacy-law-to-all-firefox-users-in-2020/
[8] "Four key takeaways to CPRA, California's latest privacy law." Mozilla Open Policy & Advocacy Blog. November 20, 2020. https://blog.mozilla.org/netpolicy/2020/11/20/here-are-four-key-takeaways-to-cpra-californias-latest-privacy-law/
[9] Kelly, M.J. "13 things to know about the GDPR." May 23, 2018. https://blog.mozilla.org/en/products/firefox/gdpr-mozilla/
[10] Hodges, Jenn Taylor. "It's Time to Pass U.S. Federal Privacy Legislation." August 24, 2022. https://blog.mozilla.org/netpolicy/2022/08/24/its-time-to-pass-u-s-federal-privacy-legislation/
[11] Soyinka, Nneka. "Practicing lean data is a journey that can start anywhere." January 26, 2022. https://blog.mozilla.org/netpolicy/2022/01/26/lean-data-practice-journey/

lean (rather than collecting, storing, and sharing indiscriminately), and build-in security (improving key security features, training for employees and vendor due-diligence). We implement these principles in our own products using our Data Privacy Principles, and to *Question 43*, we support new rules imposing data minimization limitations on companies' collection, use, and retention of consumer data. Industry efforts require complementary solutions from regulators, both in the form of guidance but also outright prohibition of egregious practices that cause significant harm to consumer interests.

It is important that any rules promulgated by the FTC govern not just the collection of data, but also the uses of that data resulting in harmful effects. The ANPR rightly seeks input on a diverse set of harms consumers are experiencing online today, from effects of addictive consumer interfaces on kids, to the use of recommendation systems, to discrimination in housing and jobs.

We have a long history of advocating for broad considerations of privacy beyond PII and in practice treat browsing history as extremely sensitive, as is evidenced[12] in our design of Firefox Sync which is based on end-to-end encryption. This leads to Mozilla having access to user-associated browsing data server-side only when consumers explicitly choose to share it with us (eg: for Rally)[13], protecting it both from malicious actors and any potential internal security lapses. Similarly, our efforts with DNS over HTTPS were designed to protect the browsing history of our users from ISPs and other surveillance risks by encrypting their DNS lookups.[14] We think that both of these approaches were industry-defining in their rollouts and establish that consumer browsing activity should be treated as sensitive data.[15] We encourage the FTC to adopt a similar approach and provide guidance to protect browsing history by default to the average consumer.

A. Opt-out mechanisms

Mozilla strongly supports the use of uniform and standardized signaling mechanisms at the platform level (such as web browsers) to allow consumers to opt-out of the sale or sharing of their personal information, including practices that track consumer activity

[12] "Private by Design: How we built Firefox Sync", Tom Ritter, Mozilla Hacks Blog,. November 13, 2018. https://hacks.mozilla.org/2018/11/firefox-sync-privacy/
[13] Mozilla Rally, https://rally.mozilla.org/
[14] Porter, Jon. "Firefox turns controversial new encryption on by default in the US." The Verge. February 25, 2020.
https://www.theverge.com/2020/2/25/21152335/mozilla-firefox-dns-over-https-web-privacy-security-encryption
[15] Green, Matthew. "Why the FBI can't get your browsing history from Apple iCloud (and other scary stories)." March 21, 2021.
https://blog.cryptographyengineering.com/2021/03/25/whats-in-your-browser-backup/

online. We encourage the FTC to use these platform-based settings as a signal of consumer choice and to require companies to honor these opt-out signals.

Firefox today blocks third-party cookie-based tracking, most recently using the progress we've made with Total Cookie Protection (TCP).[16] However, our technical protections are less suited for cases of first parties that might collect consumers' data and sell or share that data without the consumers' knowledge. As more browsers move to restrict cookies, we expect more websites to shift to this first party data collection and opaque sharing of that data behind the scenes, both of which will require regulatory intervention to address in a meaningful way.

Moreover, consumers cannot reasonably be expected to opt-out of the sale or sharing of their information individually from every party they interact with on the Internet. That is why a universal opt-out mechanism, set by the consumer, sent by the browser to all websites, and then enforced by the regulators, is so critical – as asked in *Question 81*. Mozilla has experimented with just such a setting: the Global Privacy Control (GPC). Once turned on, it sends a signal to the website's consumer visit telling them that the consumer does not want their data to be sold.

Unfortunately, the enforceability of the GPC remains ambiguous, with many businesses uncertain about the legal enforceability when they receive a signal such as the GPC. This is particularly true for companies receiving a GPC signal from consumers outside of specific jurisdictions that have codified GPC obligations in state privacy laws. Related to *Question 80* on if opt-out choices are effective in protecting against commercial surveillance, the practical impact of lack of enforceability is that businesses may simply ignore the GPC signal - especially if they have elected to use any other mechanisms to receive opt-out requests.

History shows that without a clear legal mandate, most businesses will not comply with consumer opt-out signals sent through browsers. This vacuum is the same reason that Do Not Track ("DNT") failed to gain adoption. It was eventually abandoned as a meaningful privacy mitigation by major browsers because it created a false sense of consumer protection that could not be enforced, as websites retained discretion on whether they would respect the DNT signal.

Mozilla encourages rules that expressly require business to comply with GPC - as asked in *Question 82*. Further, enforcement authorities should expect businesses to interpret the GPC as governing both the direct sale of consumer's information as well as

[16] Firefox rolls out Total Cookie Protection by default to all users worldwide, Mozilla. June 14, 2022. https://blog.mozilla.org/en/mozilla/firefox-rolls-out-total-cookie-protection-by-default-to-all-users-worldwide

the sharing of consumers' information for programmatic advertising targeting purposes. Regulators must step in to give tools like the GPC enforcement teeth and to ensure consumers' choices are honored.

B. Intra Company Data Sharing

As we've said before, regulatory interventions should limit data sharing within large technology conglomerates which have first party relationships with consumers across a variety of services.[17] Privacy regulations already require companies to be explicit with consumers about who has access to their data, how it is shared, etc. Technology conglomerates conveniently escape these rules because the individual products and services are housed within the same company. It has been argued that third party tracking identifiers are a means to counterbalance the market power of large platforms with access to masses of first party data. We believe this approach, which would seek to protect competition by leaving people's privacy at risk, is wrong. Competition regulators can better tackle dominance in first party data directly through targeted interventions restricting how data can be shared and used within the holding structures of large platforms. This approach achieves better outcomes for consumers than using regulatory authority to prop up an outdated and harmful tracking technology like third party cookies.

III. Mozilla's Thinking on Privacy Preserving Advertising

The current state of advertising online is rampant with risks for both consumers and innovation. It is a hostile place for consumer privacy, and is effectively an arms race between browser anti-tracking technologies and trackers. It's opaque by design, rife with fraud, and does not serve the vast majority of those which depend on it - from publishers, to advertisers, and of course, the people who use the open web. It is also important to note that many critical aspects of internet architecture were not created with privacy in mind.

At the same time, we believe there's nothing inherently wrong with digital advertising. It supports a large section of services provided on the web and it is here to stay, in some form. However, the ways in which advertising is conducted today - through pervasive tracking, serial privacy violations, market consolidation and lack of transparency - are not working and cause more harm than good. In particular, the linkage of data collected

[17] Tiwari, Udbhav. "Competition should not be weaponized to hobble privacy protections on the open web." April 12, 2022. https://blog.mozilla.org/netpolicy/2022/04/12/competition-should-not-be-weaponized-to-hobble-privacy-protections-on-the-open-web/

by advertisers leading to harmful decisions in other high consequence contexts (such as hiring, for example) is a growing concern (relevant to *Question 24* regarding costs and benefits of current practices).

Regulators and technology companies together have an opportunity to improve the privacy properties of online advertising—an industry that has not seen privacy advancement in many years. At Mozilla, we believe the web can do better and have been working to drive the industry in a better direction, away from pervasive and opaque web tracking. We've done so by limiting the use of third-party cookies, developing more privacy preserving ways to measure user interactions online, and advancing privacy preserving advertising. We're engaging at standards development organizations (SDOs) committing to plug the holes; this can be seen with the increasing focus on privacy at the World Wide Web Consortium (W3C). As asked in *Question 2*, the above efforts are measures that companies can implement and engage with to protect consumer data. Industry changes alone, however, are not sufficient.

Apart from technical solutions, we believe in the need for better regulation and creating more partnerships with the ecosystem within the diverse range of actors. Without a better regulatory framework, a sufficient incentive won't exist to move towards more privacy preserving techniques. And even with some of those techniques in place, various types of harm will persist that require regulatory intervention. We do not think that technical work alone will solve the problem of the dependence on data, and the risks and harms that this causes, which include:

- Disinformation (elections, politics, etc.)
- Discrimination (race, age, gender, etc.)
- Societal Manipulation (vaccines, etc.)
- Privacy Violations (leaks, breaches, etc.)

In an ideal state, a combination of new research, technical solutions, increased public awareness, and effective regulatory enforcement will reform advertising for the future of the web.

A. Advertising Business Models

Consumers are caught in a vicious cycle in which their data is collected, often without their understanding or meaningful consent, and then used to manipulate them. Much of that is driven by behavioral advertising, which incentivizes greater data collection,

creating greater privacy risk for consumers. As raised in *Questions 1 and 4* on practices companies use to surveil consumers and how they create harm, behavioral advertising allows for far more granular targeting than contextual advertising, and therefore serves to facilitate greater harm.

There is a paucity of rigorous research showing the ultimate value of behavioral advertising to the internet ecosystem. The vast majority of research assessing the benefits of behavioral advertising is produced by advertising tech (adtech) companies with a vested interest in the sale of these targeting techniques. Behavioral advertising does appear to *convert* at high rates, meaning consumers are more likely to click behaviorally targeted ads, potentially benefiting advertisers. And because websites pay for targeting tools, third-party ad-tech companies clearly benefit as well. However, the value to publishers and consumers is far less clear. Indeed, studies suggest[18] that the revenue benefit to publishers of these sophisticated behavioral advertising tools might be marginal even while advertisers pay higher rates to target ads. This indicates that behavioral advertising may be siphoning value away from publishers while also harming consumers. In fact, the actual impact that targeted advertising has on the revenue of advertisers remains quite poorly documented, and would benefit from independent research that organizations like the FTC could convene, support, and disseminate.

We believe that contextual advertising, along with limited targeting based on first party data, might serve as a viable business model for the Internet and can provide a strong signal of user intent for ad monetization. However, two main factors are preventing their wider adoption. First, there is a dire need for actual research that tests the various proposals and doesn't solely rely on claims by parties invested in maintaining the status quo to gauge their effectiveness. Second, so much of the web ecosystem today is wedded to the more sophisticated behavioral targeting tools that moving industry away from these techniques will come with a short-term cost. It is hard, as a matter of business practice and technology, to not use these tools. This is because this cost/impact is both poorly understood and there are insufficient alternatives in the marketplace to use as alternatives at a similar scale.

The Internet needs to be weaned off behavioral advertising over time and converted to a new model through a combination of new legislation and technical advances, platform changes, and regulatory enforcement. Each of these factors is equally important and must go hand in hand in moving the ecosystem forward, a role that the FTC is well positioned to encourage and oversee.

[18] The Impact of Behavioral Targeting on Ad Revenue. https://www.algorix.co/the-impact-of-behavioral-targeting-on-ad-revenue/

Moving the ecosystem away from this approach requires a concerted effort that includes:

- Browser-developers building privacy features that make it harder to create sophisticated targeting profiles based on people's browsing activity (measures to protect data, as asked in *Question 2*). We are proud to say that Firefox has been leading the way on that.
- Regulators creating costs to disincentivize bad behavior, both when data is collected without consent or when users are tricked into handing over their data, and then when data is used to manipulate or harm people (to *Question 29* – if the FTC refrained from promulgating new rules, it would perpetuate and create costs for users online).
- Advertisers and platforms providing greater transparency into targeting practices, so that regulators can understand and intervene when people are being manipulated. (As asked in *Question 83*, rules should require companies to make more information available about their surveillance practices.)

That is where our work and the work of the FTC come in. Technical changes can make it harder to violate peoples' privacy, but alone they will not solve the risks and harms in digital advertising today. Technical interventions, combined with giving regulators the framework and authority to create real costs for bad behavior, however, can shift the incentives towards less invasive, more benign forms of advertising.

IV. Deceptive Design Patterns

We welcome the FTC's ongoing focus on dark patterns (or "deceptive design patterns") and wholeheartedly agree with concerns noted in the ANPR around how such practices "trick and trap consumers"[19] and impact users' ability to make meaningful choices regarding their data, to the benefit of an online service (as raised in *Question 77*).

We know as a browser maker that deceptive design practices are pervasive and that consumers are exposed to them on a daily basis; but even while we work to protect[20]

[19] FTC Report Shows Rise in Sophisticated Dark Patterns Designed to Trick and Trap Consumers, https://www.ftc.gov/news-events/news/press-releases/2022/09/ftc-report-shows-rise-sophisticated-dark-patterns-designed-trick-trap-consumers

[20] https://support.mozilla.org/en-US/kb/how-does-phishing-and-malware-protection-work

our[21] users[22], there is little a browser can do when consumers are deceived by the parties they engage with online. Consumers are being tricked into handing over their data with deceptive patterns, then that data is used to manipulate them. This is bread-and-butter deception - the online manifestation of what the FTC was established to address - and it is critical that the FTC has the authority to take action against such deception.

Deceptive design patterns are more than a threat to peoples' privacy on the web; they affect consumers and shape markets in significant ways by influencing when, if, and how people make decisions online.[23] *Question 27* asks if a new trade regulation rule would impede or enhance competition. Continuing the FTC's important work on deceptive design patterns as part of a rulemaking would benefit both people's online privacy and enhance competition, enabling users to make informed and effective choices and supporting a vibrant ecosystem.

Deceptive design practices by operating system providers can exclude independent companies or rivals seeking to offer consumers a different experience or innovation, beyond the control of any big tech platform. When operating systems use these practices to direct consumers to their own products, they are often leveraging their platform market power in another market.

To give an example of deceptive patterns deployed to distort browser competition, operating systems often make it difficult for consumers to set alternative browsers (such as Mozilla's Firefox) as a default browser by using complex visuals and adding unnecessarily complicated steps that ordinary consumers would not remember or may not feel technically competent to perform.[24] If a consumer has overcome these hurdles, they may face further deceptive patterns to overturn their default selection. For example, prompts following updates in the Windows operating system which include confusing or misleading graphics (such as a lock icon and fingerprint) that falsely suggests that reverting to the operating system provider's browser is a security-related choice for the user.[25] By manufacturing concerns about the merits and risks of

[21] https://support.mozilla.org/en-US/kb/what-does-your-connection-is-not-secure-mean
[22] "Firefox rolls out Total Cookie Protection by default to all users worldwide." Mozilla Blog. June 14, 2022. https://blog.mozilla.org/en/products/firefox/firefox-rolls-out-total-cookie-protection-by-default-to-all-users-worldwide/
[23] Sunstein, Cass R. "Choosing Not to Choose: Understanding the Value of Choice." 1st edition. New York, NY: Oxford University Press, 2015
[24]Five Walled Gardens, A Report from Mozilla. September 2022. https://research.mozilla.org/files/2022/09/Mozilla_Five-Walled-Gardens.pdf
[25] Five Walled Gardens

third-party competitors, this affiliated preferencing tactic can undermine fair competition and diminish consumers' ability to benefit from using the applications of their choice.[26]

This use of deceptive design patterns to self-preference[27] results in consumer harms such as the below: [28]

- **Limited or frustrated choice -** an operating system provider making it difficult or impossible for a consumer to switch browsers ultimately removes their ability to choose for themselves. It also hampers existing competitors and deters new products from entering the market and providing increased choice.

- **Lower quality -** where the monetary price for consumers is zero (as is the case for browsers) providers might be expected to compete on quality. But without effective competition from independent browsers, consumers may receive products which are lower quality.[29]

- **Lower innovation -** linked to quality is innovation. Consumers miss out on developments (for example, improved features and functionality). And a reduced likelihood of disruptive innovation might be accompanied by reduced choice for consumers.

- **Poor privacy -** consumers can be left with a product which subjects them to compulsory data sharing, misuse of data or other privacy harms. These outcomes can be an indication of low quality caused by ineffective competition.

- **Unfair contracts -** without proper choice, consumers may be forced to enter into contracts which might be exploitative or unfair.

Potential FTC trade regulation rules should tackle these practices that prioritize business objectives over true consumer choice and agency.

---

[26] For instance, BEUC notes in its paper on the Digital Markets Act that "[e]ven where it is technically feasible for a consumer to switch a service, she/he may, for example, be bombarded with repeated and "intimidating" messages about the purported disadvantages or dangers of switching, or this may be made so time-consuming or complex that the consumer gives up. Such tactics can be just as effective as technical barriers". *Position paper of BEUC on DMA,* 1.4.2021, pp. 7-8.
[27] See, for example, Report of the Digital Competition Expert Panel for the UK Government, "Unlocking Digital Competition", March 2019
[28] Five Walled Gardens
[29] Feature and security imitations on Internet Explorer, before Firefox came along; or feature limitations with Amazon's Silk browser or Apple's Safari browser are such examples.

V. Automated Decision Making Systems

In reference to *Question 53*, we would like to focus on algorithmic *harms*. This category can be both a subset of algorithmic *error* where erroneous automated decision-making systems (ADMS) negatively affect individuals or groups of individuals. However, it can also be distinct from error, where ADMS that function as intended still have harmful downstream effects. The FTC should take this distinction into account and pay particular attention to how ADMS can cause or contribute to harm regardless of whether they function erroneously or not.

For years, research and investigative reporting have again and again uncovered instances of ADMS that cause or enable discrimination, surveillance, or other harms to individuals and communities.[30] Further, Mozilla's own research has shone a light on Tinder's opaque and unfair personalized pricing algorithms[31] (related to *Question 61*) as well as YouTube's ineffective user controls.[32] This can be rooted in a variety of different causes, all of which need to be addressed: from the data used to train and test machine learning-based systems to the design of an ADMS to the very purpose for which they are meant to be used. For example, research by Mozilla fellows Abeba Birhane, Deborah Raji, and others has repeatedly pointed to harmful and toxic data as well as privacy risks in datasets widely used for machine learning. For instance, they have uncovered misogynistic and racist imagery in large computer vision datasets[33] and pointed to issues of obtaining genuine consent in the construction of such datasets[34].

Harms caused by ADMS are of particularly grave consequence when these are deployed in critical areas and where they affect, for example, people's livelihoods, safety, or liberties — be it a rejected loan, wrongful termination of a job, or discriminatory pricing of goods and services. Taking effective countermeasures and

---

[30] See, for example, the AI Incident Database for a non-exhaustive catalog of such cases: https://incidentdatabase.ai/

[31] Mozilla, "New Research: Tinder's Opaque, Unfair Pricing Algorithm Can Charge Users Up to Five-Times More For Same Service," 2022, https://foundation.mozilla.org/en/blog/new-research-tinders-opaque-unfair-pricing-algorithm-can-charge-users-up-to-five-times-more-for-same-service/

[32] Ricks and McCrosky, "Does This Button Work? Investigating YouTube's ineffective user controls", 2022, https://foundation.mozilla.org/en/research/library/user-controls/report/

[33] See, for example, Birhane et al., "Multimodal datasets: misogyny, pornography, and malignant stereotypes", 2022, https://arxiv.org/pdf/2110.01963.pdf; Prabhu & Birhane, "Large Datasets: A Pyrrhic Win for Computer Vision?", 2020, https://arxiv.org/pdf/2006.16923.pdf

[34] See, for example, Paullada et al., "Data and its (dis)contents: A survey of dataset development and use in machine learning research", 2021, https://www.sciencedirect.com/science/article/pii/S2666389921001847

providing affected people with remedies — as well as questioning whether ADMS should be used in such critical contexts at all — is imperative in this context.

Companies, both those developing and those deploying ADMS, are already in a position to prevent or mitigate such harms. While neither easy nor straightforward, people across industry, academia, civil society, and the public sector have developed a range of tools and frameworks that can contribute to ensure that ADMS help rather than harm people. Yet, the incentives for broad adoption of these are lacking. Rules that change this incentive structure and make companies price in the externalities caused by the ADMS they develop or put to use is therefore clearly necessary.

In the U.S., much important work is already underway to help companies meet their responsibility. For instance, NIST's AI Risk Management Framework as well as the White House OSTP's Blueprint for an AI Bill of Rights, both of which Mozilla weighed in on[35], list a broad range of important considerations and concrete measures for actors along the ADMS value chain. At the same time, adoption of either of these instruments is purely voluntary. In a next step (and with regard to *Question 56*), enforceable rules that hold developers and deployers of ADMS to a higher standard should follow.

Such rules should build, at minimum, on the following pillars:

- **Transparency:** Enabling systemic transparency is critical to ensuring that ADMS can be effectively scrutinized by different actors and that rules can be enforced. This includes meaningful transparency vis-à-vis those interacting with or affected by ADMS as a precondition for seeking remedy (see below) — only if they are aware that (unfair) outcomes were affected by an ADMS can people take action to correct them. Therefore, there should be disclosure obligations for deployers of ADMS where ADMS materially contribute to decisions with a potentially significant impact on consumers.

  Additional measures could aim at enabling more independent research on ADMS (as well as on testing and training datasets) and/or create and encourage (or require) registration in a public registry of ADMS used in critical areas, as proposed in the EU's Artificial Intelligence Act.[36] Both of these measures could

[35] Gahntz, M and Hodges, J. "Managing AI Risks — Mozilla Files Comments to NIST." June 8, 2022. https://foundation.mozilla.org/en/blog/managing-ai-risks-mozilla-files-comments-to-nist/; Alotta, J Bob. "The White House AI Blueprint is a Welcome Signal." Oct 6, 2022. https://foundation.mozilla.org/en/blog/the-white-house-ai-blueprint-is-a-welcome-signal/

[36] "How to make sure the EU's AI Act delivers on its promise." April 2022. https://foundation.mozilla.org/en/blog/how-to-make-sure-the-eu-ai-act-delivers-on-its-promise/

enhance public oversight and the public's understanding of how ADMS function and how they are used.

- **Accountability:** The effectiveness of any regulatory framework requires both an effective allocation of responsibility and robust oversight. To this end, regulators should pay close attention to who should be held accountable for which obligations, from developers to deployers and across the lifecycle of ADMS. At minimum, the FTC should mandate developers and deployers to keep basic technical documentation about ADMS design, development, and deployment, including model and dataset documentation (also with regard to data provenance, collection, and curation). This would raise the standard of due diligence across industry and provide the FTC with a better starting point for investigations without unduly overburdening companies.

  Additionally, the FTC should be equipped with both the resources and expertise necessary to ensure adherence to the rules. Specifically, the FTC should build internal capacity to engage in regulatory audits, building on the work of and collaborating with outside experts from the emerging field of algorithmic auditing. Mozilla is itself funding work aimed at developing new and better auditing tools, both through the Open Source Audit Tooling project led by Mozilla fellow Deborah Raji and through the Mozilla Technology Fund.[37] We hope the FTC will join us in this effort to advance the AI auditing ecosystem and encourage companies to undergo rigorous external audits of their ADMS, particularly of ADMS used in critical areas.

- **Redress:** Where people have already been harmed, they need access to mechanisms for redress. For this reason, any regulatory framework for ADMS should ensure that both affected individuals and groups of individuals can take action against those that have caused such harm. The FTC should therefore introduce a formal complaint mechanism for consumers and organizations representing their interests to lodge complaints with the FTC. Especially if tied to effective transparency measures (as mentioned above), this holds the promise to bring cases of potentially unfair or deceptive ADMS to the FTC's attention and point them into promising directions for prospective investigations.

[37] Raji. "It's Time to Develop the Tools We Need to Hold Algorithms Accountable." 2022. https://foundation.mozilla.org/en/blog/its-time-to-develop-the-tools-we-need-to-hold-algorithms-accountable/ ; Mozilla Technology Fund: Auditing Tools for AI Systems,https://foundation.mozilla.org/en/what-we-fund/awards/mozilla-technology-fund-mtf/

By following these suggestions, the FTC could contribute to making ADMS used in consumer products and services more trustworthy and lay the ground for more targeted rules down the line.

VI. Data Access for Effective Enforcement of Trade Rules

As part of, and in addition to, the rulemaking process, we encourage the FTC to strengthen the mechanisms that empower policymakers and trusted experts to better understand what is happening on major internet platforms. Regulators like the FTC require greater visibility and data access to address the harmful practices that abound online today.

A large amount of harm happens on major tech platforms outside the view of regulators and the public. These platforms offer highly sophisticated targeting tools that allow content producers to narrowly segment their audience, tailor content accordingly, and reach people most susceptible to their messages. Each consumer has their own individualized, potentially misleading or deceptive experience. This highly personalized experience means that harm enabled by platforms through their targeting systems is not easily identified by regulators, watchdog groups, or researchers. Because the experience is so personalized, harm can only be shown anecdotally, when a particular piece of content appears to be harmful and when the regulator is somehow made aware of that content. There is dangerously little insight into what people experience, what ads are presented to them and why, and what content is recommended and why. This creates an asymmetry of information between those who produce harmful content and those seeking to understand it.

Simply put, any new trade rules will fail if the FTC does not have sufficient visibility to systematically identify violations of those rules, as is the case today. To address this, we need greater access to platform data (subject to strong user privacy protections), greater research tooling, and greater protections for researchers. This is why Mozilla has invested in building tools for researchers and why we support legislative solutions to provide greater insight into online disinformation, discrimination, and deception currently hidden from the public and from regulators.

Key pieces of this agenda, which the FTC should look to both support and take advantage of, include:

A. Bolstering Ad Disclosure Regimes

At Mozilla, we want to see a much stronger disclosure regime for online targeting. This is why we advocated for Article 30 of the Digital Services Act, which requires large online platforms to publicly disclose the ads, along with ad metadata, running on their platforms.[38] These disclosure tools may serve as a global regulatory resource and, as they come online, the FTC should seek ways to take advantage of them.

Similarly in the U.S., it is important to consider how ad disclosure might also be incorporated in an FTC rulemaking. To provide transparency into ad-based harm, we believe companies should be required to provide data that includes content of ads, who is paying for ads, how much is spent, how they are targeted, how algorithms optimize for the "best ads," and other specifications. This is something we have been driving at Mozilla by publicly disclosing our ads and the targeting parameters in our advertising campaigns. You can find all of our ads and targeting parameters in the "Targeted Advertising Disclosures" section of Mozilla's bi-annual transparency report.[39] Such practices should be used across the industry to provide insight regulators need to identify algorithmic discrimination.

Current users facing ad transparency mechanisms - those that would allow a consumer to click an icon in an advertisement and learn how they were targeted - are not enough. These approaches do not provide sufficient information to regulators to identify potential discrimination systematically. That is why our preferred approach is one that would require companies to disclose their targeting parameters publicly, in machine formats readily available to everyone, just as we have done in our own campaigns.[40] We encourage the FTC to consider transparency requirements that would provide greater insight into this targeting – as asked in *Question 83*. Publishers, consumers and oversight bodies should be able to understand how or why they are being targeted.

B. Building and Supporting Data Platforms for Oversight

Unfortunately, the data and tools made available by major platforms to understand topics like deception and discrimination online remain inadequate. Most of the voluntary

[38] Mozilla position paper on the EU Digital Services Act, May 2021. https://blog.mozilla.org/netpolicy/files/2021/05/Mozilla-DSA-position-paper-.pdf
[39] Mozilla Transparency Report, Reporting Period: January 1, 2022 to June 30, 2022 https://www.mozilla.org/en-US/about/policy/transparency/jan-jun-2022/
[40] Shepherd, Lindsay. "Our approach to advertising on Facebook platforms." May 27, 2021. https://blog.mozilla.org/en/mozilla/mozilla-approach-to-advertising-on-facebook-platforms/

public-facing measures by major platforms have failed. Researchers have found, for example, that ad transparency tools are often nearly unusable.[41]

We need far more robust research tools along with a deep pool of subject matter experts capable of taking advantage of these tools. Efforts like the Markup's CitizenBrowser project, New York University's Ad Observer, Who Targets Me, and Mozilla's own Rally and Youtube RegretsReporter[42] Projects all aim to pry open closed systems and provide actionable insights to regulators. The research produced by these teams has begun to shift the understanding and perceptions of legislators, regulators, and the public.

Mozilla Rally has worked with journalists, academics, and non-profit researchers to develop and launch projects to understand the public's experiences online through user contributed browser and interaction data. Current Rally studies[43] are providing insight into topics such as how consumers get information about COVID-19, browser search defaults, and local news consumptions. Browsing data contributed by Mozilla Rally users in one study helped The Markup uncover hospitals sending sensitive patient data to Facebook, in violation of HIPPA laws.[44] Mozilla's RegretsReporter project is the world's largest crowdsourced investigation into YouTube's recommendation algorithm. Data contributed by RegretsReporter users allowed Mozilla researchers to discover and document YouTube's recommendation algorithm routinely recommending videos that violated the platform's own community guidelines[45] (and which were later removed from the platform) and how YouTube's user control features do not work as advertised by the company[46]. The project was ultimately cited in European legislation as rationale for stricter transparency requirements for user-generated content recommendation systems.

These tools are still embryonic and require continued development. It is critical that the FTC find opportunities to support the development of these tools, as it did when it

[41] Edelson, Laura. "Facebook's political ad spending numbers don't add up." Medium, October 12, 2020. https://medium.com/online-political-transparency-project/transparency-theater-facebooks-political-ad-spending-numbers-don-t-add-up-d7a85479a002
[42] https://foundation.mozilla.org/en/youtube/regretsreporter/
[43] Mozilla Rally current studies, https://rally.mozilla.org/current-studies/
[44] "Facebook Is Receiving Sensitive Medical Information from Hospital Websites", Todd Feathers, Simon Fondrie-Teitler, Angie Waller, and Surya Mattu. The Markup, June 16, 2022. https://themarkup.org/pixel-hunt/2022/06/16/facebook-is-receiving-sensitive-medical-information-from-hospital-websites
[45] https://assets.mofoprod.net/network/documents/Mozilla_YouTube_Regrets_Report.pdf
[46] https://foundation.mozilla.org/en/research/library/user-controls/report/

weighed in[47] on the ongoing dispute between Facebook and the NYU Ad Observer, and to take advantage of the data these tools provide to support its regulatory actions.

C. Mandating a Safe Harbor to Protect Public Interest Researchers

Mozilla supports creating a legal safe harbor for public-interest research and journalism that respects user privacy. This is key to shedding light on hidden harms and disinformation.

Experts engaged in research need to be protected and free from threat of legal action. Mozilla often hears of researchers[48] who are concerned that companies or governments may take legal action against them for their legitimate research – including civil or criminal penalties under laws such as the Computer Fraud and Abuse Act (CFAA), violations of Terms of Service, and more. Facebook, for example, has blocked research tools and threatened legal action against researchers seeking to investigate election integrity and misinformation online.

These actions by platforms not only put researchers themselves at legal risk, but also stifle vital transparency into real world harm by deterring individuals and institutions doing critical work. Indeed, research tools and initiatives most vital to public interest - most capable of identifying patterns of harm or threats to information integrity on major tech platforms - may receive the greatest scrutiny and be subject to the greatest legal exposure. This is the pattern we have seen with New York University's Ad Observatory project, which has offered research tools effective at identifying harms on tech platforms, and as a result, has had to withstand sustained legal attack.

A safe harbor would protect and promote research in the public interest as long as researchers handle data responsibly and adhere to professional and ethical standards, such as those developed to support the vetting process described above. There is enormous value this can provide to the public. Mozilla has one of the earliest Bug Bounty programs in software. We make clear that we will not threaten or bring any legal action against anyone who makes a good faith effort to comply with our vulnerability notification policy because this encourages security researchers to investigate and disclose security issues. Their research helps make the internet a safer place.

---

[47] "Letter from Acting Director of the Bureau of Consumer Protection Samuel Levine to Facebook," August 5, 2021.
https://www.ftc.gov/blog-posts/2021/08/letter-acting-director-bureau-consumer-protection-samuel-levine-facebook

[48] New Approaches to Platform Data Research, February 2021.
https://publicinfrastructure.org/2021/02/09/new-approaches-to-platform-data-research/

VII. Conclusion

Mozilla supports the FTC's effort to collect comments on the questions posed in the ANPR and is encouraged by the prospect of legislation[49] or regulation in this space. As set out above, the practices surrounding consumer data on the internet today, and the resulting societal harms, have put people and trust at risk. The future of privacy online requires industry to step up to protect and empower people, and demands that lawmakers and regulators implement frameworks that create the ecosystem and incentive for a better internet ahead.

Contact for Additional Information

Jenn Taylor Hodges, Director of US Public Policy and Government Relations, Mozilla Corporation - jhodges@mozilla.com

[49] Hodges. "It's Time to Pass U.S. Federal Privacy Legislation."