ROB BONTA
Attorney General of California
ANYA M. BINSACCA, SBN 189613
Supervising Deputy Attorney General
NICOLE KAU, SBN 292026
ELIZABETH K. WATSON, SBN 295221
Deputy Attorneys General
455 Golden Gate Avenue, Suite 11000
San Francisco, CA 94102-7004
Telephone: (415) 510-3847
E-mail: Elizabeth.Watson@doj.ca.gov
Attorneys for Defendant

IN THE UNITED STATES DISTRICT COURT

FOR THE NORTHERN DISTRICT OF CALIFORNIA

SAN JOSE DIVISION

NETCHOICE, LLC d/b/a NetChoice,

Plaintiff,

v.

ROB BONTA, ATTORNEY GENERAL OF THE STATE OF CALIFORNIA, in his official capacity,

Defendant.

5:22-cv-08861

DECLARATION OF SERGE EGELMAN, PH.D. IN SUPPORT OF DEFENDANT'S OPPOSITION TO PLAINTIFF'S MOTION FOR PRELIMINARY INJUNCTION

I, Serge Egelman, Ph.D., declare and state as follows:

1. I submit this declaration in support of Defendant's Opposition to Plaintiff's Motion for Preliminary Injunction.

BACKGROUND & QUALIFICATIONS

2. I am the Research Director of the Usable Security & Privacy Group at the International Computer Science Institute (ICSI), which is a non-profit research institute affiliated with the University of California, Berkeley. I also hold a position as a research scientist within the Electrical Engineering and Computer Sciences (EECS) Department at the University of California, Berkeley. I received my Ph.D. from Carnegie Mellon University's School of Computer Science. My research has been cited over 11,000 times, and my h-index—the most common metric for scientific impact—is 50.[1,2]

3. I have been performing research into online privacy for nearly twenty years. My research focuses on the interplay of online privacy and security and human factors; in short, I study consumer privacy and security decision making, consumer privacy preferences, privacy and security expectations, and how those expectations comport with reality. I have served as an invited expert for several web standards efforts that pertained to privacy and security, and have received over a dozen research awards (including best paper awards from two European data protection authorities, AEPD in Spain and CNIL in France; the USENIX Security Symposium Distinguished Paper Award, from one of the top academic computer security conferences; and seven paper awards from the Special Interest Group on Computer-Human Interaction [SIGCHI], the top human-computer interaction conference).

4. Over the past decade, my laboratory has been studying the mobile app ecosystem, which has included building tools to detect when personal information is accessed by mobile apps and the third parties with whom they share it. We have used these tools in peer-reviewed published research studies about consumer privacy, including examining mobile apps' compliance with various privacy regulations and platform policies.

[1] https://en.wikipedia.org/wiki/H-index
[2] https://scholar.google.com/citations?user=WN9t4n0AAAAJ&hl=en

5. One research study performed by my laboratory demonstrated that a majority of child-directed Android apps appeared to be violating COPPA,[3] which led to major policy shifts by both Google and Apple, makers of the two leading mobile platforms. I have since been invited to give keynotes at several international conferences on child development and technology as an expert on online privacy as it pertains to children. I have also testified before the U.S. Senate on how COPPA can be improved to match the realities of modern technology, and have been asked to provide feedback on draft legislation from members of both houses of Congress.

6. My *curriculum vitae*, which sets forth my experience and credentials more fully, is attached as Exhibit A.

7. I have testified as an expert in the following cases:

- *Hart v. TWC Prod. & Tech. LLC*, 526 F. Supp. 3d 592, Case No. 20-cv-03842-JST (N.D. Cal. 2021)
- *District of Columbia v. Town Sports International, LLC*, Case No. 2020 CA 003691 B (D.C. Sup. Ct. 2020)
- *In re Vizio, Inc., Consumer Privacy Litig.*, 238 F. Supp. 3d 1204, (C.D. Cal. 2017)
- *In re Linkedin User Privacy Litigation*, 932 F. Supp. 2d 1089 (N.D. Cal. 2013)
- *In re Netflix Privacy Litigation*, Case No.: 5:11-CV-00379 EJD (N.D. Cal. 2012)

8. I am being compensated in the above-entitled case at an hourly rate of $400/hour for preparing this declaration. My compensation is not in any way dependent on the outcome of this or any related proceeding.

9. The opinions in this declaration are my expert opinions, which are based on my education and training, my peer-reviewed published research and the research of others, my knowledge of relevant technologies (including my reading of the public technical documents offered by NetChoice's members about their capabilities), as well as my reading of the legislation.

[3] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. *"Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale.* Proceedings on Privacy Enhancing Technologies (PoPETS), 2018(3):63–83.

10. I have reviewed AB 2273, the California Children's Age Appropriate Design Code (AADC) Act. In my expert opinion, this law is necessary to address realities of modern technology that have resulted in the exploitation of minors; its provisions are reasonable and technically feasible to adopt (i.e., the technologies necessary to comply are already in widespread use by NetChoice's members), and I believe that they are substantially similar to policies in other jurisdictions within which NetChoice members operate. The law only applies to services that are likely to be used by children (rather than all online services), and only requires that companies take steps to limit harm to children, allowing them and their parents to make more informed decisions about their online activities and dissemination of their personal information. Services not likely to be used by children are unlikely to be impacted by this legislation; child-directed services can comply by simply limiting privacy-invasive tracking and considering potential harms to children. Similar laws already exist in other sectors, which society has accepted: that convenience stores cannot sell tobacco products and alcohol to minors is not viewed as tyrannical overreach or limitations on "freedom to innovate," but instead as a commonsense safeguard.

COLLECTION & USE OF PERSONAL INFORMATION ONLINE

11. The "free" Internet is subsidized through the collection of users' personal information for both advertising and analytics purposes. In the case of advertising, this means showing Internet users ads that are specifically tailored to their inferred interests. In the case of analytics, this means observing how users interact with the service in order to maximize its profitability (e.g., strategically placing in-app purchase opportunities based on users' in-app behaviors, identifying the users most likely to buy expensive items based on their inferred demographics, manipulating users into spending more time using a service, etc.). In other cases, this may mean straight up selling the user data to third parties so that they may perform these activities and other yet-unknown use cases.

12. Because so much of the Internet is supported by advertisements, one key metric that online services use is known as "engagement," which refers to the amount of time that consumers spend using a service or the frequency of interactions that consumers have with that service. That is, the more time consumers spend using a service that displays ads, the more ads

that consumers are likely to be shown. Thus, many services collect analytics data to measure engagement and then use this data to develop features that are likely to lead to greater levels of engagement. For example, social media platforms have discovered that emotionally manipulating consumers based on what content they are shown results in greater levels of engagement,[4] and therefore many of these platforms are optimized for this purpose.[5] Facebook researchers previously showed that they can manipulate users' content feeds to intentionally upset people, and that this emotional manipulation can then be spread to users' connections.[6] Another study found that moral outrage resulted in more "retweets" on Twitter (i.e., greater engagement due to users resharing a post),[7] and others have found similar results on other social media platforms.[8] For this reason, conspiracy theories also result in greater levels of engagement and spread quickly online,[9] since they make their believers angry (and cater to confirmation bias). Thus, platforms are incentivized to make their users angry so that there is more "engagement," which results in more advertisements being viewed (due to more time spent on the platform), resulting in more revenue.

13. Advertisements are targeted at users based on inferences about those users' interests. Individual users' interests are inferred based on data automatically collected from them: the services they use, how they use them, from where they use them, and so forth. In short, online and offline activities are tracked, which allows companies to maintain detailed profiles of individual user behavior, which in turn is used to predict users' interests, preferences, and even demographics. The collected information may be used to predict a consumer's religion, health

[4] Gilad Edelman, "Facebook Quietly Makes a Big Admission." *Wired,* August 31, 2021. https://www.wired.com/story/facebook-quietly-makes-big-admission-political-content/

[5] Filippo Menczer, "How 'engagement' makes you vulnerable to manipulation and misinformation on social media." *The Conversation,* September 20, 2021. https://theconversation.com/how-engagement-makes-you-vulnerable-to-manipulation-and-misinformation-on-social-media-145375

[6] Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock. "Experimental evidence of massive-scale emotional contagion through social networks." *Proceedings of the National Academy of Sciences* 111.24 (2014): 8788-8790.

[7] William J. Brady *et al.* "Emotion shapes the diffusion of moralized content in social networks." *Proceedings of the National Academy of Sciences* 114.28 (2017): 7313-7318.

[8] Rui Fan *et al*. "Anger is more influential than joy: Sentiment correlation in Weibo." *PloS one* 9.10 (2014): e110184.

[9] Soroush Vosoughi, Deb Roy, and Sinan Aral. "The spread of true and false news online." *Science* 359.6380 (2018): 1146-1151.

conditions, sexual orientation, or political affiliation; some of this information may be revealed by the phone's location alone (e.g., where they live, who they live with, where they work, etc.), or even by just the name of the app that is being used (e.g., revealing sexual orientation, religion, age, or socioeconomic status).

14. Tracking of users' online behaviors is made possible by "persistent identifiers." An identifier is any piece of information that allows an individual—or device—to be uniquely identified. "Persistent" identifiers are identifiers that tend to not change over time.[10] For example, motor vehicles have persistent identifiers in the form of license plates: a license plate uniquely identifies a vehicle and vehicles tend to have the same license plates over time. Thus, if someone records all the license plates at a particular place over time, they can determine how many times in that period any individual vehicle was there (and thus infer their operators' activities). Similarly, if license plates are recorded at many different locations and that data is combined, one could reconstruct the movements of individual vehicles. Thus, combining a persistent identifier with information about where that identifier was observed (e.g., a website or mobile app) allows a data recipient to reconstruct an individual's activities. Using this knowledge, one could infer information about their routines, preferences, demographics, and even relations and social connections. It is for this reason that persistent identifiers, including ones that identify personal devices—because they tend to be used by one individual—are categorized as personal information under various privacy laws (e.g., CCPA,[11] COPPA,[12] HIPAA,[13] GDPR,[14] GLBA[15]).[16]

15. Online advertisements need not use consumers' personal information: while the *behavioral* or *targeted* advertising described in the prior paragraphs relies on collecting personal information to infer users' interests, *contextual* advertising does not. Contextual advertising refers to choosing ads based on what the user is doing in the moment: the type of website or online

[10] https://www.nnlm.gov/guides/data-glossary/persistent-unique-identifier
[11] Cal. Civ. Code § 1798.140(15).
[12] 15 U.S.C § 6501(8)(F).
[13] 45 C.F.R. § 164.514(b)(2)(i).
[14] GDPR Art. 4 (1).
[15] 16 C.F.R. § 313.3.
[16] See, e.g., https://www.federalregister.gov/documents/2021/12/09/2021-25736/standards-for-safeguarding-customer-information

service that the user is currently visiting, which is where the ad is to appear, and not based on a collected profile or tracking information. For example, a mattress review website does not need to collect personal information to know that visitors might be receptive to ads for mattresses or bedding. By definition, contextual advertising does not require the collection of consumers' personal information, because it does not rely on the tracking of their online activities.

16. In addition to questionable economic benefits, over half a century of published research on consumer behavior and preferences has demonstrated that consumers are opposed to this type of tracking by businesses. For example, when Westin performed consumer surveys on public privacy perceptions going back to the 1970s,[17] he consistently found that a majority of the U.S. public are either "very" or "somewhat" concerned with how their personal information is collected and used by businesses. In 2001, one study found that as many as 64% of consumers refused to shop online due to privacy concerns.[18] A Pew survey from 2020 found that more than half of Americans have refused to use certain products or services due to privacy concerns.[19] In the past two decades, as more and more aspects of daily life have moved online, many consumers have also simply become resigned to having their information used in objectionable ways.[20] A 2019 Pew survey of consumers found that 62% of Americans do not believe it is possible to "go through daily life without companies collection data about them," 79% are very or somewhat concerned about this, and 81% believe the risks of collecting this data outweigh the benefits.[21]

17. While consumers are overwhelmingly opposed to this type of tracking and the profiling and resale of their information that it supports (one study of U.S. consumers found that

[17] Ponnurangam Kumaraguru and Lorrie Faith Cranor. "Privacy indexes: a survey of Westin's studies." *Carnegie Mellon University Tech Report* CMU-ISRI-5-138, 2005.

[18] M J. Culnan and Milne, G. R. "The Culnan-Milne Survey on Consumers & Online Privacy Notices: Summary of Responses." In *Interagency Public Workshop (Ed.) Get Noticed: Effective Financial Privacy Notices*, Washington, D.C., 2001.

[19] Andrew Perrin, "Half of Americans have decided not to use a product or service because of privacy concerns." *Pew Research Center,* August 14, 2020. https://www.pewresearch.org/fact-tank/2020/04/14/half-of-americans-have-decided-not-to-use-a-product-or-service-because-of-privacy-concerns/

[20] Nora A. Draper and Joseph Turow. "The corporate cultivation of digital resignation." *New media & society* 21.8 (2019): 1824-1839.

[21] Pew Research Center. "Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information." Nov. 15, 2019. https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/

up to 86% do not want ads that are tailored based on their online activities),[22] consumers nonetheless continue to engage with services that appear to conflict with their stated privacy preferences. This is known as the "privacy paradox." Some stakeholders like to point out this disconnect and use it to disingenuously claim that it means that consumers do not "really" care about privacy. But the published research on the privacy paradox demonstrates that this argument is incorrect, and that there are several rational explanations for the privacy paradox, which include lack of awareness of data collection methods, poor usability, mismatched incentives, and perceived lack of agency.

18. In many cases, consumers simply do not understand when they are making decisions that will impact their privacy. For example, in a series of studies that I co-authored,[23,24,25] we presented subjects with different search engine interfaces, including one that annotated search results with privacy information; subjects were instructed to use the search engine to buy items from merchants of their choice. While all subjects expressed strong privacy preferences in a survey administered prior to the study (i.e., subjects were specifically screened for strong privacy preferences, so that we could explicitly test whether interface design impacted their ability to act on those preferences), we observed that without information about privacy practices presented in an easily-accessible manner, subjects made purchases from the cheapest merchants. Whereas when search results were annotated with privacy ratings, subjects were significantly more likely to make purchases from merchants with more agreeable privacy policies (i.e., better aligned with participants' stated privacy preferences), even paying more money to do so. These and other studies demonstrate that people often act in ways that seem contrary to their

[22] J. Turow, J. King, C. J. Hoofnagle, A. Bleakley, and M. Hennessy (2009). "Americans Reject Tailored Advertising and Three Activities That Enable It." https://doi.org/10.2139/ssrn.1478214

[23] Janice Y. Tsai Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. "The effect of online privacy information on purchasing behavior: An experimental study." *Information systems research* 22, no. 2 (2011): 254-268.

[24] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. "Timing is everything? The effects of timing and placement of online privacy indicators." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 319-328. 2009.

[25] Julia Gideon, Lorrie Cranor, Serge Egelman, and Alessandro Acquisti. "Power strips, prophylactics, and privacy, oh my!." In *Proceedings of the Second Symposium on Usable privacy and security*, pp. 133-144. 2006.

stated privacy preferences when they are not fully aware of a business's privacy practices (e.g., due to the well-documented problems with the "notice and consent" framework, i.e., expecting consumers to read and understand privacy policies, which I describe in subsequent sections).

19. In other cases, convoluted user interfaces make it difficult for consumers to understand how to make privacy-protective decisions. This poor usability often results in consumers sharing personal information without ever being aware of it. For example, while studies have shown that consumers have concerns about sharing personal information with the wrong audiences on social media, they nonetheless continue to overshare,[26] which has been shown to be the result of difficult-to-use privacy settings interfaces (or mismatches between the design of those interfaces and users' mental models).[27] One early study on the use of Facebook found that while participants expressed strong privacy preferences, they nonetheless shared sensitive information because more than one-in-five did not understand what Facebook's privacy settings did or how to use them, and therefore did not change them from the overly-permissive defaults.[28] In a study of file-sharing software, researchers discovered that due to convoluted privacy settings interfaces, many users were inadvertently sharing their entire hard drives.[29] In a study of tools provided by the advertising industry to opt out of behavioral advertising on websites, the researchers observed:

"Participants found many tools difficult to configure, and tools' default settings were often minimally protective. Ineffective communication, confusing interfaces, and a lack of feedback led

[26] Maritza Johnson, Serge Egelman, and Steven M. Bellovin. 2012. Facebook and privacy: it's complicated. In Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS '12). Association for Computing Machinery, New York, NY, USA, Article 9, 1–15. https://doi.org/10.1145/2335356.2335369

[27] Jennifer King, Airi Lampinen, and Alex Smolen. 2011. Privacy: is there an app for that? In Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS '11). Association for Computing Machinery, New York, NY, USA, Article 12, 1–20. https://doi.org/10.1145/2078827.2078843

[28] Alessandro Acquisti and Ralph Gross. "Imagined communities: Awareness, information sharing, and privacy on the Facebook." In *Privacy Enhancing Technologies: 6th International Workshop*, PET 2006, Cambridge, UK, June 28-30, 2006, Revised Selected Papers 6, pp. 36-58. Springer Berlin Heidelberg, 2006.

[29] Nathaniel S. Good and Aaron Krekelberg. 2003. "Usability and privacy: a study of Kazaa P2P file-sharing." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. Association for Computing Machinery, New York, NY, USA, 137–144. https://doi.org/10.1145/642611.642636

many participants to conclude that a tool was blocking [online behavioral advertising] when they had not properly configured it to do so. Without being familiar with many advertising companies and tracking technologies, it was difficult for participants to use the tools effectively."[30]

20. Incentives are also important when studying privacy tradeoffs. Privacy decisions are not made in a vacuum: that consumers engage with services that violate their privacy preferences is often an indictment of the lack of market choice rather than an indication that consumers are behaving hypocritically. Similarly, privacy is often not the only consideration: if the costs of protecting one's privacy are unreasonably high (e.g., time invested learning to correctly use privacy settings, monetary costs, abstaining from social life, etc.), many consumers will engage with privacy-violative services because they cannot afford the alternatives. For example, I value my free time, but that I still show up to work does not make me a hypocrite. Similarly, when faced with the choice between protecting their privacy or engaging with their peers online, many younger people will choose the latter, despite the known privacy risks. Many studies have shown that despite the known privacy risks, many young people continue to use social media due to the fear of missing out.[31,32,33]

21. Finally, many consumers simply do not believe they have agency when it comes to making online privacy decisions: because many believe that their privacy preferences will not be honored no matter the actions that they take, many choose to engage with privacy-violative services to extract benefits, believing that they will end up paying the privacy costs regardless. A 2015 consumer survey concluded the following:

---

[30] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. "Why Johnny can't opt out: a usability evaluation of tools to limit online behavioral advertising." In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 589-598. 2012.

[31] Vittoria Franchina, Mariek Vanden Abeele, Antonius J. Van Rooij, Gianluca Lo Coco, and Lieven De Marez. "Fear of missing out as a predictor of problematic social media use and phubbing behavior among Flemish adolescents." *International journal of environmental research and public health* 15, no. 10 (2018): 2319.

[32] Dmitri Rozgonjuk, Cornelia Sindermann, Jon D. Elhai, and Christian Montag. "Fear of Missing Out (FoMO) and social media's impact on daily-life and productivity at work: Do WhatsApp, Facebook, Instagram, and Snapchat Use Disorders mediate that association?." *Addictive Behaviors* 110 (2020): 106487.

[33] Ine Beyens, Eline Frison, and Steven Eggermont. "'I don't want to miss a thing': Adolescents' fear of missing out and its relationship to adolescents' social needs, Facebook use, and Facebook related stress." *Computers in Human Behavior* 64 (2016): 1-8.

"[A] majority of Americans are resigned to giving up their data—and that is why many appear to be engaging in tradeoffs. Resignation occurs when a person believes an undesirable outcome is inevitable and feels powerless to stop it. Rather than feeling able to make choices, Americans believe it is futile to manage what companies can learn about them. Our study reveals that more than half do not want to lose control over their information but also believe this loss of control has already happened."[34]

22. A study specifically on young people and the privacy paradox observed: *"Based on focus group interviews, we considered how young adults' attitudes about privacy can be reconciled with their online behavior. The "privacy paradox" suggests that young people claim to care about privacy while simultaneously providing a great deal of personal information through social media. Our interviews revealed that young adults do understand and care about the potential risks associated with disclosing information online and engage in at least some privacy-protective behaviors on social media. However, they feel that once information is shared, it is ultimately out of their control. They attribute this to the opaque practices of institutions, the technological affordances of social media, and the concept of networked privacy, which acknowledges that individuals exist in social contexts where others can and do violate their privacy."*[35]

23. Similarly, users continue to use apps that they find "creepy" due to a sense of learned helplessness: they do not believe that they have the power to control who receives their personal information when they participate in the digital economy.[36]

[34] Joseph Turow, Michael Hennessy, and Nora Draper. "The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation." *Available at SSRN 2820060* (2015).

[35] Eszter Hargittai, and Alice Marwick. ""What can I really do?" Explaining the privacy paradox with online apathy." *International journal of communication* 10 (2016): 21.

[36] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: perceptions of privacy and mobile app use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. Association for Computing Machinery, New York, NY, USA, 2347–2356. https://doi.org/10.1145/2556288.2557421

TOOLS FOR LIMITING COLLECTION & USE OF PERSONAL INFORMATION

24. **Privacy Policies.** Internet users have few tools to control their online privacy. Since the dawn of the Internet age, the primary framework for managing online privacy has been the "notice and consent" framework, whereby online services post privacy policies ("notice") and consumers can choose whether to engage with services based on their understanding of those policies ("consent"). Unfortunately, this framework is fundamentally detached from reality: decades of research have demonstrated that consumers do not read these privacy policies, do not understand what they mean (when they do read them), and worse, privacy policies often do not accurately describe their services' behaviors.

25. In one study, in which participants were asked to explicitly confirm that they read and agreed to a website's privacy policy, 80% clicked a box to affirm that they had done so despite not actually accessing or reading the policy.[37] This number likely represents a lower bound, given the presence of "demand characteristics" (i.e., participants were in a laboratory setting and therefore were likely to pay more attention to the instructions than they likely would have in the real world), as well as the fact that most online services do not present users with interstitial messages demanding that they read and agree to their privacy policies: most privacy policies are accessed through discreet links outside the user's field of focus. Another study found that privacy-concerned users were influenced by the mere presence of a privacy policy link, despite few reading the policies.[38] This suggests that the mere presence of a privacy policy erroneously signals "good" privacy practices.

26. Nonetheless, if users do opt to read privacy policies, it is often a significant time investment. In 2008, McDonald and Cranor showed that if users read the privacy policies for every website they accessed, they would need to spend up to 300 hours per year doing so

[37] Nili Steinfeld. "'I agree to the terms and conditions': (How) do users read privacy policies online? An eye-tracking experiment." *Computers in Human Behavior* 55 (2016): 992-1000.

[38] Jensen, Carlos, Colin Potts, and Christian Jensen. "Privacy practices of Internet users: Self-reports versus observed behavior." *International Journal of Human-Computer Studies* 63, no. 1-2 (2005): 203-227.

annually (based on average policy lengths, number of websites visited, and reading speeds).[39] Of course, their estimate is based on data from 2008 that showed the average Internet user visits around 1,500 unique websites annually; 15 years later, the number of websites has proliferated, as has the amount of time that consumers spend online, which suggests that the time investment to read and understand privacy policies has only increased.

27. It is also not clear that the time investment to read privacy policies is worthwhile for most consumers: several studies have shown that the privacy policies found on popular websites are written at the college level and therefore may not be understood by a significant proportion of the population (much less children).[40,41,42]

28. Even when policies are noticed, read, and understood, they generally do not explain a service's data practices in sufficient detail for consumers to make informed decisions. For example, despite CCPA and CalOPPA requiring that services post privacy policies, there are no requirements that force those services to name the specific third parties with whom they share data—they are only required to specify the broad categories of data recipients. Even though those third parties may have their own data practices that are documented in their own privacy policies, it is nearly impossible for consumers to inform themselves about those practices if they are unable to locate those additional privacy policies because they do not know the identities of the companies. Similarly, it is nearly impossible for consumers to understand the privacy practices of large companies that offer multiple services, as their privacy policies are often written in a manner that aggregates their practices across all of their offered services (e.g., Google's privacy policy[43] describes their data collection practices across all of their services and does not convey what data may be collected by Google Maps vs. Gmail vs. Docs vs. Search).

[39] Aleecia M. McDonald and Lorrie Faith Cranor. "The cost of reading privacy policies." *I/S: A Journal of Law and Policy for the Information Society*, 4 (2008): 543.

[40] Yuanxiang Li *et al*. "Online privacy policy of the thirty Dow Jones corporations: Compliance with FTC Fair Information Practice Principles and readability assessment." *Communications of the IIMA* 12.3 (2012): 5.

[41] Carlos Jensen and Colin Potts. "Privacy policies as decision-making tools: an evaluation of online privacy notices." *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* 2004.

[42] George R. Milne, Mary J. Culnan, and Henry Greene. "A longitudinal assessment of online privacy notice readability." *Journal of Public Policy & Marketing* 25.2 (2006): 238-249.

[43] https://policies.google.com/privacy?hl=en-US

29. **Blocking Cookies and Fingerprinting.** In addition to reading privacy policies, there are some technologies that consumers can use in futile attempts to better protect their privacy. "Cookies" are data that websites store in consumers' web browsers, which are then transmitted back to websites when visited in the future. This allows a website to recognize a user over time, without having to log in again (as well as allowing the website to "remember" other settings, such as a default language). Because cookies have been historically abused for invasive tracking and profiling,[44] modern web browser software allows users to delete stored cookies or to block cookies set by third-party trackers altogether.

30. However, deleting or blocking cookies is no longer an effective strategy, as tracking now occurs using other means that consumers cannot control.[45,46] For example, unique "fingerprints"—the aggregation of several data points to create a unique identifier—can be constructed based on seemingly-benign information that is automatically transmitted to online services without user consent: software versions (e.g., the web browser and operating system), language settings, time zones, screen resolution, battery levels, etc.[47,48] Even what fonts are installed on a computer, which are available to websites, can be used to uniquely identify a website visitor.[49] Apps on mobile devices have additional data points available for constructing unique fingerprints to identify their users, all without the use of cookies, and with few actions that users can take to prevent this from occurring. Perversely, whether a user has configured privacy settings away from the defaults is often used as a data point for further tracking (i.e., while some

[44] J. R. Mayer and J. C. Mitchell, "Third-Party Web Tracking: Policy and Technology," *2012 IEEE Symposium on Security and Privacy*, San Francisco, CA, USA, 2012, pp. 413-427, doi: 10.1109/SP.2012.47.

[45] N. Nikiforakis, A. Kapravelos, W. Joosen, C. Kruegel, F. Piessens and G. Vigna, "Cookieless Monster: Exploring the Ecosystem of Web-Based Device Fingerprinting," *2013 IEEE Symposium on Security and Privacy,* Berkeley, CA, USA, 2013, pp. 541-555, doi: 10.1109/SP.2013.43.

[46] R. Upathilake, Y. Li, and A. Matrawy, "A classification of web browser fingerprinting techniques," *2015 7th International Conference on New Technologies, Mobility and Security (NTMS)*, Paris, France, 2015, pp. 1-5, doi: 10.1109/NTMS.2015.7266460.

[47] Peter Eckersley. "How unique is your web browser?." In *Privacy Enhancing Technologies: 10th International Symposium, PETS 2010, Berlin, Germany, July 21-23, 2010. Proceedings 10*, pp. 1-18. Springer Berlin Heidelberg, 2010.

[48] https://amiunique.org/

[49] David Fifield and Serge Egelman. "Fingerprinting web users through font metrics." *Financial Cryptography and Data Security: 19th International Conference, FC 2015*, San Juan, Puerto Rico, January 26-30, 2015, Revised Selected Papers 19. Springer Berlin Heidelberg, 2015.

web browsers can transmit a user-configurable "do not track" signal to websites, many websites choose not to honor this and instead use it as another source of entropy to identify and track users).[50,51]

31. Every device connected to the Internet has an Internet Protocol (IP) address, which is used to route information to and from it. While IP addresses must be transmitted to send and receive data, they can also be used to track users over time. Since devices behind a firewall (e.g., a household WiFi router) will appear to the outside world to share the same IP address, the collection of IP addresses is often used as a way of performing "cross-device tracking," which allows data recipients to infer when the same individual has moved from using a mobile device to a desktop computer to a smart TV; it also allows data recipients to infer when multiple individuals reside within the same household. For example, Meta's privacy policy states that they collect "information about the network you connect your device to, including your IP address" to target advertisements and provide "business services" to unnamed partners.[52] There is little that consumers can do to prevent this, without substantially degrading their online experiences. Worse, there is no way for consumers to know when this type of tracking is even occurring.

32. **Machine-Readable Privacy Policies.** Over 20 years ago, due to the privacy concerns regarding cookies, online tracking, and the acknowledgement that natural language privacy policies are woefully inadequate, several proposals were put forth to create machine-readable privacy policies. The idea behind these proposals was that consumers could use an interface to save their privacy preferences within their web browsers (or other software under their control), websites could post machine-readable policies, and then web browsers could act on consumers' behalf to either alert them when encountering a website with a disagreeable privacy policy (determined by the browser's automatic parsing of a website's machine-readable policy),

[50] Geoffrey A. Fowler, "Think you're anonymous online? A third of popular websites are 'fingerprinting' you." *The Washington Post*, October 31, 2019. https://www.washingtonpost.com/technology/2019/10/31/think-youre-anonymous-online-third-popular-websites-are-fingerprinting-you/

[51] Michael Simon, "Apple is removing the Do Not Track toggle from Safari, but for a good reason." *Macworld*, February 6, 2019. https://www.macworld.com/article/232426/apple-safari-removing-do-not-track.html

[52] https://www.facebook.com/privacy/policy/

or take some other action (e.g., automatically negotiating a better policy, blocking cookies or other transmissions, etc.). One of these proposals became a web standard: the Platform for Privacy Preferences Project (P3P),[53] was a web standard developed by the World Wide Web Consortium. (I served on the standards committee as an invited expert.)

33. The P3P standard gained traction, with many industry stakeholders adopting it by posting "P3P policies" on their websites so that web browsers could automatically parse them and alert users when they encountered websites that violated those users' stated privacy preferences. Microsoft's Internet Explorer (IE) browser was the first major web browser to adopt P3P, and by default, IE would block third-party tracking cookies unless the website posted a P3P policy (and then would block third-party cookies in accordance with the user's stated privacy preferences). In response, many companies (e.g., Amazon, Facebook, and Google) posted P3P policies that did not actually describe their privacy practices, but nonetheless tricked the IE browser into accepting their tracking cookies, due to the presence of a valid P3P header.[54] One study of over 33,000 websites observed that more than one third were transmitting P3P policies that appeared to be designed to circumvent IE's cookie blocking (and did not accurately describe their sites' actual privacy practices).[55] (The same study found that many of these websites were certified participants in TRUSTe's[56] EU Safe Harbor industry self-regulation program, and concluded that such certified sites were no more likely to comply with the P3P standard than websites not certified.) Some of these P3P policies can still be found today when accessing the websites that include trackers from NetChoice members.[57] For example, as of March 28, 2023, Google Ads[58]

---

[53] https://en.wikipedia.org/wiki/P3P

[54] Lorrie Faith Cranor, "Necessary but not sufficient: Standardized mechanisms for privacy notice and choice." *J. on Telecomm. & High Tech. L.* 10 (2012): 273.

[55] Pedro Giovanni Leon, Lorrie Faith Cranor, Aleecia M. McDonald, and Robert McGuire. 2010. Token attempt: the misrepresentation of website privacy policies through the misuse of p3p compact policy tokens. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society (WPES '10).* Association for Computing Machinery, New York, NY, USA, 93–104. https://doi.org/10.1145/1866919.1866932

[56] TRUSTe is now known as "TrustArc."

[57] Lorrie Faith Cranor, "Internet Explorer privacy protections also being circumvented by Google, Facebook, and many more." *Technology Academics Policy,* February 18, 2021. https://www.techpolicy.com/Cranor_InternetExplorerPrivacyProtectionsBeing Circumvented-by-Google.aspx

[58] https://adservice.google.com/adsid/google/ui

transmits a P3P policy header, but the body of the policy is as follows:

CP="This is not a P3P policy! See g.co/p3phelp for more info."

34. Thus, I have come to the conclusion that voluntary online standards that aim to give consumers more control over their privacy are futile, as they are likely to be coopted.

SPECIAL CONCERNS REGARDING CHILDREN'S PRIVACY

35. This data monetization free-for-all is even more concerning when the data comes from children, who are unlikely to understand that this is happening, much less consent to it, but who could potentially face enormous impacts due to future usage of this data. This data may be used for manipulative marketing campaigns, but also may feed biased and unaccountable algorithms that use it to make decisions about a child's future, not to mention outright malicious uses of the data (e.g., non-custodial parents purchasing location data to geolocate a child).

36. In 2016 my research team decided to look at how well mobile apps directed at children appeared to be complying with COPPA, which has been in effect since 2000. We wrote bespoke instrumentation for the Android platform that allows us to run mobile apps and monitor exactly what personal information those apps access and with whom they share it.[59,60,61,62,63] We also used our instrumentation to determine whether transmissions containing personal information were performed securely and confidentially.

[59] P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. Wagner, and K. Beznosov. "Android permissions remystified: A field study on contextual integrity." In *Proceedings of the 24th USENIX Security Symposium (USENIX Security 15)*, pages 499–514, Washington, D.C., Aug. 2015. USENIX Association.

[60] P. Wijesekera, A. Baokar, L. Tsai, J. Reardon, S. Egelman, D. Wagner, and K. Beznosov. "The feasability of dynamically granted permissions: aligning mobile privacy with user preferences." In *Proceedings of the 2017 IEEE Symposium on Security and Privacy*, Oakland '17. IEEE Computer Society, 2017.

[61] P. Wijesekera, J. Reardon, I. Reyes, L. Tsai, J.-W. Chen, N. Good, D. Wagner, K. Beznosov, and S. Egelman. "Contextualizing privacy decisions for better prediction (and protection)." In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems,* CHI '18, pages 1–13, New York, NY, USA, 2018. Association for Computing Machinery.

[62] J. Reardon, A. Feal, P. Wijesekera, A. E. B. On, N. Vallina-Rodriguez, and S. Egelman. "50 Ways to Leak Your Data: An Exploration of Apps' Circumvention of the Android Permissions System." In *Proceedings of the 24th USENIX Security Symposium, USENIX Security '19,* Berkeley, CA, USA, 2019. USENIX Association.

[63] We wrote our tools for Google's Android platform only because it is open source: having the source code for the operating system allowed us to modify it for this purpose; at the time, we didn't look at Apple's iOS simply because we didn't have the source code to add the same level of instrumentation.

37. Starting in late 2016, we began downloading as many free apps in the "Designed for Families" (DFF) program as we could find, which ended up being just under 6,000 apps.[64] The DFF program is a section of the Play Store, Google's centralized Android app market, which is exclusively for apps that are directed to children. Mobile app developers must participate in the program when they upload their app and disclose to Google that it is directed at children. As part of the program, they must affirm to Google that their app is in compliance with COPPA. Our goal was to evaluate whether that appeared to be the case in practice.

38. Of the child-directed apps that we tested, more than half appeared to be violating COPPA in one way or another: 5% collected location or other contact information and 19% collected personal information without verifiable parental consent and shared them with third parties whose public disclosures indicated they would use them for prohibited purposes (e.g., behavioral advertising); 40% transmitted personal information insecurely. Separately, 39% appeared to be violating Google's platform policies (i.e., an example of industry self-regulation) surrounding the collection of persistent identifiers for advertising and analytics purposes.[65]

39. We also examined mobile apps that had been certified by the COPPA Safe Harbor programs, meaning that the app developer claimed to participate in a private FTC-approved compliance-certification program.[66] (We found it extraordinarily difficult to identify which mobile apps had actually been certified; none of the programs we contacted were willing to share lists of apps with us, and most of their websites did not provide this information.) Of the 237 apps we found that claimed to be Safe Harbor certified, 64% appeared to violate Google's policies on transmitting identifiers for advertising/analytics purposes, 33% transmitted personal information to prohibited third parties, and 32% transmitted personal information insecurely. We concluded that the apps that we examined, which claimed to be certified as COPPA-compliant by Safe Harbor programs, were no more likely to protect children's personal information than apps that

[64] Reyes *et al.*, *supra* note 3.
[65] *Ibid.*
[66] 16 C.F.R. § 312.11.

had not been certified by these programs.[67] (This result is consistent with prior research on adverse selection in industry self-regulatory certification programs.)[68]

40. Thus, based on this research, I have come to the conclusion that voluntary industry self-regulatory programs are ineffective, and do not lead to better outcomes for consumers.

41. Similarly, through this research, I identified several additional gaps in regulation (beyond the inadequacy of the Safe Harbor programs), that I recommended be fixed in my U.S. Senate testimony.[69] Particularly relevant here are COPPA's "internal operations" exemption[70] and "actual knowledge" standard.[71]

42. Generally, websites and other online services must obtain verifiable parental consent before disclosing children's personal information to third parties, unless it is to support the service's internal operations and is not used for any other purpose. However, from a technical standpoint, most internal operations do not strictly require the collection of persistent identifiers that can be used to track children's activities across different services. In fact, both major platforms provide guidelines on how software developers can perform these activities *without* collecting advertising identifiers or non-resettable device identifiers.[72,73] For example, by definition, "contextual advertising" involves showing consumers ads *without* using data previously collected about them, and therefore no personal information is needed to show contextual ads. To prevent one user from being shown the same ad repeatedly (known as "frequency capping"), a session-based or installation-based identifier should be used, such that the collected data cannot be used to track the user across other services.

---

[67] Reyes *et al.*, *supra* note 3.

[68] Benjamin Edelman. "Adverse selection in online" trust" certifications." In *Proceedings of the 11th International Conference on Electronic Commerce*, pp. 205-212. 2009.

[69] U.S. Congress. Hearing of the Subcommittee on Consumer Protection, Product Safety, and Data Security of the Committee on Commerce, Scient, and Transportation. Hearing on "Protecting Kids Online: Internet Privacy and Manipulative Marketing." Testimony of Serge Egelman, 2021. https://www.commerce.senate.gov/services/files/0DC78E9D-88B2-4D54-8F4A-AE7B4C7D0EF6

[70] 15 U.S.C. § 6501(4)(A).

[71] 15 U.S.C. § 6501(4)(B).

[72] Google, "Best Practices for Unique Identifiers." April 6, 2023. https://developer.android.com/training/articles/user-data-ids

[73] Apple, "User Privacy and Data Use." 2023. https://developer.apple.com/app-store/user-privacy-and-data-use/

43. Nonetheless, in the course of my research, I have noticed that many privacy policies associated with child-directed services use the phrase "internal operations," when describing the flow of children's personal information to third parties. In many of these cases, these third parties are advertisers whose public disclosures indicate that they may use the data for COPPA-prohibited purposes. Thus, I have concluded that for many developers, the phrase "internal operations" appears to be a shibboleth used to justify privacy-invasive practices.

44. Secondly, COPPA's "actual knowledge" standard, by which it must be shown that an individual within these third-party organizations knew that they received data from children, incentivizes data recipients to simply look the other way if and when they receive children's personal information, even when those third-party transmissions also include the names of the apps or websites that are transmitting them the data. Many of these data recipients are advertising and/or analytics companies that publicly advertise their abilities to target ads based on inferring the demographics of users of the services sending them data. Furthermore, there are many commercial services that purport to provide the target demographics of a given mobile app or a website, and thus determining whether or not a service is directed at children is readily ascertainable.

45. For example, ironSource is a targeted advertising company that we observed receiving personal information from child-directed apps.[74] Their privacy policy stated they did not knowingly receive personal information from children under 13, a point which was reiterated to my laboratory in a letter from their general counsel.[75] In my response, I pointed out that all developers wishing to use ironSource's services must provide a company name at sign-up, and we observed companies with the following names sending them personal information: "Arial & Babies," "Androbaby," "Babies Funny World," "BabyBus Kids Games," "For Little Kids," "GameForKids," and "KidsUnityApps." From these developer names provided to ironSource, the resulting data was likely coming from children. However, ironSource can deny actual knowledge,

[74] Reyes *et al.*, *supra* note 3.

[75] Serge Egelman, "We get letters." The AppCensus Blog, May 10, 2018. https://blog.appcensus.io/2018/05/10/we-get-letters/

so long as no human within the company looks at the data that they are soliciting from developers who use their services.

CALIFORNIA CHILDREN'S AGE-APPROPRIATE DESIGN CODE ACT

46. From my understanding of the California Age-Appropriate Design Code Act (AADC), I believe that several of the privacy problems I have identified in my research will be addressed, and that technology to comply with the AADC is already in widespread use (including by NetChoice's members).

47. I understand that the AADC requires that businesses that provide services, products or features likely to be accessed by children perform Data Privacy Impact Assessments, which includes identifying potential risks to children and how to mitigate those risks, as well as requiring that privacy notices be accessible and that user-configurable privacy settings are set to a high level (unless the business has a compelling reason not to). It also prohibits those same businesses from misleading their users about their policies and procedures, profiling child users (unless it is in the best interest of the child), collecting location data for purposes beyond determining AADC applicability, and sharing personal information with third parties for secondary purposes.

48. I understand that the AADC requires that DPIAs consider whether algorithms could result in harm to children. An algorithm is simply a sequence of operations: there is often an input, calculations are performed on that input, and then the results of those calculations are provided as output. Within the context of online services, algorithms are used for everything from recommending content to users to inferring a user's preferences and traits for purposes such as targeted advertising. There is no such thing as a "neutral" algorithm: algorithms are designed for specific purposes. One algorithm might be designed to show ads that maximize ad revenue, whereas another might be designed to optimize engagement through content recommendations; other algorithms might be used for more mundane tasks, such as sorting items chronologically or alphabetically. For example, in determining the tweets that appear in a user's feed (of the hundreds of millions sent per day), Twitter weighs factors such as the number of likes, retweets,

social relations, recency, perceived topic relevance, and use of embedded media, among other factors.[76]

49. While some algorithms might make objective decisions (e.g., correctly sorting a list of items by date), others are subjective and therefore less straightforward to audit for correctness (e.g., recommending content and choosing advertisements to display).[77] Algorithms are increasingly being used to make decisions about individuals that can have profound consequences, such as extending credit, housing, insurance, employment, or school admissions; in many cases there is little transparency or recourse surrounding these decisions, as they are made automatically and opaquely, and may also use incorrect or biased data.[78] Most adults do not understand if, when, and how these decisions are being made, children less so.

50. Algorithms that are optimized for increasing user engagement can also result in harm to consumers. For example, there was public outrage when the public learned that Facebook was using its content recommendation algorithms to intentionally cause emotional distress among its users. (Facebook researchers found that emotionally-charged posts were more likely to lead to user engagement; Facebook thus has an incentive to use its algorithms to prioritize showing users posts that are likely to evoke emotional responses.)[79]

51. The AADC regulates the use of so-called "dark patterns." Dark patterns are design choices that are used to "nudge" the user into making a decision that is advantageous to the business. For example, making it easier to sign up for a service than cancel it is a dark pattern, as is the use of artificial scarcity (e.g., countdown timers to convey a sense of urgency or "limited

[76] Josiah Hughes, "How the Twitter Algorithm Works [2023 Guide]." Hootsuite, December 14, 2022. https://blog.hootsuite.com/twitter-algorithm/

[77] Zeynep Tufekci, "Algorithmic Harms beyond Facebook and Google: Emergent Challenges of Computational Agency," Colorado Technology Law Journal 13, no. 2 (2015): 203-218.

[78] Danielle Keats Citron and Pasquale, Frank A., "The Scored Society: Due Process for Automated Predictions" (2014). Washington Law Review, Vol. 89, 2014, p. 1-, U of Maryland Legal Studies Research Paper No. 2014-8, Available at SSRN: https://ssrn.com/abstract=2376209

[79] Kashmir Hill, "Facebook Manipulated 689,003 Users' Emotions For Science." Forbes, June 28, 2014. https://www.forbes.com/sites/kashmirhill/2014/06/28/facebook-manipulated-689003-users-emotions-for-science/

time" offers).[80] Research shows that these techniques are prevalent in child-directed online services,[81] and that children are likely to be more susceptible to manipulations than adults.[82]

52. Given the problems with privacy policies and the lack of consumer understanding explained above, I believe the AADC addresses this issue by requiring the language to be understandable by target audiences (when their online services are likely to be accessed by children).

53. I understand that the Plaintiff in this case argues that they are unable to estimate the approximate ages of their users. However, the law does not appear to be proscriptive as to how services used by children should perform age estimation. Many such technologies exist, which all have benefits and drawbacks. For example, France's data protection agency, CNIL, published a guide to choosing appropriate technologies.[83] The report recommends that to balance user privacy with age estimation accuracy, services should not perform age estimation themselves, but instead should use independent third parties who can confidentially make guarantees to relying child-directed services without revealing additional personal information.

54. The report[84] also links to a prototype "implementation of an age-verification system that allows accessing restricted websites without sharing other personally identifiable data."[85] The recommended system is based on "zero-knowledge proofs," a concept in

[80] Sara Morrison, "Dark patterns, the tricks websites use to make you say yes, explained." Vox, April 1, 2021. https://www.vox.com/recode/22351108/dark-patterns-ui-web-design-privacy

[81] J. Radesky, A. Hiniker, C. McLaren, E. Akgun, A. Schaller, H. M. Weeks, S. Campbell, & A. N. Gearhardt (2022). "Prevalence and Characteristics of Manipulative Design in Mobile Applications Used by Children." JAMA network open, 5(6), e2217641. https://doi.org/10.1001/jamanetworkopen.2022.17641

[82] Dale Kunkel, Brian L. Wilcox, Joanne Cantor, Edward Palmer, Susan Linn, and Peter Dowrick. "Report of the APA task force on advertising and children." *Washington, DC: American Psychological Association* 30 (2004): 60.

[83] CNIL, "Online age verification: balancing privacy and the protection of minors." September 22, 2022. https://www.cnil.fr/en/online-age-verification-balancing-privacy-and-protection-minors

[84] *Ibid.*

[85] CNIL, "Demonstration of a privacy-preserving age verification process." June 23, 2022. https://linc.cnil.fr/demonstration-privacy-preserving-age-verification-process

cryptography that has been well-known for almost 40 years now,[86,87] which allows an entity to prove the validity of a statement without revealing additional details about that statement. As the CNIL report explains, this technology could easily be used to prove to relying online services that a user is above or below the age of 18 without revealing additional personal information about that user.

55. I understand that Plaintiff implies that it is not possible to reliably determine Internet users' geographic locations in order to determine which regulations apply. This is incorrect. There are many widely-used methods for identifying where in the world an Internet user is physically located. At the most basic level, public and private databases exist that map IP addresses—again, these are transmitted with every Internet connection—to physical locations. This technology is known as "geoIP" and is used by many Internet services to automatically determine where in the world their users come from. For example, MaxMind provides a free database for this purpose that claims 99.8% accuracy in determining a user's country and 80% accuracy for state/region.[88] Private databases, such as those maintained by several of NetChoice's members, are likely to be more accurate.

56. For example, Meta is already using geoIP data to automatically determine which Internet users should receive protections under CCPA/CPRA. Their documentation explains: "we will determine if a person is in California or not based on certain available signals which may include IP address or advertising ID, when those are available."[89] Google similarly automatically detects when users are located in California for the purposes of CCPA/CPRA compliance: "you can select the advertising partners that are eligible to receive bid requests for users Google determines are in California."[90]

[86] S. Goldwasser, S. Micali, and C. Rackoff. 1985. The knowledge complexity of interactive proof-systems. In Proceedings of the seventeenth annual ACM symposium on Theory of computing (STOC '85). Association for Computing Machinery, New York, NY, USA, 291–304. https://doi.org/10.1145/22145.22178

[87] U. Fiege, A. Fiat, and A. Shamir. 1987. Zero knowledge proofs of identity. In Proceedings of the nineteenth annual ACM symposium on Theory of computing (STOC '87). Association for Computing Machinery, New York, NY, USA, 210–217. https://doi.org/10.1145/28395.28419

[88] https://support.maxmind.com/hc/en-us/articles/4407630607131-Geolocation-Accuracy

[89] https://www.facebook.com/business/help/1151133471911882

[90] https://support.google.com/adsense/answer/9560818?hl=en

57. Both companies named above also allow their customers to specifically target ads to Internet users located within California. For example, here is a true and correct screenshot from Google Ads', https://ads.google.com/, accessed on March 28, 2023, targeting configuration interface, which allows advertisers to show ads to people specifically located within California:

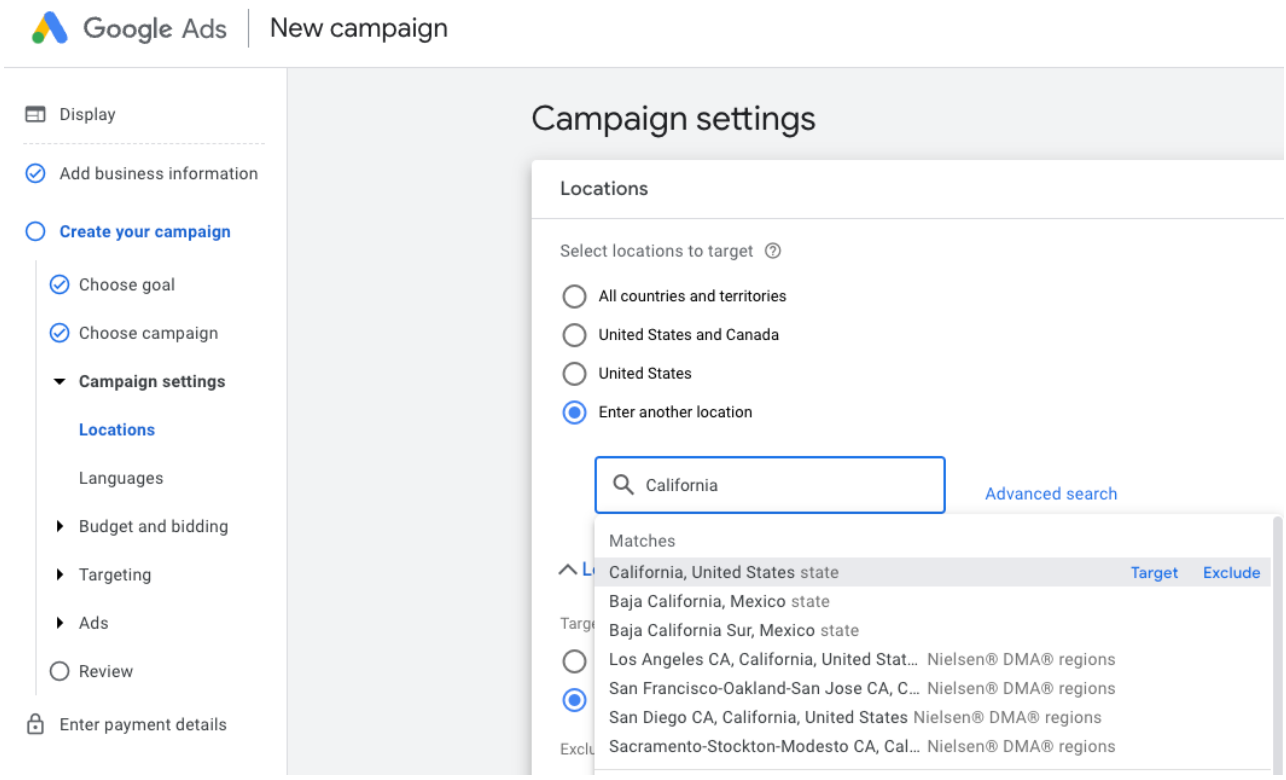

58. Below is a true and correct screenshot from Meta's Business Help Center website, https://www.facebook.com/business/help/365561350785642?id=176276233019487, accessed on March 28, 2023, describing how their customers can target ads to residents of specific states:

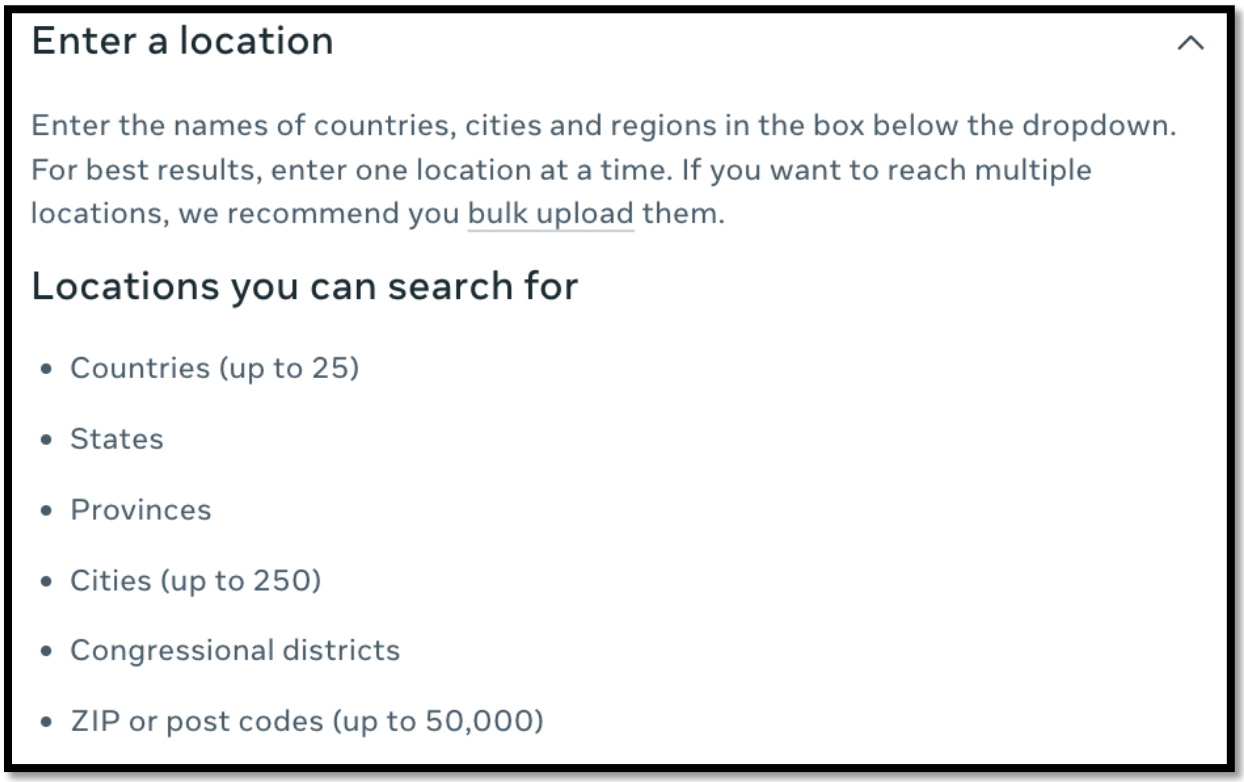

59. Yahoo!, another NetChoice member, also allows their customers to target ads to Internet users in specific states, even using California as an example. Below is a true and correct screenshot from Yahoo!'s Developer Network website, https://developer.yahoo.com/dsp/docs/lines/targeting-geos.html#target-geographic-areas, accessed on March 28, 2023:

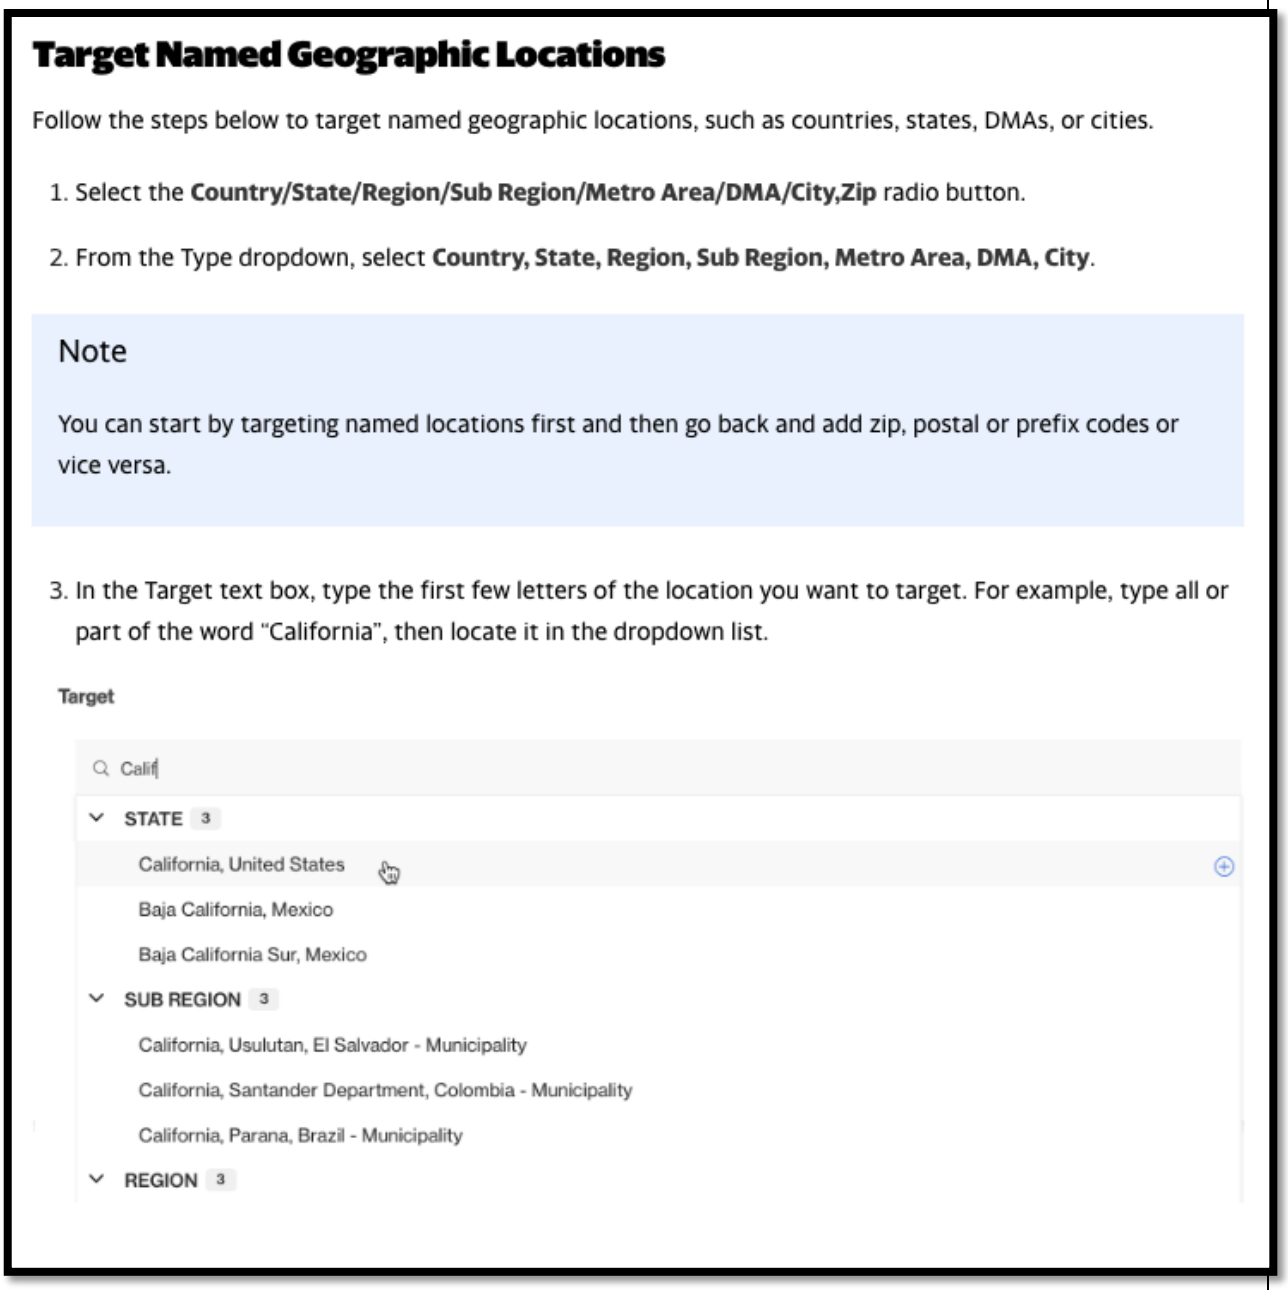

Target Named Geographic Locations

Follow the steps below to target named geographic locations, such as countries, states, DMAs, or cities.

1. Select the **Country/State/Region/Sub Region/Metro Area/DMA/City,Zip** radio button.

2. From the Type dropdown, select **Country, State, Region, Sub Region, Metro Area, DMA, City**.

Note

You can start by targeting named locations first and then go back and add zip, postal or prefix codes or vice versa.

3. In the Target text box, type the first few letters of the location you want to target. For example, type all or part of the word "California", then locate it in the dropdown list.

Target

Calif

STATE 3

California, United States

Baja California, Mexico

Baja California Sur, Mexico

SUB REGION 3

California, Usulutan, El Salvador - Municipality

California, Santander Department, Colombia - Municipality

California, Parana, Brazil - Municipality

REGION 3

60. In addition to geoIP lookups using available tools (many of which are already in use by NetChoice's members, and in many cases geolocating users to California for the purpose of determining CCPA/CPRA applicability), other methods exist for geolocating users, such as access to GPS hardware or other device sensors. For example, mobile apps running on the Android platform have access to Google's Geolocation services, which use nearby cellular towers and WiFi networks to determine the user's location, including providing the accuracy radius.[91]

[91] https://developers.google.com/maps/documentation/geolocation/overview

Apple's iOS platform offers similar functionality, which also make use of nearby cellular networks, WiFi hotspots, and other sensor data.[92]

61. Similarly, all of the major web browsers support functionality to geolocate their users,[93] which usually makes use of multiple methods, including using WiFi network information, GPS hardware, geoIP databases, and other data sources. Using these methods, the operators of online services have the ability to identify their users with street-level accuracy.

62. Thus, the technology to identify California consumers within a reasonable degree of accuracy already exists and is already in use by many of NetChoice's members.

OPINIONS

63. For the reasons I set out in this declaration, I believe that the AADC takes a reasonable approach to children's online safety. Based on my research and experience, consumers broadly believe that they are being protected by privacy laws that simply do not exist. Requiring online services to disclose policies in a manner accessible to their users and that they enforce those policies would go a long way towards helping consumers make informed decisions about their personal privacy.

64. The technologies needed to comply with the AADC's requirements already exist and are already in widespread use. Behaviors that the AADC prohibits have already been prohibited by major platforms. For example, child-directed Android apps are prohibited from collecting location data or performing behavioral advertising.[94]

65. As demonstrated above, consumers overwhelming want the practices this law requires for services that are likely to be accessed by children: limiting privacy-invasive tracking, providing safe defaults, and considering the harm to their users.

66. Finally, I believe that it is reasonable for services likely to be used by children to consider the harm they may have on their users. In fact, I think it's not unreasonable to ask that the offeror of any product or service consider the harm they might be causing to others.

[92] https://developer.apple.com/documentation/corelocation
[93] https://developer.mozilla.org/en-US/docs/Web/API/Geolocation_API
[94] https://support.google.com/googleplay/android-developer/answer/9893335?hl=en

I declare under penalty of perjury that the foregoing is true and correct. Executed on this 20th, day of April, 2023 in Berkeley, California.

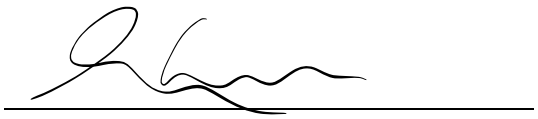

Serge Egelman, Ph.D.