Advertising in the Age of Data Protection

Background for a Proposed Risk-Utility Framework for Novel Advertising Solutions (v 1.0)

DISCUSSION DRAFT

v 1.0 April 2024

Public Comments Due May 26, 2024

Acknowledgements

Lead author: Stacey Gray (sgray@fpf.org), Senior Director, U.S. Policy, Future of Privacy Forum, with thanks to Dr. Rob van Eijk (FPF), Dr. Aaron Massey (FPF), Ryan Campbell (FPF US Policy Intern), and a growing number of advertising and data protection stakeholders.

Public Comment

This Discussion Draft is available for public comment until May 26, 2024. We welcome any reactions or thoughts from interested stakeholders, including suggested additions or changes that align with the goal of making this a practical resource for practitioners and policymakers.

- Send comments by email to info@fpf.org (subject line: "FPF Novel Advertising Framework Comments - [Name or Organization]").
- Please note that all changes remain at the sole discretion of the authors.
- A digital copy can be accessed by visiting www.fpf.org/adtech

About FPF

The Future of Privacy Forum (FPF) is a non-profit organization that serves as a catalyst for privacy leadership and scholarship, advancing principled data practices in support of emerging technologies. Learn more about FPF by visiting fpf.org.

Executive Summary

The goal of this Background Paper and its associated **Risk-Utility Framework (Appendix 1)** is to provide a **comprehensive rubric** for navigating the many tradeoffs inherent in the evolving digital advertising landscape and the technology it is built upon. The aim is for a policymaker, advertiser, or privacy leader at a company, to have tools to evaluate the impact of a given digital advertising proposal, system, or regulatory approach. The Proposed Framework includes a comprehensive list of factors that are relevant to assessing a novel advertising solution or system, including (1) advertising utility; (2) data protection and privacy; (3) competition and economic factors; and (4) other social impact.

In 2024, the digital advertising industry is in the midst of a sea change. Regulators are now more focused than ever on strengthening individual privacy rights and preventing many of the harms associated with the use of personal information in advertising. Meanwhile, large platforms such as Apple and Google have taken significant steps in to limit access to advertising-related data about their users, through efforts like App Tracking Transparency (ATT), Intelligent Tracking Prevention (ITP), and the deprecation of third party cookies. Each change has rippled effects throughout the economy.

In reaction to these regulatory and platform pressures, businesses are actively seeking new tools and solutions to maintain identity and addressability, or to provide greater privacy safeguards, ideally (in their view) doing so while sustaining as much business utility as possible. Many solutions involve privacy-enhancing technologies (PETs), while others involve a significant shift in business models, such as a return to contextual advertising, the use of solely first-party data, or a shift to client-side processing.

Among competing solutions, there is heated debate over the extent to which PETs-driven advertising solutions necessarily represent an improvement in privacy. Meanwhile, solutions that do create greater technical limitations on the transfer of data may eliminate or reduce many aspects of business utility, or be viewed as detrimental to other social values, such as market concentration. Access to data is also relevant for many other values, such as the protection of children, the mitigation of bias, and the ability to conduct social research on issues such as political influence. FPF's Proposed Framework aims to provide a means to assess all of these competing aspects and their associated tradeoffs.

We welcome **public comments until May 26, 2024,** including on how best to expand upon this Framework, how to demonstrate and clarify the inherent tradeoffs, and make it a useful resource for practitioners and other experts.

Table of Contents

1. Introduction - Expanding the Risk-Utility Debate

The goal of this **Proposed (Draft) Risk-Utility Framework** is to provide a common understanding and rubric for practitioners and policymakers to navigate the many aspects of data protection, utility, and other social values impacted by novel advertising strategies. In offering a comprehensive rubric, we aim to provide a practical tool, as well as a demonstration of the connections and tradeoffs inherent in the evolving digital advertising landscape and the technology it is built upon.

In focusing on the risks and utility of advertising, this Framework is an acknowledgement to the fact that digital advertising stands at a unique, and vitally important, intersection within the realm of data protection. Absent meaningful legal restrictions for most of its development in the last three decades, advertising has become increasingly data-driven, and incentivized large-scale data collection, profiling, and targeting methods that are often considered disproportionate, invasive, or inappropriate. At the same time, digital advertising drives a $600 billion global industry that impacts every sector of the economy, drives market competition, educates consumers about products, and allows organizations to communicate their value propositions. Digital advertising has also subsidized free and lower cost digital content and services, a societal shift not without its criticism, but that has nonetheless transformed access to information relied upon by people around the world.

In 2024, the **digital advertising industry is in the midst of a sea change.** Around the world, privacy regulators have become far more critical of mainstream advertising business models. Both lawmakers and enforcers of existing laws are now more focused on strengthening individual privacy rights and specifically preventing many of the harms associated with the use of personal information in advertising.[1] In particular, as we describe below, there is a growing awareness of the ways in which personal information collected for advertising often ends up being used for a **wide range of secondary purposes,** both socially beneficial and potentially harmful or surveillance-oriented.[2]

Meanwhile, large platforms such as Apple and Google have taken significant steps in recent years to limit access to advertising-related data about their users, through efforts like App Tracking Transparency (ATT), Intelligent Tracking Prevention (ITP), and an ongoing process to deprecate third party cookies in Google Chrome.[3] Each change has rippled effects throughout the economy, changing the way advertisers do business and often impacting social values.

---

[1] *See, e.g.*, *FTC Cracks Down on Mass Data Collectors: A Closer Look at Avast, X-Mode, and InMarket*, Federal Trade Commission (Mar. 2024), https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/03/ftc-cracks-down-mass-data-collectors-closer-look-avast-x-mode-inmarket; *CCPA Enforcement*, California Attorney General's Office (last visited Mar. 18, 2024), https://oag.ca.gov/privacy/ccpa/enforcement.

[2] *See, e.g.*, https://www.nytimes.com/2024/01/25/us/politics/nsa-internet-privacy-warrant.html

[3] *See, e.g.*, *App Tracking Transparency*, Apple, https://developer.apple.com/documentation/apptrackingtransparency; *Safari Privacy Overview*, Apple (Nov. 2019), https://www.apple.com/safari/docs/Safari_White_Paper_Nov_2019.pdf (explaining Intelligent Tracking Prevention); Anthony Chavez, *The next step toward phasing out third-party cookies in Chrome*, Google (Dec. 2023), https://blog.google/products/chrome/privacy-sandbox-tracking-protection/.

In reaction to these regulatory and platform pressures, businesses are **actively seeking new tools and solutions** to maintain identity and addressability, or to provide greater privacy safeguards, ideally (in their view) doing so while sustaining as much business utility as possible.[4] Many solutions involve privacy-enhancing technologies (PETs), while others involve a significant shift in business models, such as a return to contextual advertising, the use of solely first-party data, or a shift to client-side processing (models such as the Apple SKAdNetwork API, and Google's Privacy Sandbox). Even advertising experts may not know where to begin, and regulators face a steep challenge in crafting well-informed policy.[5]

Among competing solutions, there is heated debate over the extent to which PETs-driven advertising solutions **necessarily represent an improvement in privacy**. For example, ad tech proposals claiming to be reliant on PETs have attracted criticism related to perceived "privacy washing" and "PETs washing," referring to the superficial application of privacy branding, while in reality failing to meaningfully restrict or limit the collection and use of personal information.[6] In contrast, solutions that do create greater technical limitations on the transfer of data may eliminate or reduce many aspects of business utility, or be viewed as detrimental to other social values, including market concentration.[7] Meanwhile, we observe that access to data is relevant for many other values, such as the protection of children, the mitigation of bias, and the ability to conduct social research on issues such as political influence.

FPF's Proposed Framework aims to provide a means to assess all of these competing aspects and their associated tradeoffs. We **do not assign values to each aspect of utility, risk, or social impact,** but rather aim to holistically identify the metrics for assessing a given advertising solution.

---

[4] E.g. *Identity Solutions Guidance Draft*, *IAB Tech Lab* (October 2023), available at https://iabtechlab.com/wp-content/uploads/2023/10/Identity-Solutions-Guidance-Draft-CLEAN-for-Public-Comment.pdf; Todd Parsons, *The Multifaceted Solution for Addressability: What Marketers Need to Know*, Criteo (July 2023), https://www.criteo.com/blog/the-multifaceted-solution-for-addressability-what-marketers-need-to-know.

[5] *See Verve Group, Identity Decoded: A Guide to Navigate the Identity Discussion* 8 (2023) (Noting that "A majority (57%) of those who haven't adopted any cookieless solution say it's because they don't know how the solutions work, while others say there are too many choices, making it hard to know where to begin").

[6] E.g., Lee McGuian et al, *The after party: Cynical resignation in Adtech's pivot to privacy* (2023), *available at* https://journals.sagepub.com/doi/10.1177/20539517231203665

[7] See, e.g., Lukasz Olejnik, *Data Protection assessment of Privacy Sandbox's Protected Audience API*, LukasOlejnik.com (Jan. 2024), https://blog.lukaszolejnik.com/data-protection-assessment-of-privacy-sandboxs-protected-audience-api/; Report, *CMA Q3 2023 update report on implementation of the Privacy Sandbox commitments*, UK Competition and Markets Authority (Oct. 2023), https://assets.publishing.service.gov.uk/media/653a58e2e6c968000daa9b94/__Q3_2023_update_report_____.pdf.

2. Novel Solutions and Risk Mitigation Strategies

In recent years, a range of novel solutions have begun to emerge in response to regulatory changes and platform-level changes, including the ongoing deprecation of third-party cookies. This Proposed (Draft) Risk-Utility Framework is intended to be applicable to any given advertising technology, system, solution, or business strategy.

Specifically, we are most interested in technical solutions that seek to replace elements of functionality being lost due to the deprecation of third-party cookies, many of which involve privacy-enhancing technologies (PETs). These most prominently include **alternative IDs** for identity maintenance, **data clean rooms** (DCRs), as well as uniquely PETs-driven solutions such as **on-device processing** or client-side processing.

In addition, we include in this section a brief discussion of strategies that are not necessarily novel at a technical level, but represent a significant shift in strategic trends, including: **contextual advertising**, reliance on **first party data**, and compliance with decentralized **individual opt-out mechanisms.** Each novel solution or business strategy serves to satisfy different business and/or regulatory goals, while creating its own set of limitations and trade-offs, summarized below. In all cases, the specific impact on data protection, utility, and other social values (Framework, Appendix I) will depend on the details of implementation.

- ***Alternative IDs***

In response to loss of third party cookies, many advertising providers are relying on "alternative IDs," such as the Trade Desk's Unified ID (UID) 2.0, RampID, ID5, and Panorama ID, to conduct various forms of reporting, measurement, attribution, profiling, or targeting activities that rely on identifying unique browsers or devices over time and/or across unaffiliated sites or services.

In general, alternate IDs have two forms: **deterministic IDs** that are derived from information that a user provides, such as a hashed email address or phone number, and **probabilistic IDs** that are derived from passive or active device signals. Deterministic IDs, because they rely on user-provided information, may lend themselves to a greater opportunity to provide notice and/or consent. At the same time, they may be more limited in the scope of their application. For example, a survey found that 70% of publishers have no greater than 30% logged-in users, and that 44% of those publishers have been able to authenticate only 10% of those users.[8] Accordingly, these forms of IDs may be most valuable to advertisers that seek a degree of certainty about the identity of the prospective ad recipients, as with highly personalized ad campaigns to known customers.

In contrast, probabilistic identification encompasses a range of methods that rely on passive or active signals from the device or browser (such as IP address, operating system details) to predict, within a range of useful probability, unique browsers or devices over time and across

[8] ID5, *State of Digital Identity Report 2021,* https://id5.io/resources/state-of-digital-identity-report-2021/.

sites. Given the inherently imprecise nature of such models, probabilistic IDs can have varied utility – for example, offering certain short-term measurement capabilities, while offering less utility for campaigns requiring high degrees of accuracy. At the same time, such details can determine the impact on privacy and data protection: whether the signals are passive or active; derived from signals that are difficult or expensive to change, such as hardware information; the extent to which predictions are accurate; and the length to which identification can persist over time.

- ***Data Clean Rooms***

A growing number of advertising solutions involve the use of intermediaries, or Data Clean Rooms (DCRs), that allow for the matching or linkage of personal information across datasets owned by two or more unaffiliated organizations, usually with technical or contractual measures in place to ensure the data is not used for unrelated purposes.

Depending on the technical and administrative configuration, the privacy implications of a business arrangement marketed as a "Data Clean Room" can vary significantly. In older or more traditional marketing co-op arrangements, the protections for pooled personal information could be limited to solely administrative or contractual agreements between participating companies. However, ad tech providers are increasingly developing DCR models that involve greater technical safeguards intended to limit companies' ability to access, share, or re-use the underlying pool of personal information.

In many cases, advertising-related use cases for DCRs involve either enhancing the data that a company brings (e.g., as a custom audience), or facilitating some form of 1:1 matching between disparate datasets (e.,g., in order to deliver or measure the effectiveness of ads between systems that use different forms of personal identifiers). These kinds of uses of DCRs may still trigger data protection legal requirements, or lead to the delivery of personalized content, even if the use of PETs limits the nature and volume of data that participating companies receive. In more privacy-protective arrangements, a DCR or other intermediary might provide only aggregated insights that do not allow for inferences or actions related to individual devices or browsers. Such arrangements may involve the use of privacy-enhancing technologies such as differential privacy (the addition of noise and limited queries) or secure multi-party computation.

- ***"On Device" Processing (Topics API, SKAd)***

A significant shift in recent years, particularly among leading web browsers (e.g., Apple Safari, Google Chrome, Microsoft Edge) and other large platforms (e.g., Meta) has been the introduction of a range of solutions involving "on device" data processing. In general, these kinds of solutions are responsive to business or other incentives to derive insights from information that would otherwise either 1) be kept fully private; or 2) was previously made available, but is now limited or restricted in response to regulatory or platform pressures, causing platforms to seek alternative solutions for advertising needs.

In the first instance, platforms may seek to apply privacy-enhancing technologies to gain access to aggregate insights or machine learning models derived from data that would

otherwise be kept fully private. This is the case, for example, for entities that apply a combination of **federated learning and differential privacy** to train models on the basis of data from communications or other private datasets.[9] Federated learning relies on a decentralized approach to training machine learning models, in which processing occurs locally on many different users' devices, and updates to the central model are sent using privacy-protective methods. The resulting model can benefit from insights and learning derived from many different devices, without sharing or centrally processing the raw data from each user's device.

In other cases, a range of advertising solutions are emerging that offer an **alternative means** to fulfill advertising use cases that may have previously relied on the unrestricted transfer and sale of personal data. This includes, for example, Google's Privacy Sandbox solutions, Apple's SkadNetwork, and others. Such solutions are designed to allow third parties to have access to either no personal information, or enough information to fulfill advertising use cases, but without having access to a significant amount of information about individuals' browsing behavior over time. In reality, such solutions may involve a mix of on-device and off-device processing, with the central goal of offering advertising solutions while "reduc[ing] cross-site and cross-app tracking."[10] [11]

For example, in 2018, Apple introduced SKadNetwork, an API-based solution for ad measurement and attribution that allows advertisers to generate summary reports on ad campaigns conducted in the App Store and Safari, in the absence of third-party identifiers.[12] Meanwhile, Google's Privacy Sandbox involves a range of solutions being developed to offer advertisers the ability to target audiences and measure ad campaigns in Google Chrome without the use of third-party cookies, which are planned for deprecation.[13] This includes tools for generating measurement and attribution reports, targeting custom audiences (Protected Audience API), and targeting content to interest-based cohorts of browsers based on individual browsing history (Topics API). Other browsers are following suit, with Microsoft Edge recently introducing an Ad Selection API.[14]

Overall, these kinds of advertising solutions involve a significant shift towards greater client-side or on-device processing, targeting, and segmentation, including through the use of privacy-enhancing technologies. This shift is intended to alleviate many of the significant privacy and data protection concerns associated with the unrestricted onward transfers of personal information associated with open web advertising exchanges and real-time bidding

---

[9] *See Applying federated learning to protect data on mobile devices*, Meta (June 2022), https://engineering.fb.com/2022/06/14/production-engineering/federated-learning-differential-privacy/.

[10] Privacy Sandbox overview https://developers.google.com/privacy-sandbox/overview

[11] *See, e.g.*, *Apple Advertising and Privacy*, Apple (Apr. 2023), https://www.apple.com/legal/privacy/data/en/apple-advertising/ ("[w]e may also use local, on-device processing to select which ad to display, using information stored on your device, such as the apps you frequently open")

[12] SKAd Network, Apple Developer Documentation (last visited Mar. 28, 2024), https://developer.apple.com/documentation/storekit/skadnetwork/.

[13]

[14] Microsoft Advertising Blog, *The future of private advertising* (Mar. 5, 2024), https://about.ads.microsoft.com/en-us/blog/post/march-2024/the-future-of-private-advertising

(RTB). At the same time, however, they may still involve a degree of first-party data collection, processing, and individual profiling that remains concerning to privacy advocates, and may come at the cost of other values, such as transparency, accountability, advertising utility, or market concentration. Each solution should be carefully evaluated for its data protection and privacy impact, as well as impact on other values.

- ***Contextual Advertising (e.g. Search, AdSense)***

In response to regulatory pressures, many advertisers are returning to a focus on **contextual advertising**, which refers to the placement of advertisements that are tailored to the content of the website or app in which they are viewed, rather than to the individual viewer.[15] For example, contextual advertising might include search engine ads (e.g., a product or business relevant to search results in a single session), or ads based on key phrases on a website or app (e.g. an ad for reading glasses on a website for book reviews). More broadly, contextual advertising includes ads that are relevant to the likely audience of a publication at a particular point in time (e.g., an ad for luxury vacations placed in Forbes.com on the basis that its readers are likely to be affluent, or an ad for candy placed in advance of Valentine's Day).

Notably, in most cases contextual advertisements are still delivered programmatically, i.e., through a data-driven ad network, and therefore may rely on personal information for aspects such as ad delivery and reporting (ADR), or brand quality control, discussed in more detail below. Contextual advertising may also not be sufficient for certain aspects of advertising utility, such as the ability to reach small or niche audiences. At the same time, it's worth noting that contextual advertising can be highly informed by inferences drawn from large datasets of personal information, for example relying on less obvious correlations between human behavior and external factors, such as weather.[16] Depending on the availability of data to be used for these additional purposes (making correlations, and measuring and reporting on ad effectiveness), the actual utility of contextual advertising solutions can vary widely.

- ***Reliance on First Party Data (e.g., Social Media, Large Publishers)***

As large platforms increasingly restrict access to information about their users (e.g. Apple's ATT), and in light of regulations that primarily restrict "sharing" and "sale," many advertising providers are witnessing a significant shift towards greater reliance on "first party data."

First party data refers to personal information that a business collects directly from their own customers through direct engagement. This can include **known or identifiable customers**, e.g. customers that purchase goods and services, subscribe to content, or hold authenticated accounts with a business, such as for social media or streaming services. It can also include **less identifiable customers**, such as not-logged-in visitors to a company's websites.

In general, companies with either very large or high-value customer bases tend to be most successful in conducting personalized marketing and advertising for their own customers, or

---

[15] Peter Wallace, *Contextual: don't call it a comeback*, New Digital Age, https://newdigitalage.co/advertising/contextual-dont-call-it-a-comeback/.
[16] *What is contextual advertising? Everything you need to know*, The Weather Company (Sept. 2021), https://www.weathercompany.com/blog/what-is-contextual-advertising-everything-you-need-to-know/.

facilitating the placement of advertising from others to those customers, i.e., functioning as a Demand Side Platform (DSP). This would include, for example, social media platforms, premium publishers and their affiliates, and any other consumer-facing company with sufficiently valuable first-party data.[17]

Depending on the scope of the first-party relationships, the privacy implications of first-party data collection and use can vary. It's worth also noting that first party data is frequently "enriched" or appended with additional information to enhance the value of audiences and add to the information that is known about individuals, for example by purchasing data from data brokers and public sources.

- ***Individual Opt-Outs through Decentralized Signaling (e.g., Global Privacy Control) or Authorized Agents (e.g., Permission Slip)***

Another strategy, which has been largely responsive to regulatory efforts in the United States, has been compliance with various forms of individual "opt out" mechanisms. Although operationalizing individual preferences is not new to the advertising industry,[18] recent legal developments in the United States are requiring businesses to honor decentralized, browser-based opt-out preference signals received from individuals.[19] A significant benefit of such signals for individuals is that it allows them to communicate a preference to opt out of particular processing activities at scale, i.e., to all websites or apps with whom they interact, rather than approaching each company individually. At the same time, communication of such preferences does not necessarily prevent data collection at the source, and there are ongoing serious challenges related to how businesses ought to comply, and how compliance can be measured or audited.

In a similar way, there are a growing number of organizations that offer an "opt out at scale" service to users, including by acting as an "authorized agent" under relevant laws and communicating an individual's desire to exercise various rights (to opt out of personalized advertising, or to delete their data) to a large number of companies at once. While similarly offering a degree of post-hoc individual control over data processing and a way for individuals to exercise certain data protection rights at scale, these tools do not prevent initial data collection or necessarily guarantee legal compliance.

---

[17] *See* Eric Benjamin Seufert, *Everything is an Ad Network* (Nov. 8, 2021) https://mobiledevmemo.com/everything-is-an-ad-network/.

[18] *Your Ad Choices Give You Control*, Digital Advertising Alliance, https://youradchoices.com/control.

[19] Currently, eight state privacy laws require a universal opt out mechanism (UOOM): California, Colorado, Connecticut, Delaware, Montana, New Jersey, Oregon, and Texas. *See* Samuel Adams and Stacey Gray, *Survey of Current Universal Opt-Out Mechanisms*, Future of Privacy Forum (Oct. 2023), https://fpf.org/blog/survey-of-current-universal-opt-out-mechanisms/.

Table 1. Examples of Novel Advertising Solutions and Strategies

Novel Solution or Strategy	Examples
1. Alternative IDs	Trade Desk's Unified ID (UID) 2.0, RampID, ID5
2. Data Clean Rooms	PAIR
3. Client-side processing through APIs	PS, SKAd
4. On-device segmentation (cohorts) based on browsing behavior and federated learning	Topics API, MS Ad Selection
5. Contextual advertising	Varies* (*Many leading providers, e.g. Criteo, Google, The Trade Desk, and others, offer contextual advertising solutions)
6. Reliance on First Party Data	Social media (Meta, X, Reddit); large publishers (NYTimes, WaPo)
7. Individual Opt-Outs through Decentralized Signaling or Authorized Agents	LAT, DNT, GPC, Permission Slip

3. Utility: How and Why Advertisers Use Data (Framework Part A)

In order to begin evaluating the many emerging proposals (including, but not limited to, those discussed above), it is helpful to first have a basic understanding of the current role that personal information plays in various aspects of utility, and why advertisers consider access to personal information so vital.

Digital advertising operates as a complex data ecosystem. Programmatic advertising in the open web relies on a real-time exchange marketplace for buying and selling advertising inventory on the basis of personal information. Information about individual browsers or devices is derived from a user's visits and interactions across a large number of websites, apps, and platforms, typically linked together through **persistent, cross-site identifiers** such as cookie IDs, mobile advertising identifiers (MAIDs), or a growing number of "alternate IDs" that seek to offer a replacement for third-party cookies (described above). Information related to online and mobile devices is often tied to "offline" information as well, such as information on retail purchases or visits, including information on an individual's real identity.

The following areas describe the use of personal information in different aspects of programmatic advertising.

- ***Advertising Delivery and Reporting (ADR)***

A certain amount of basic information is required to facilitate the delivery of advertisements within an ad campaign (e.g., in the right formats, with a rotation of creative content, over a particular period of time) and to provide a report to the advertiser and bill them accordingly. In industry terms, this is referred to as Ad Delivery and Reporting (ADR).

For the placement of ads (Ad Delivery), information from devices and browsers is used to allow advertisers to specify details for their campaign, such as advertising for a particular type of browser, device, time of day, or location, as well as capping frequency (number of times an ad is shown to the same device or browser), and sequencing of advertising creatives.[20] Meanwhile, reports provided to an advertiser can vary and include, but are not limited to, "statistical reporting, traffic analysis, analytics . . . [and] billing."[21] Typically, these activities are facilitated with the use of persistent, cross-site identifiers or other forms of personal information, for example to deliver ads across multiple devices owned by user(s) in the same household, or to deliver ads in a particular order. However, a growing body of research may develop ways to facilitate certain ADR activities, such as frequency capping, without the use of personal information that allows for tracking over time.[22]

- ***Measurement and Attribution***

A range of information, including personal information, is also used in measuring campaign effectiveness by linking advertising campaigns to **specific business goals** (e.g., overall brand lift or sales growth) or to **individual actions** (clicks, conversions, downloads, or purchases), or both. In the past, a simple metric such as "click through rate" (CTR) may have informed the effectiveness of an advertising campaign. In contrast, modern campaigns typically rely on other forms of measurement, such as app downloads, purchases, or conversion rates (the rate at which viewers download a service or sign up for an account, becoming known customers).

More sophisticated attribution models often rely on the processing of personal information over time, in order to assess factors such as dynamic attribution, or attributing an action based on multiple advertisements or "touchpoints" with an individual, potentially across different devices or platforms that must necessarily be linked to the same individual.[23] Each of these factors influence payment models for advertisers, as well as decisionmaking on how best to achieve return on investment (ROI). In general, the use of data for measurement and attribution does not create the same privacy concerns associated with profiling and

---

[20] *See* definition of "Ad Delivery and Reporting" in the Network Advertising Initiative's (NAI) Digital Advertising Glossary, https://thenai.org/about-online-advertising/glossary/.
[21] *Id.*
[22] See Jonathan Mayer and Arvind Narayanan, *Tracking Not Required: Frequency Capping* (April 3, 2012), http://webpolicy.org/2012/04/23/tracking-not-required-frequency-capping/.
[23] *See* Stacey Gray and Jules Polonetsky, *Cross Device: Understanding the State of State Management, Future of Privacy Forum* (Nov. 2015), https://fpf.org/wp-content/uploads/2015/11/FPF_FTC_CrossDevice_F_20pg-3.pdf.

behavioral advertising. However, it's worth noting that absent limitations on storage, transfer/sale, and other secondary uses, information collected for measurement and attribution may still lead to many of the data protection impacts arising from downstream uses described below (Part 4).

- ***Re-Targeting (Remarketing)***

Retargeting, or remarketing, is an industry term for advertising to an individual that has previously visited a given website or otherwise interacted with a particular business. In a typical example, a retailer may wish to reach unknown individuals that visited their site, browsed or placed items into a shopping cart, and perhaps left before completing a purchase. Currently, this can be facilitated through the use of pixel tracking and third-party cookies, by which a website owner places code (Javascript or a 1x1 pixel) on a portion of their website that triggers the placement of a third party cookie containing a unique identifier. Through the use of a third party ad network, the website owner can then reach that visitor in another unaffiliated context, such as through ads placed on other websites or platforms. This common scenario is frequently attributed with the experience of ads that "follow" a person around.

Although relatively simple in principle, pixel-based tracking and retargeting can raise a number of privacy and data protection risks. For example, data collection can be considered particularly high-risk if it occurs on websites that may reveal sensitive information, such as a retail site for prescription drugs. Ad networks or other platforms that facilitate retargeting for large numbers of publishers may be in a position to collect large amounts of information from across many publishers that individuals would not expect.

- ***Behavioral Advertising***

Behavioral advertising, online behavioral advertising (OBA), interest based advertising (IBA), or simply "targeted advertising," are all terms that refer to a range of activities that facilitate advertising to individuals on the basis of their inferred interests based on online activities over time and across unaffiliated sites or platforms.[24] In addition to the basic cross-site advertising example described above (reaching a website visitor on another site through retargeting), behavioral advertising typically also includes advertising to "audience segments," or groups of browsers or devices that have been segmented into perceived or inferred interest categories.

Although interests can be inferred using solely first-party data (for example, a user's activities within a single social media platform), there is a broad industry of advertising intermediaries dedicated to inferring interest categories on the basis of an individual's activities across unaffiliated platforms. Audience segments can often be appended or enhanced with "offline data," or information about an individual's retail purchases, in-person visits to locations, or information from public records.

[24] For a leading industry definition of "interest based advertising," see, e.g., https://thenai.org/glossary/interest-based-advertising-iba/.

4. Data Protection: The Impact of Digital Advertising on Individual Privacy and Processing of Information (Framework Part B)

The impact of an existing or novel advertising solution on privacy and data protection can vary tremendously – from relatively benign, to highly invasive or capable of causing significant harm. In a further complication, the concept of privacy itself, and what constitutes privacy and data protection harms, have different, culturally contingent meanings, with different aspects holding different values for different stakeholders.

Even within this complex debate, however, privacy harms related to digital advertising have been well-documented. In addition to the concrete, tangible harms that can occur through misuse of information, there is a widespread lack of trust, as in a 2019 survey finding that 55% of Americans object to companies "spying" on them for the purposes of advertising.[25]

In **Appendix I,** we aim to separate out the many elements of mainstream privacy and data protection principles that underpin our current legal and policy norms, as well as regulations such as the EU's GDPR and ePrivacy Directive, as well as U.S. state privacy laws like the CCPA.[26] This section is intended to allow an evaluator to comprehensively examine different aspects of privacy and data protection principles and bring their own values to the resulting analysis. While Appendix I contains a comprehensive checklist, we highlight a few of the most significant aspects here:

- ***Lawfulness, Fairness, and Transparency***

As a first step in any assessment of privacy and data protection, an evaluator should consider the basics of whether, and how, personal information relating to individuals is collected. Both US and EU laws consider "fairness," a fundamental aspect of data processing that involves collecting data in an open and transparent manner that is not misleading or unexpected. In the context of advertising, this may mean asking whether information about how the system operates is accessible, intuitive, easy to understand, and collected pursuant to consent or a known functionality or direct relationship with the controller.

- ***Purpose Limitation (Secondary Uses)***

A common and significant criticism of modern digital advertising practices is that information collected for advertising can frequently end up being used for additional purposes that are

---

[25] Brooke Auxier, *54% of Americans Say Social Media Companies Shouldn't Allow Any Political Ads*, Pew Research Center (Sept. 24, 2020), https://www.pewresearch.org/short-reads/2020/09/24/54-of-americans-say-social-media-companies-shouldnt-allow-any-political-ads/.

[26] *See, e.g.*, *Quick Guide to the Principles of Data Protection*, Ireland Data Protection Commission, https://www.dataprotection.ie/sites/default/files/uploads/2019-11/Guidance%20on%20the%20Principles%20of%20Data%20Protection_Oct19.pdf; *A Guide to the Data Collection Principles*, UK Information Commissioner's Office, https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-protection-principles/a-guide-to-the-data-protection-principles/.

unexpected, surprising, harmful, or otherwise incompatible with the original purpose for which the data was collected. For example, information derived from advertising exchanges is frequently purchased by government agencies for national security purposes, as well as routinely accessed by federal, state, and local law enforcement.[27]

- ***Data Minimization and Storage Limitation***

As in all industries with data protection implications, digital advertising tools can be evaluated on the basis of their level of data minimization, which refers to collecting and processing the minimum amount of personal data required to meet the specific needs of the processing. Data minimization also incorporates the idea that information can be altered, redacted, or transformed to reduce its risk of identifiability (e.g., pseudonymization), and that there should be clear limitations on any long-term storage of information. Data minimization and purpose limitation are often considered together since the ability to consider the scope of collection and retention is limited by the purpose or purposes for which the data is likely to be used.

- ***Integrity and Confidentiality (Security)***

The digital advertising system also has distinct vulnerabilities which may be exploited by bad actors, threatening both individual privacy and national security. For example, in the real-time bidding (RTB) process through which advertisements are facilitated, information or identifiers such as location data or time-stamps can be used to identify specific individuals. Reports by both the EU and US signalled that this information, such as individuals within given advertising categories, could be exploited to identify military and intelligence personnel, and to gather sensitive information about them.[28]

- ***Accountability (Including Limiting Unrestricted Onward Transfers)***

Many of the harms, including secondary uses of information, that arise from current digital advertising practices are a direct result of an inherent lack of accountability associated with real-time bidding processes. Insofar as RTB advertising exchanges rely on the relatively unrestricted sharing and sale of information on individual web browsing behavior (subject to contractual agreements) with hundreds, or even thousands, of third parties, it is next to impossible to ensure that data will only be used for appropriate purposes. As a result, to the extent that a novel advertising solution offers an alternative to this form of advertising, it may be considered successful from at least one significant data protection perspective.

---

[27] Press Release, Wyden Releases Documents Confirming the NSA Buys Americans' Internet Browsing Records, Calls on Intelligence Community to Stop Buying U.S. Data Obtained Unlawfully from Data Brokers Violating Recent FTC Order (Jan. 25, 2024), https://www.wyden.senate.gov/news/press-releases/wyden-releases-documents-confirming-the-nsa-buys-americans-internet-browsing-records-calls-on-intelligence-community-to-stop-buying-us-data-obtained-unlawfully-from-data-brokers-violating-recent-ftc-order.

[28] E.g., Emma Woollacott, *Web Browsing Data Is Serious Security Threat to EU and US*, Forbes (Nov. 14, 2023), https://www.forbes.com/sites/emmawoollacott/2023/11/14/web-browsing-data-is-serious-security-threat-to-eu-and-us/?sh=1a2510633929.

- ***User Control (Opt-Ins, Opt-Outs, "Do Not Sell") and Design***

Another relevant factor in data protection analysis is the extent to which individuals retain control over the processing of their information. This can include the ability to ***opt-in*** or consent to collection or specific processing, as well as the ability to ***opt-out*** or object to particular data processing. Other forms of individual control include the ability to exercise rights to ***access, correct, or delete*** information.

The exercise of individual control and individual rights over information is an integral part of data protection analysis, although it can sometimes be at odds with other elements, such as the minimization of data or reduction in identifiability. For example, data that has been pseudonymized, while it may still create some privacy risk to individuals, may no longer being sufficiently linkable to allow an individual to exercise an access or correction request. Finally, all opportunities for an individual to opt in or out of data processing, will necessarily involve some form of choice interface, and therefore an opportunity for "nudging," or suggestive design or impartial presentation of choices. In particular, coercive or heavily weighted choice interfaces are typically viewed as negating an individual's meaningful consent.[29]

- ***Comprehensiveness and Sensitivity of Profiling***

Among the most prevalent concerns around advertising technology is that websites and advertisers can track users across the internet to form detailed profiles about individuals, including to evaluate, analyze or predict personal aspects of an individual. In a sense, all advertising involves a kind of profiling, since advertisements are placed in ways that are designed to reach a particular target audience – for example, by advertising running shoes in a fitness magazine, or by running a political ad on a late night comedy channel in order to reach young voters. Advertising profiles or "segments" can often be benign, such as when they are based on a perceived interest in a retail product. They can also be invasive, harmful, or offensive, such as when advertisers seek to reach low-income struggling audiences with harmful content, such as payday loans.[30]

Audience segments can also range widely in accuracy, since advertisers are in the business of making predictions about likely interest, as well as shaping those interests. For all of these reasons, data protection experts and regulators typically look to the sensitivity of the profiling information (for example, whether it involves information on an individual's race, religion, sexual orientation, political beliefs, health information, or related sensitive categories), as well as the volume and comprehensiveness of the profiling.

---

[29] [I.e., "dark patterns"]

[30] *See* Federal Trade Commission, *Data Brokers: A Call for Transparency and Accountability* (May 2014), https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf ; and The Markup, From Heavy Purchasers of Pregnancy Tests to the Depression-Prone, We Found 650,000 Ways Advertisers Label You (June 8, 2023), https://themarkup.org/privacy/2023/06/08/from-heavy-purchasers-of-pregnancy-tests-to-the-depression-prone-we-found-650000-ways-advertisers-label-you.

- ***Impact on Marginalized or Vulnerable Populations***

Digital advertising may also result in or perpetuate discrimination, resulting in advertising that is targeted or restricted based on protected characteristics. For example, a 2018 case brought against Facebook (now, Meta) alleged that Facebook and online advertisers were engaging in gender discrimination by disproportionately targeting men with job advertisements.[31] The Department of Justice also brought a case against Meta for violations of the Fair Housing Act, alleging that Meta's advertising algorithms filtered who was eligible to see certain housing ads based on users' race and sex.[32]

5. Economy and Society: The Effects of Digital Advertising on Competition, Political Influence, and Other Social Values (Framework Parts C and D)

Finally, this Proposed Framework aims to expand the traditional debates that frame utility and privacy as the only relevant tradeoffs to be considered in evaluating an advertising system. In recent years, it has become increasingly apparent that advertising impacts every sector of the economy, including both large and small players. Changes in mainstream advertising models and solutions can have an outsized impact on small businesses, small and medium sized publishers and independent and local journalism, and other economic metrics.

Many of these considerations are fundamentally questions of economic power or competition, either between large and small advertising providers, or between closed systems ("walled gardens") and the more open, decentralized sharing of information that occurs on the "open web." It is also the case that the collection and use of data can favor existing dominant players in the advertising industry.

Finally, we observe that beyond utility, privacy, and competition, there are a range of social values related to understanding political influence, giving access to researchers, and "social good" questions that frequently arise in debates about digital advertising. For example, niche and hyper-local marketing campaigns are often used by a range of businesses and non-profits, including government agencies, to communicate public health messages or other social benefits. The same technology that enables micro-targeting for advertising retail products can enable recruiting for clinical trials for rare conditions. And similarly, evaluators should consider the impact on the non-profit sector, and the ability of social causes to fundraise and reach audiences interested in social or political causes.

---

[31] Janet Burns, *ACLU Suit Says Facebook Excluded Women, Older Men from Targeted Job Ads*, Forbes (Sept. 20, 2018), https://www.forbes.com/sites/janetwburns/2018/09/20/aclu-suit-says-facebook-excluded-women-older-men-from-targeted-job-ads/?sh=33ddae757a6e.

[32] U.S. Department of Justice, Justice Department and Meta Platforms, Inc. Reach Key Agreement as They Implement Groundbreaking Settlement, Press Release (Jan. 27, 2023), https://www.justice.gov/opa/pr/justice-department-and-meta-platforms-inc-reach-key-agreement-they-implement-groundbreaking.

What's Ahead: Next Steps

The authors actively welcome feedback on the proposed Risk-Utility Framework presented in Appendix 1 (below). In a future iteration, we aim to ensure that we have captured all relevant questions, presented them in a useful and practical way, and demonstrated many of the inherent tradeoffs between relevant factors.

As regulators and practitioners evaluate new systems and both craft and enforce new regulations, we anticipate that the era of the "advertising sea change" will continue for many years to come. A more privacy-protective era will necessarily require re-visiting many of the assumptions of previous decades, and returning to key principles of transparency, data minimization, purpose limitation, and individual consent. Depending on which aspects of privacy and data protection are most salient in our laws and policies, these will necessarily involve tradeoffs to many aspects of utility, competition, and other social values. Our goal is for this rubric to provide a resource for a holistic and fair analysis of new systems as they continue to develop.

- *Are there any additional utility, privacy, competitive, or other social consequences that should be considered in a holistic analysis of a novel advertising solution or system?*
- *How can we effectively demonstrate, including through visual elements, the relevant tradeoffs and interconnectivity between different factors?*
- *What are the circumstances in which non-data protection elements (such as utility, competition, or social impact) ought to be considered as part of Data Protection Impact Assessment or other impact analysis?*

Email comments to info@fpf.org by May 26, 2024 or contact us to discuss further.

Appendix 1

Proposed Risk-Utility Framework: A Rubric for Evaluating an Advertising Solution's Impact on Utility, Privacy and Data Protection, and Other Values

The following questions are intended to facilitate an holistic assessment of the impact of an advertising system on the basis of its advertising utility, impact on core data protection principles, and other societal values, including economic efficiency, competition, safety, political influence, and related research purposes.

Questions are intended to be addressed largely in order, although not all may be relevant to every advertising solution. We do not assign values to each question, leaving the reader to assess the relative importance of different assessments. Unless otherwise noted, the terms "personal information" and "individual" are not intended to refer to legal concepts, but instead should be conceptualized broadly to include any information collected in connection with an individual, their device(s), their browser(s), or their household, for the purposes of running an advertising campaign.

A. Advertising Utility

Some privacy-enhancing limitations to an advertising system will necessarily decrease the utility of that system by decreasing the availability of information for use in targeting, measurement, and attribution (for example, using non-persistent or short-range forms of identification rather than persistent or globally unique identifiers). The following questions are designed to assess the impact of a system primarily in terms of its utility to advertisers.

Ad Delivery and Reporting

1. Does the solution enable the **technical delivery** of an advertisement (e.g., in the correct size and format appropriate to the publisher's specifications and the user's browser, device, or operating system)?
2. Does the solution enable **frequency capping** (limiting the number of times that an advertisement is displayed to the same user or household)?
3. Does the solution enable **ad sequencing** (delivering creatives in a specific order to the same user or household)?

Measurement and Attribution

4. Does the solution enable an advertiser to know, with sufficient precision, the number of viewers or impressions an advertisement received?
5. Does the solution enable an advertiser to know, with sufficient precision, how many viewers converted to known customers (i.e., made a purchase, signed up for a service, or downloaded an app)?

Targeting and Audience Reach

6. Does the solution enable an advertiser to reach small or niche audiences (e.g. Sci-fi readers who live in Tampa, FL)?
7. Does the solution enable an advertiser to retarget or remarket to individuals who are not known customers, but may be interested in a product or service because they visited a website or clicked on a link?
8. Does the solution enable an advertiser to retarget or remarket to individuals on the basis of "offline" information, such as visiting a retail location?
9. Does the solution enable cross-device tracking?
10. Does the solution enable lookalike modeling?

Quality and Brand Safety

11. Does the solution enable advertisers to prevent their advertisements from being displayed in inappropriate contexts? (E.g. an airline promotion displayed next to an article about an airline crash)
12. Does the solution enable advertisers to prevent their advertisements from being displayed on particular publications (e.g., to avoid particular platforms or publishers)?
13. Can the processes and data flows be audited for compliance and ethical operation?

B. Data Protection and Privacy

The following questions relate to personal information (PI) collected, shared/sold, or processed for purposes of delivering and measuring advertising campaigns. Broadly, they are categorized on the basis of the principles and values derived from the GDPR, US state consumer privacy laws, and closely related data protection values (such as non-discrimination and the protection of children).

Lawfulness, Fairness, and Transparency

14. Is the PI collected as part of a direct relationship with the individual?
15. Does the individual have access to a concise, easily accessible, and easy to understand description of the controller's data practices, including the purposes for which data is processed?

16. Does the individual about whom the PI relates have the ability to object to the processing of the data or prevent it from occurring?
17. Does the individual have the ability to easily access the underlying data, including any inferences made about them?
18. Is information about the technology easy or intuitive for the user to locate?

Purpose Limitation (Secondary Uses)

19. If the PI was collected for a non-advertising purpose, has it been re-used for advertising without the individual's awareness, understanding or involvement?
20. If PI is collected for the purpose of advertising, does the collecting entity retain the right or ability to re-use PI for additional related purposes (such as improving targeting algorithms, identifying fraud or malware, or identifying bias or discrimination)?
21. If PI is collected for the purpose of advertising, does the collecting entity retain the right to re-use PI for additional unrelated purposes (such as sales for non-advertising business analytics use cases, law enforcement, or unrestricted rights to use data)?

Data Minimization and Storage Limitation

22. Is information collected necessary to achieve the specific advertising purposes and use cases communicated to the user?
23. Are individuals directly or indirectly identifiable on the basis of the data collected?
24. Is the PI tied to any universally unique or commonly used device identifier, such as a MAID, IMEI, phone number, or email address?
25. Is the PI collected associated with any persistent form of identification such that the same individual can be recognized over time and/or across devices, beyond the scope of the initial point of data collection?
26. If Yes to above: Is the persistent identifier used by other unaffiliated controllers such that it can be used to identify behavior across unrelated websites or platforms?
27. How long is information retained in a readily identifiable format?
28. How long is information retained in any individualized format (even if steps are taken to de-identify the data)?
29. Are time limits specific and limited (e.g. "three months"), rather than vague, indefinite, or adjustable ("as business needs require")?

Accuracy

30. Does the controller have procedures for periodically reviewing, correcting, or erasing any inaccurate information?
31. Does the individual about whom the PI relates have the ability to request that any inaccurate information about them be corrected?

Accountability (including Unrestricted Onward Transfer)

32. Is PI shared with unrelated or unaffiliated controllers in order to facilitate advertising purposes? (I.e., excluding genuine service providers acting solely at the behalf of a single controller, not re-using, sharing/selling, or retaining data for any other purposes)
33. Is PI shared with a *very large number* of unrelated or unaffiliated controllers in order to facilitate advertising purposes?
34. Is PI shared with unrelated or unaffiliated controllers as part of an unrestricted sale or monetization arrangement?
35. If Yes (to any of above): Are the unrelated or unaffiliated controllers able to access data easily and without restrictions, or are they subject to meaningful technical, legal, or contractual restrictions?

User Control (Opt-Ins, Opt-Outs, "Do Not Sell") and Design

36. Does the advertising system prevent the sale or unrestricted transfer of data?
37. Does the individual have the ability to enable, or disable, the data collection at its source, e.g. through blocking of sharing of data?
38. Does the individual have the ability to communicate an automated opt-out preference signal to all or most entities in an automated way (e.g., the Global Privacy Control) that will be respected by recipients?
39. If Yes (above), does the opt-out preference signal contain a mechanism for recipients to communicate back to the individual that the signal has been honored or respected?
40. If the advertising technology includes user permissions and signals (like Do Not Sell or user consent), does the system ensure that they are adhered to throughout the value chain, beyond contractual obligations?
41. Does the individual have the ability to communicate a direct opt-out or deletion request to a company in an easily accessible manner, without encountering an overly burdensome process or unnecessary requests for information to verify the request?
42. Are options presented symmetrically (e.g. "I consent" and "I do not consent"), and free of deceptive or manipulative design?

Comprehensiveness and Sensitivity of Profiling

43. Does the controller have the ability to buy and/or append "offline" data, or data from unrelated sources and combine it with existing profiles?
44. Does the system facilitate the creation of highly comprehensive behavioral profiles of an individual's behavior?
45. Does the system allow for direct inference of sensitive characteristics or attributes of an individual (such as race, ethnicity, political affiliation, sexual orientation, transgender identity, union membership, etc.)
46. Does the system allow for indirect inference of sensitive characteristics or attributes of an individual (such as race, ethnicity, political affiliation, sexual orientation, transgender identity, union membership, etc.)

47. If Yes (to either of above), are there protections or safeguards in place to prevent sensitive characteristics from being shared, misused, or used to target content?

Protection of Children

48. Does the system allow for targeting of content directed to children?
49. If intended for general audiences, does the system allow for indirect inference that an individual is from or about a child? If so, does the system have dedicated mechanisms to prevent the identification of, analysis or sharing of information about, or targeting of content to, children?
50. Does the advertising technology support parental controls or age verification?

Impact on Marginalized or Vulnerable Populations

51. Does the system enable the targeting of individuals with content on the basis of protected characteristics (e.g. race/ethnicity, religious beliefs, political affiliation)?
52. Does the system allow for detecting and correcting biases, especially those inherent in AI models? How adequate are these systems? (What is the data's provenance? How representative is the data? Is the data balanced? Has intersectionality been considered? Are sensitive attributes present in the data?)

C. Economic and Competitive Implications

Restrictions on data, including as a result of regulations, will necessarily impact the economic landscape for advertisers, intermediaries, and publishers – "favoring" some and "disfavoring" others through access to data. Notably, the actual economic impact of some changes will be dependent on alternatives that continue to exist in the market or are developed as innovative solutions to regulatory pressures. The following questions are designed to help address the extent to which changes create undue hardships or competitive disadvantages:

53. Does the advertising system favor "closed" or account-restricted systems over the unrestricted sharing and dissemination of information on the "open web"?
54. To what extent does the advertising system create reliance on the advertising measurement capabilities of a single or smaller number of large or dominant platforms to facilitate delivery and measurement of advertising?
55. To what extent does the advertising system create greater reliance on "first party data," or disproportionately benefit large or well-established brands or publishers?
56. Does the advertising system create undue economic barriers, such as diminished advertising revenue, for small and medium publishers or content creators?
57. Does the advertising system create undue economic barriers, such as higher prices, for small and medium businesses that advertise to small or niche audiences?
58. Does the technology simplify stakeholder cooperation and compliance or add complexity layers in a tightly interwoven ad tech ecosystem?

D. Social Implications (Research, Political Influence, Social Causes)

An advertising solution can also be assessed on the basis of its impact on social and other factors, including the ability to facilitate public and regulatory oversight of political advertising, influence, and disinformation.

General

59. What is the effect on detection and mitigation of advertising fraud, e.g. bots and invalid network traffic?
60. What is the effect on detection and mitigation of advertisement-delivered malware?
61. Are there measures in place to detect any form of misuse at scale?
62. Is the solution open or accessible to outside researchers to verify or test privacy claims or otherwise ensure the system's transparency and accountability?

Political and Foreign Influence

63. Does the solution allow for detecting and mitigating ads that promote hate speech, abusive content, or illegal content?
64. Does the solution allow for the placement of specific or issue-based political advertisements?
65. Does the solution contain safeguards for the kind of entities that can place political ads, in order to prevent, assess, measure, or mitigate influence from foreign nations?
66. Does the solution contain protections to prevent the sale, transfer, or access, of information to foreign adversaries?

Social Causes and Democracy

67. What is the effect, if any, of the mainstream use of the advertising solution on small, local, and/or independent journalism and news media?
68. What is the effect, if any, on support for public health and other government agency communications and advertising?
69. What is the effect, if any, on non-profits fundraising for social causes?
70. What effect, if any, on recruiting for specific clinical trials for rare conditions?