

Testimony of Ashkan Soltani¹
Independent Privacy Researcher and Consultant

United States Senate, Committee on Commerce, Science, and Transportation
Hearing on
The State of Online Consumer Privacy

March 16, 2011

Chairman Rockefeller, Ranking Member Hutchison, and the distinguished members of the Committee, thank you for the opportunity to testify about online consumer privacy and the state of tracking on the Web today.

My name is Ashkan Soltani. I am a technology researcher and consultant specializing in consumer privacy and security on the Internet. I have more than 15 years of experience as a technical consultant to Internet companies and federal government agencies. I received my Master's degree in Information Science from the University of California at Berkeley, where I conducted extensive research and published two major reports on the extent and means of online tracking. Last year, I served as a staff technologist in the Division of Privacy and Identity Protection at the Federal Trade Commission on investigations related to Internet technology and consumer privacy. I have also worked as the primary technical consultant on the Wall Street Journal's *What They Know* series investigating Internet privacy issues on the ground.

I have been asked to testify about the current state of online tracking from a technical perspective. I will describe the basics of how online tracking works and discuss some of my research that demonstrates how pervasive tracking is online today. I will then discuss the extent to which consumers are actually aware that they are being tracked online and whether they are able to meaningfully control unwanted tracking with existing industry-provided and browser-based mechanisms. Finally, I will discuss the Do Not Track proposals in light of these findings.

A. How Online Tracking Works

As an illustrative example to explain how consumers are tracked online, we can step through a typical Web browsing session. A user wants to look up information about cholesterol on WebMD, so he types "www.webmd.com" into his browser's location bar and navigates to a specific page on WebMD's site focused on cholesterol. The browser contacts the WebMD server to retrieve the contents of the page. Much of the page's content will be provided directly by WebMD itself, but some of the content may originate from other entities, such as an advertisement provided by an online advertising service such as Google's DoubleClick. As a result, although the browser's location bar will show "www.webmd.com," many other third party entities may have a presence on the website, and often it is unclear to the user which content comes from which provider.

¹My oral and written testimony here today to the Committee represents my own personal views, and does not reflect the views of any of the organizations I have consulted or worked for in the past.

A useful analogy may be to imagine a picture frame that has slots to display a number of different photos. WebMD provides the “frame” and a few of the “photos,” while the rest of the “photos” are provided by third parties that WebMD has partnered with. This practice of embedding content from third party entities is nearly universal on the Web today. As I will explain below, it is primarily these third party entities that are capable of tracking users as they browse the Web.

In this example, the WebMD page on cholesterol includes a third party online advertisement that is displayed at the top of the page. As the web browser fetches the ad, two things relevant to tracking typically occur. *First*, the company providing advertisements can attempt to uniquely identify the browser using a variety of technical mechanisms, which I will discuss below. The simplest and most common technique is to use a browser cookie. In this context, a cookie is a file containing a unique identifier that is placed on the user’s computer by the third party ad service and is transmitted back to the service upon each subsequent ad request². *Second*, the ad service can record detailed information about this interaction. The ad service may log the date and time of the ad request, which ad was displayed, and perhaps the details about the content of the WebMD page on which the ad was shown. Most importantly, the ad service can *link* all this information to the unique identifier, and collect this information together in a consumer database.

Some time later, the user checks the weather by browsing to “www.weather.com.” It turns out that the same third party ad service used by WebMD is also providing ads for the Weather Channel’s site. As an ad loads in the margins of the Washington DC forecast page, the ad service can again uniquely identify the user’s browser, using the same cookie file that was previously stored. The ad service can now tie the user’s browsing activity between the two sites together—the same browser that previously accessed health information about cholesterol also looked up the weather forecast in Washington, DC. As the user continues to browse, this ad service can continue to follow the user’s activity on the websites on which it has a presence. These activities are the essence of online tracking.

Web browsing interactions are generally described as being in one of two categories, *first party* or *third party*. A first party is typically defined as an entity whose site the user knowingly visits and whose Web address appears in the browser’s location bar—in the scenario above, WebMD and then later, the Weather Channel. Users typically interact with a first party by directly typing its Web address into the location bar or by browsing to it from another site, for instance, by following a link from a search engine or a social network.

A third party is an entity that provides content that is included on a first party site, like the ad service in our earlier scenario. While some third party interactions are visible to the user, such as a displayed ad or an embedded video, it may not be clear that this content is being provided by someone other than the site they are visiting. However, other third party interactions may be

²Cookies are text files that can store various types of information. For the purposes of tracking, they typically contain unique descriptors such as *user=1234567890* or *email=john.doe@host.com*.

invisible to the user. For example, a “web bug” is an imperceptible image placed on first party sites, but operated by third parties, for the express purpose of invisibly tracking users.³ These third party tracking objects can only appear on a site with the knowledge and consent of the first party. As an example, ads from Google DoubleClick will only appear on Weather Channel pages if the Weather Channel explicitly decides to include DoubleClick on its site.

Note also that the same business entity can be both a first party or a third party, depending on the context. For instance, if a user browses directly to “www.youtube.com” to watch online videos, YouTube is a first party. But, if a first party site such as CNN.com embeds a YouTube video into one of its stories, YouTube is now a third party.

In our scenario, the ad service uses a standard browser cookie to link together two separate user interactions—one on WebMD and the other on the Weather Channel. Even though the cookie by itself does not usually identify the user by name, third party trackers are able to build a “browsing profile” that consists of data from numerous Web interactions over time from the same user.⁴ This browsing profile has the potential to reveal quite a bit of information about the user’s real world identity.⁵

Despite some claims that these collected browsing profiles are “anonymous,” recent computer science research suggests that it is often quite easy to re-identify datasets that contain user information.⁶ As the number of data points in a browsing profile increases, so too does the possibility that it can eventually be re-identified to reveal the user’s actual identity, such as a name, e-mail address, or other personally identifiable information. For example, when a user purchases a product online, the merchant could decide to share the user’s e-mail address—collected in the billing process—with a third party ad service that is present on the purchase page. This issue can also arise with the use of social networks, whereby identifying information

³Web bugs are sometimes also referred to as tracking pixels or web beacons. Web bugs are typically used to provide websites with information that will help them understand and optimize web usage, and typically track users.

⁴Of course, some browsers may be shared by multiple users, but often browsers will be used primarily by a single user. This is particularly salient in the case of mobile phones, where the sharing of devices is less common.

⁵Each data point may also reveal the time of each site access and in many cases the user’s approximate geographic location based on his IP address. More advanced tracking techniques on a single page may be able to determine exactly how the user moves his mouse on the page or what text on the page gets highlighted and copied.

⁶Narayanan, A., & Shmatikov, V. (2008). How to Break Anonymity of the Netflix Prize Dataset. In Proc. of 29th IEEE Symposium on Security and Privacy, Oakland, CA, May 2008, pp. 111-125. *and* Ohm, P. Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization (2009, August 13). University of Colorado Law Legal Studies Research Paper No. 09-12. Available at SSRN: <http://ssrn.com/abstract=1450006>

may leak to third party ad services.⁷

1. The State of Online Tracking

The practice of using third party services to add tracking and other functionality to a website is quite common. In our Berkeley *KnowPrivacy* study, we found an average of 12 trackers present on each of the top 100 most popular websites, with one having as many as 100 different trackers over the course of a month.⁸ This means that when a user visits that website, potentially 100 entities—nearly all unseen by the user—will learn about the visit.

The very reason why online tracking is both effective and why it raises privacy concerns is that third party entities can track consumers across multiple unrelated first party websites. In our Berkeley study, we also found that some third party trackers have an extensive “reach” across a large number of first party sites. One advertising company was able to monitor activity on 91 of the top 100 most popular sites, as well as 88% of 350,000 sites sampled in our dataset, as of March 2009.⁹ In 2010, a leading social network announced that their third party sharing widgets were present on 2.5 million websites¹⁰ and growing at a rate of 10,000 sites per day.¹¹ In both these examples, the presence of third party objects generates a steady stream of data that flows to a single entity. These uniquely pervasive positions give these companies the capacity to infer a great deal more than just a user’s interest in automobiles or sporting goods. Their tracking technologies reach the vast majority of every user’s Web browsing activity.

It is important to point out that online tracking is not limited to Web browsers. Consumers are connecting to the Internet using a variety of devices that extend beyond what we consider a typical PC-and-browser setup. Mobile phones, televisions, set top boxes (such as a Tivo or a cable box), video game consoles and even some automobiles are now equipped with Internet connectivity and can leverage Web services which include online advertisement. Some of these platforms also allow applications written by third parties, the most prominent example

⁷Krishnamurthy, B. and Willis, C. (2009). On the leakage of personally identifiable information via online social networks. In Proceedings of the 2nd ACM workshop on Online social networks (WOSN '09). ACM, New York, NY, USA, 7-12. DOI=10.1145/1592665.1592668 from <http://doi.acm.org/10.1145/1592665.1592668>

⁸Gomez, J., Pinnick, T., and Soltani, A. (2009, June 1). *KnowPrivacy* available at http://knowprivacy.org/report/KnowPrivacy_Final_Report.pdf, p.26.

⁹Id. p.27.

¹⁰Constine, J. (2011, February 27). All of Facebook’s Like Buttons on Third-Party Sites Now Publish a Full News Feed Story. Inside Facebook - Tracking Facebook and the Facebook Platform for Developers and Marketers from <http://www.insidefacebook.com/2011/02/27/like-button-full-story/>

¹¹Parr, B. (2010, October 26). 10,000 Websites Integrate with Facebook Every Day. Social Media News and Web Tips – Mashable – The Social Media Guide. from <http://mashable.com/2010/10/26/10000-websites-integrate-with-facebook-every-day/>

being “app stores” on mobile smartphones.¹² Mobile devices, in particular, raise unique privacy concerns because consumers carry them nearly all of the time.¹³ As such, applications and services running on the phone may have the ability to access precise geolocation information, using GPS technology, to learn even more intimate details about a consumer’s physical habits.

2. Existing Privacy Tools are Easily Circumvented

Every major Web browser includes privacy enhancing technologies that can be used by consumers to limit the extent to which they are tracked online. Unfortunately, these built-in tools, which include “private browsing modes” and cookie controls, only protect users from some tracking technologies, and do not provide consumers with the privacy protections they may reasonably expect.¹⁴

As one example, cookie blocking features in the major Web browsers do not always work in the same way, and even sophisticated users do not fully understand these intricacies.¹⁵ This may cause consumers to have misplaced beliefs about the extent browsers are protecting them from tracking. But even when consumers do understand how these features work, sites have consistently devised new ways to track users and evade the protections of existing privacy tools.

In a study called *KnowPrivacy* published by my Berkeley colleagues and I in 2009,¹⁶ we found that several ad services had deployed a new stealthy technique to *resurrect* tracking cookies, even after the user had used the available cookie deletion tools built into his browser. Ad services developed a way to “remember” the cookie file using another technology—Adobe’s

¹²The Wall Street Journal reported that 47 of the 101 third party mobile applications tested transmitted location to third parties. 56 of the same apps transmitted unique device identifiers (UDIDs) which act similar to permanent cookies, and which users currently have no control over. See Thurm, S. (2011, December 17). iPhone and Android Apps Breach Privacy - WSJ.com. The Wall Street Journal from <http://online.wsj.com/article/SB1>

¹³Three in five mobile phone owners say they carry their phones at all times, even inside the home. See: Stanton, D. (2008, September 8). New Study Shows Mobile Phones Merging New, Established Roles. Knowledge Networks from http://www.knowledgenetworks.com/news/releases/2008/091808_mobilephones.html

¹⁴Soghoian, C. (2010, December 9). Why Private Browsing Modes Do Not Deliver Real Privacy, Internet Architecture Board, Web Privacy Workshop, from http://www.iab.org/about/workshops/privacy/papers/christopher_soghoian.pdf

¹⁵Not all browsers implement third party cookie blocking in the same way. Typically browsers allow third party cookies by default but if a user elects to configure their browser to block third party cookies, 3 of the 4 major browsers allow the third party cookies to be read if they were previously set, such as in a first party context. This is a small technical nuance, but it allows certain players to proceed as normal with regards to online tracking and potentially cause confusion for consumers as to the degree their privacy is protected. Additionally, it significantly effects whether certain players, *i.e.*, those that consumers have a first party relationship with, receive a competitive advantage over the lesser known websites.

¹⁶Soltani, A., Cauty, S., Mayo, Q., Thomas, L., and Hoofnagle, C., Flash Cookies and Privacy (2009 August 10). Available at SSRN <http://ssrn.com/abstract=1446862>

Flash Player—such that they could restore the cookie later, even after the user deleted it. This tracking technology—commonly called Flash cookies—is even *more* difficult for users to manage with existing privacy tools, when compared to standard cookie controls.¹⁷

Further, some ad services have shifted to new, cutting-edge tracking techniques, many of which are beyond the control of consumers¹⁸. While these are less well known, they are no less powerful—and in some cases more powerful—in their ability to track users’ browsing activities. From a technical perspective, browser vendors—and thus consumers—are losing the game of privacy Whack-a-Mole. The ongoing development of new, hidden tracking techniques is far outpacing the ability of browser vendors to develop and deploy adequate defenses. As a result, consumers and the privacy controls available to them will likely fail to keep up.

B. Existing Consumer “Notice and Choice” Mechanisms

The current system of industry self-regulation stresses two complementary approaches regarding online tracking: *notice*, through privacy policies and in-ad enhancements, and *choice*, through ad preference managers and industry-provided opt-out tools.

1. Privacy Policies

For more than a decade, websites have routinely included privacy policies, typically linked to from the bottom of the front page. These documents are often long and difficult to read—most likely because they are written by lawyers, for lawyers—and have not helped consumers to stay informed about the degree of tracking online.¹⁹ Research has also shown that the majority of Americans incorrectly believe that the phrase “privacy policy”—and its mere presence on

¹⁷Adobe has denounced the use of its Flash technology in order to restore tracking cookies. Although not yet widely deployed, the company has recently taken steps to work with major browser vendors in order to move Flash cookie privacy controls directly into the browser settings and allow users to manage them in a similar way as standard cookies. See Albanesius, C. (2011, March 8). Adobe Flash Player 10.3 Beta Adds Greater Control Over 'Flash Cookies' PC Magazine. from <http://www.pcmag.com/article2/0,2817,2381650,00.asp>

¹⁸In the past year, I have confirmed tracking by third party companies on widely used websites using mechanisms including but not limited to browser fingerprinting (<http://radar.oreilly.com/2011/03/device-identification-bluecava.html>), cache cookies (<http://www.wired.com/epicenter/2009/08/flash-cookie-researchers-spark-quantcast-change/>), CSS history profiling (<http://blogs.forbes.com/kashmirhill/2010/11/30/history-sniffing-how-youporn-checks-what-other-porn-sites-youve-visited-and-ad-networks-test-the-quality-of-their-data/>), domain masquerading (<http://doi.acm.org/10.1145/1592665.1592668>), UDIDs (<http://online.wsj.com/article/SB10001424052748704694004576020083703574602.html>), and HTML5 storage (<http://www.wired.com/threatlevel/2010/09/html5-safari-exploit/>) to track consumers in ways that are difficult or even impossible to control.

¹⁹McDonald, A. and Cranor, L. (2008) The Cost of Reading Privacy Policies. I/S: A Journal of Law and Policy for the Information Society 2008 Privacy Year in Review issue. [Paper originally presented at TPRC 2008, Sept 26-28, 2008, Arlington, VA.] *and* Privacy Leadership Initiative. Privacy Notices Research Final Results. Conducted by Harris Intereactive, (2001 Dec) from <http://www.ftc.gov/bcp/workshops/glb/supporting/harris%20results.pdf>.

websites—signifies that their information will be kept private.²⁰

While there is much data to suggest that consumers do not actually read or understand privacy policies, even if they did, many existing privacy policies often provide confusing or even conflicting information. In our *KnowPrivacy* study, we found that, among the top 50 most popular websites, many sites that claim to not share information with “third parties” later disclaim that they do share information with “affiliates”, which sometimes number well over 2000 companies.²¹

2. Enhanced Notice for Online Ads

One emerging self-regulatory measure is “enhanced” or “robust” notice for online ads. The purpose of enhanced notice is to increase transparency—directly within the ad—into why the particular ad was chosen and what the attached terms and policies are. Although this is a commendable step forward, the question is how many users will notice. One self-regulatory firm noted that, during the first few months of the industry’s initiative, the notice on only 0.004% of “enhanced” ads were clicked by users actually clicked through to the detailed explanatory text.²² While the initiative is in its early days, this calls into question whether enhanced notice will be sufficient to deliver meaningful transparency.

3. Ad Preferences Managers

The advertising industry has also created online tools that allow users to view and modify marketing inferences made about them within “ad preferences managers.” For example, an ad preferences tool may show the inferences made about the user’s demographic information (such as age, income range, education, or geographic location), shopping interests (such as sports, technology, or politics), or even significant life events (such as “getting married soon” or “having a baby”) based on the user’s browsing activity. In many cases, these tools also allow consumers to opt-out of certain consumer marketing sectors from which they do not wish to receive targeted ads.

Like enhanced notice, ad preference managers improve transparency into the online ad serving ecosystem. But, these managers only present a high-level summary of the information collected by the ad service. Given their vantage point, third party ad services have the capability to make

²⁰Turow, J., Mulligan, D., and Hoofnagle C. (2007 Oct), Consumers Fundamentally Misunderstand the Online Advertising Marketplace, from http://groups.ischool.berkeley.edu/samuelsonclinic/files/annenberg_samuelson_advertising.pdf

²¹Of the top 50 sites, all stated they collect IP address, 48 collect contact information such as name and e-mail address, and 39 collect click stream information. Bank of America had over 2,300 “affiliates”. See Gomez *et. al.* p24 (previously cited) and *KnowPrivacy*, <http://knowprivacy.org/profiles/bankofamerica>

²²Evidon served over 11 billion impressions in their first full scale months. Among those who click on the icon (on .004% of ads served), about 3% of users opt out of one or more provider. See Smith, S. (2011, March 11). MediaPost Publications Browsing Privacy's Next Steps 03/11/2011 from http://www.mediapost.com/publications/?fa=Articles.showArticle&art_aid

inferences or use the data for other, non-advertising-related purposes, that are not shown in the ad preference managers.²³ I'm not implying that specific companies are engaged in this practice, just that collection, retention, and correlation of this behavioral data *provides the capacity* for these inferences to be made. More transparency is needed—outside the realm of online targeted ads—about the information that is collected by third parties and how they are used.

4. Cookie-based Choice Mechanisms

In addition to notice and transparency, many ad services provide users with the ability to opt-out. Currently, most opt-outs work using special opt-out cookies—one for each ad service—stored in the user's Web browser. The cookie-based opt-outs have been plagued by a number of problems, some of which have been addressed in recent years and others which persist today.

Once consumers realize they are being tracked, they must then begin the process of obtaining opt-out cookies from each tracking company. One self-regulatory technology firm has identified 600 companies involved in collecting or using tracking data about customers on their sample of 7 million domains.²⁴ Another lists 323 tracking companies publicly.²⁵ Given the value of this marketplace and the speed with which new entrants emerge, I suspect the actual number of companies engaged in tracking may be actually be even larger. Even still, identifying 600 hidden trackers and obtaining an opt-out is a daunting task for even the most sophisticated privacy-conscious consumer.

Seeking to ease the process of obtaining opt-out cookies, industry self-regulatory groups such as the Network Advertising Initiative (NAI) have created one-stop websites where consumers can obtain opt-out cookies for multiple firms. However, these opt-out sites do not comprehensively cover all online tracking since only a fraction of approximate 600 companies discussed are covered.²⁶ This problem exists in the mobile space as well. Currently, nine of the 16 mobile ad companies do not offer an opt-out,²⁷ and data collected on mobile phones may be particularly sensitive, since it is often accompanied by hardware identifiers that users cannot

²³Similar to sports and shopping habits, a user's browsing habits could allow an observer to make inferences about a user's race, sex, sexual orientation, health status, financial health, and political affiliation, even though these categories are typically excluded from online preference managers.

²⁴Steel, E. (2011, March 4). Council of Better Business Bureaus to Enforce Online Tracking Principles - Digits. WSJ Blogs - WSJ from <http://blogs.wsj.com/digits/2011/03/04/council-to-enforce-online-tracking-principles/>

²⁵PrivacyChoice Tracker Index (Mar 14 2011) from <http://www.privacychoice.org/companies/all>

²⁶At the time of this writing, the NAI opt-out (http://www.networkadvertising.org/managing/opt_out.asp) currently allows consumers to opt-out of behavioral advertising by 68 member companies. AboutAds opt-out applies to 61 companies (<http://www.aboutads.info/choices/>) and even the most comprehensive list of trackers, offered by the independent group PrivacyChoice only allows opt-out of 160 (<http://www.privacychoice.org/privacymark>)

²⁷Brock, J. (2011, March 16). Mobile Tracking Privacy: Three thoughts. PrivacyChoice Blog. from <http://blog.privacychoice.org/?p=2882>

change or geographic location information.

Most importantly, even when opt-outs are available, many firms only allow the user to opt-out of the receipt of targeted advertising, not the online tracking itself. *Advertisers continue to collect and retain data in order to build a profile on the user*, even in the presence of an opt-out cookie.

Finally, cookie-based opt-out mechanisms are *inherently brittle*. Users are frequently taught to delete their browser cookies on a periodic basis to better protect their online privacy. But, when the user clears her browser cookies, she will also inadvertently clear her *opt-out* cookies, which will—counter-intuitively—opt the user back *in* to tracking.

C. Do Not Track Proposals

Last July, this Committee held a hearing on the topic of online privacy during which the idea of “Do Not Track” was discussed. Ever since, there has been a significant amount of public discussion and debate regarding the possibility of a Do Not Track mechanism. While the name—Do Not Track—sounds much like the highly successful Do Not Call list,²⁸ the only substantive similarity is that they both give consumers a single point of control to express their privacy preferences. While consumers can register their phone number in a FTC registry for Do Not Call, the single point of control for Do Not Track is likely to be a preference setting in the consumer’s Web browser or mobile platform.

Two primary technical approaches to Do Not Track have been proposed and implemented by major Web browser vendors. The first method is called the *header* approach, and the second is called the *blocking* approach. Two browser vendors have already taken steps to include these mechanisms in upcoming releases of their products.²⁹

1. The Header Approach

In the *header* approach, the consumer can toggle a Do Not Track setting in his Web browser privacy preferences. When this setting is enabled, the browser transmits a special signal to each remote server that the consumer has expressed his preference to not be tracked.³⁰ The idea is to give users the ability to send a clear, persistent and technology-neutral signal to websites regarding their tracking preference. Of course, in order this mechanism to be effective, it will depend upon a clear set of rules defining what websites should do when they receive this signal.

²⁸The Do Not Call list is an FTC enforced initiative based on legislation that creates a centralized registry of numbers that telemarketers may not call, under monetary penalty.

²⁹Mozilla’s Firefox 4.0 and Microsoft’s Internet Explorer 9 (MSIE9) have announced support for the header mechanism. MSIE9 also supports the blocking method as well via their Tracker Protection Lists product.

³⁰Current proposals involve sending a Do Not Track signal using a browser header within the HTTP protocol.

Under this approach, the onus is on the server to agree to respect the consumer's preference. It is possible that the server could ignore the user's request and continue to engage in tracking anyway, even once best practices are established. Thus, consumers will need a method to verify that servers are complying with the header, so they can keep firms honest about their commitment to respect user tracking preferences. Publisher sites and US brands that advertise could choose to favor ad services that respect the header preference.

2. The Blocking Approach

In the *blocking* approach, the consumer maintains (perhaps with the help of a trusted third party) a list of servers that are known to engage, or are suspected of engaging, in unwanted tracking behavior. Once a user has enabled this feature, his Web browser will automatically block all connections to the servers on the list which could also result in the blocking the display of advertisement.

As opposed to the header approach, the responsibility to prevent tracking is solely on the consumer, that is, to obtain an up-to-date list of suspected tracking servers and to block them. Servers are under no express obligation to abstain from tracking, so if one is not blocked by a consumer's browser, it is free to continue tracking as usual.

One concern with this approach is that it is sometimes difficult for consumers-at-large to determine whether a domain is engaging in tracking behavior and whether to add that domain to the block list. Additionally, there are many technical mechanisms that exist today that could be used to circumvent such blocking measures.³¹

3. Other Considerations

For any consumer choice mechanism to work, we need to clearly define what "tracking" means and what obligations are placed on tracking companies when consumers elect to opt-out of tracking. Consumer groups and privacy researchers have published proposals that attempt to define "tracking,"³² but the online advertising industry has not yet committed to respect the header nor follow any of the proposed definitions. For example, some in the industry have suggested that, like the current opt-out system, third parties be allowed permitted to continue to collect information. Others have proposed that third party services should refrain from collecting and retaining any information about consumers if they elect to not be tracked. This latter approach, while more privacy-preserving, may impact advertisers' abilities to deliver even non-targeted advertisements and includes numerous exceptions to tracking which may defeat the spirit of a privacy mechanism.

³¹In particular, domains can "spoof" the first party transactions that are whitelisted in browsers, or effectively act as first parties. This means that they are bypassing any third party-specific controls used in the browser. See Krishnamurthy *et. al.* (previously cited)

³²What Does 'Do Not Track' Mean? A Scoping Proposal" by the Center for Democracy & Technology (2011, Jan 31) from <http://cdt.org/files/pdfs/CDT-DNT-Report.pdf> and

A potential way forward may be to agree upon a definition of “tracking” that balances these conflicting priorities. One of the key components that enables tracking today is the use of unique identifiers. As such, it may be wise to consider a definition of tracking that focuses on these identifiers, in which third party services make a good faith effort to strip any unique identifiers associated with the user, browser or client device making the Web request once the request has been processed and the service delivered. By focusing on the identifiers, these companies would then be free to retain the remaining data associated with the user’s request, providing that it cannot be re-identified (following current best practices in the space). This approach will likely be good for both business and consumers, since it allows businesses to observe how their websites are being used and secure their servers, while preventing the creation of individual profiles.

Finally, it is important to consider whether creating more effective choice mechanisms for consumers may have perverse effects and ultimately drive websites to predicate access to content based on whether or not a consumer has consented to tracking. Websites could require that consumers allow tracking by third parties the website is affiliated with in order to gain access to it’s content. In our original example, WebMD could require that their affiliates, such as DoubleClick, be allowed to track consumers in order to gain access to useful health information on the website. This trend could potentially favor large first parties over smaller, independent sites or allow companies to engage in even more invasive tracking upon receiving affirmative consent. This is not a reason to abandon efforts to improve consumer choice, but certainly a reason for Congress to consider the issue carefully.

D. Conclusion

My research has shown that online tracking is pervasive. It is likely to be much more extensive than users might reasonably expect as they casually browse the Web. Many of these third party tracking activities are carefully tucked away from the view of the average user, and even in cases where the user realizes he is being tracked, the privacy tools he has available are often ineffective at stopping the most advanced forms of tracking.

Consumers need more transparency into who is tracking them online, what data is being collected, and how this data is being used, shared or sold. Today’s technical defenses to online tracking are not able to stop the leading tracking technologies, and consumers often do not have meaningful ways to control them. To be effective, privacy protections for consumers online will likely require both a technical and policy component, working in tandem, and I believe these discussions here today are a great step in making that union a reality.

Internet-related debate involves issues that are deeply technical in nature and I am grateful that this Congressional committee has allowed technologists to participate. Thank you for inviting me to testify here today, and I look forward to helping the committee understand the technical issues that make online tracking such an interesting, yet complex, issue. I will be happy to answer any further questions.

