STANFORD
UNIVERSITY

February 18, 2011

Federal Trade Commission
Office of the Secretary
600 Pennsylvania Avenue, NW
Washington, DC 20580

Re: Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers

Dear Commissioners and Staff:

We commend the Commission's staff for its incisive and farsighted draft report on consumer privacy, and we thank the Commission for the opportunity to provide input in advance of the final report.

We write to share our views on Do Not Track—the result of over half a year of research and outreach to online stakeholders. Additional materials are available at http://donottrack.us, and we would be glad to address any further inquiries the Commission may have.

Sincerely,

Jonathan Mayer

Arvind Narayanan, Ph.D.

Stanford Security Laboratory
Stanford University Department of Computer Science
353 Serra Mall MC 9045
Stanford, CA 94305

The views expressed in this comment are solely those of the authors.

Table of Contents

I. Do Not Track should apply to all third-party tracking, not just behavioral advertising.[1]

Third-party web tracking is pervasive: the average top website incorporates sixty-four independent mechanisms for tracking visitors over time and across websites.[2] Third-party web tracking is also unpopular: numerous studies have shown the vast majority of Americans oppose the practice.[3]

Do Not Track should be a consumer choice mechanism encompassing *all* forms of third-party tracking, whether for advertising, analytics, or any other purpose. As many privacy scholars have remarked, behavioral advertising just happens to be the most visible instance of third-party tracking:

> It is important to note that OBA [Online Behavioral Advertising] has borne the brunt of what might actually be a wider debate about the monitoring of user activity online, and even more widely, the aggregation of personal information for a variety of purposes. Because OBA has a public face in the form of ads, it attracts more attention than the less obviously visible user tracking that is essential to the business of research and analytic companies and certain content delivery firms. That said, the outcome of OBA regulatory efforts could have profound consequences on what counts as legitimate practice in online monitoring and beyond.[4]

The Facebook "Like" button is a prominent example of non-advertising third-party tracking. Facebook can monitor all the pages you visit that incorporate the button, whether or not you click it and whether or not you have an account.[5] Such "social plugins" may be embedded on particularly sensitive sites; England's National Health Service, for example, includes a Like button on its condition pages.[6]

More concerning yet are the multitude of third-party trackers that are completely invisible to users. As the Wall Street Journal's "What They Know" series has explored in depth, whole markets have sprung up around consumer profiling.[7]

Future proofing also cuts against a behavioral advertising focus. Five years ago behavioral advertising was a rarity; the Like button was introduced less than two years ago. It would be a mistake to narrow Do Not Track solely to current instances of third-party tracking.

---

[1] The substance of this section is drawn from Arvind Narayanan, *Do Not Track Isn't Just About Behavioral Advertising*, CENT. FOR INTERNET & SOCIETY (Dec. 20, 2010), http://cyberlaw.stanford.edu/node/6573.
[2] Julia Angwin, *The Web's New Goldmine: Your Secrets*, WALL ST. J., July 30, 2010.
[3] *E.g.*, Joseph Turow et al., *Americans Reject Tailored Advertising and Three Activities that Enable It* 15 (Sept. 29, 2009), *available at* http://ssrn.com/abstract=1478214; Lymari Morales, *U.S. Internet Users Ready to Limit Online Tracking for Ads*, GALLUP (Dec. 21, 2010), http://www.gallup.com/poll/145337/internet-users-ready-limit-online-tracking-ads.aspx.
[4] Solon Barocas & Helen Nissenbaum, *On Notice: The Trouble with Notice and Consent*, PROC. ENGAGING DATA F. (2009), *available at* http://www.nyu.edu/projects/nissenbaum/papers/ED_SII_On_Notice.pdf.
[5] *See* Arnold Roosendaal, *Facebook Tracks and Traces Everyone: Like This!* (Nov. 30, 2010), *available at* http://ssrn.com/abstract=1717563.
[6] *E.g.*, *Seasonal Flu*, NATIONAL HEALTH SERVICE, http://www.nhs.uk/conditions/Flu/Pages/Introduction.aspx.
[7] Angwin, *supra* note 2.

II. Do Not Track should be defined by the scope of third-party tracking.

Do Not Track is a response to third-party tracking; it should cover no more and no less. Defining Do Not Track thus devolves into defining "third-party tracking," which in turn requires definitions of "third party" and "tracking." The following sections propose standards for these definitions and argue the FTC should have authority to interpret the standards into bright-line rules.

A. The distinction between first and third parties should be guided by consumer expectations.

In our view, the privacy distinction between first parties and third parties is shorthand for user expectations. **An entity acts in a first-party capacity if a user reasonably expects to interact with it; it acts in a third-party capacity if a user does not.**[8] Relevant factors for user expectations include domain names, branding, and business relationships. In most cases resolving the standard is straightforward. Some real-world examples:

- A user visits The New York Times' website; Google's Doubleclick ad network collects user data. Google is a third party because it operates at a different domain, uses a different brand, and only has an advertising relationship with The New York Times.
- A user visits Amazon.com; data is collected with the Amazon Web Services platform, located at amazonaws.com. Here Amazon Web Services is a first party because, though domain names differ, Amazon Web Services is functionally a business unit of Amazon.com and is branded as an Amazon.com product.
- A user visits the ESPN website at espn.go.com; Omniture, an analytics provider, collects data at the domain w88.go.com.[9] Omniture is a third party because, though it shares a second-level domain, it is branded independently and only has an advertising relationship with ESPN.

Difficult distinctions arise where entities share more than a purpose-limited business relationship. Some hypotheticals:

- A user visits the Delicious social bookmarking site, acquired by Yahoo! in 2005. Yahoo! embeds tracking content on Delicious. The two share neither domain name nor branding, and Delicious is operated more as an independent business than as a business unit of Yahoo!. On the other hand, the logo reads "Delicious from Yahoo!" and Delicious accepts Yahoo! logins.
- A user visits the Fox News site, which embeds a Wall Street Journal tracking object. The two share neither domain name nor branding and they are operated as separate businesses, but both are owned by News Corporation.

---

[8] The concept of reasonable expectations is well established in American privacy law. *See, e.g.*, Katz v. United States, 389 U.S. 347 (1967).

[9] *See* Balachander Krishnamurthy & Craig E. Wills, *Privacy Diffusion on the Web: A Longitudinal Perspective*, PROC. 18TH INT'L WORLD WIDE WEB CONF. 541, 548 (2009), *available at* http://www2009.org/proceedings/pdf/p541.pdf.

Drawing on the above examples, we submit several observations about the consumer expectations standard. First, the first vs. third party distinction applies to roles, not businesses. In the New York Times example, Google was a third party. But when a user checks her Gmail, Google is clearly a first party. Second, a bright line at domain name boundaries would be both overinclusive (the Amazon Web Services example) and underinclusive (the Omniture example). Third, though close calls will arise, they will be rare and of much narrower scope than the easy calls.

B. Tracking should encompass all data collection, retention, and use.

As explained in the previous section, third-party activities violate a user's reasonable privacy expectations. The user's remedy should, in the first instance, be coextensive with that violation: **Do Not Track should prohibit all data collection, retention, and use.** The online ecosystem is quite complex, and we recognize that there will be a number of exceptional scenarios where privacy concerns must reasonably give way to greater interests. The following section details a standard for arbitrating such exceptions.

C. Exceptions are warranted when narrowly tailored to legitimate commercial interests that substantially outweigh privacy and enforcement interests.

We recognize that exceptions to Do Not Track may be warranted when there is significant commercial need and privacy concerns and enforcement impact[10] are minimal. We believe a two-step standard best captures this policy: **First, legitimate commercial interests must substantially outweigh privacy and enforcement interests, and second, the means of achieving the commercial interests must have no greater privacy and enforcement impact than necessary.**[11] A number of tools are available for minimizing the privacy and enforcement effects of an exception, including client-side storage, dropping parts of data, secure hashing, retention periods, internal business controls, limited sharing agreements, trusted intermediaries, and audits. As guidelines for enacting this standard, we propose several example exceptions:

- Unique browser identification for financial services fraud prevention,[12] provided limited retention periods and strong internal controls. The commercial interest in detecting financial services fraud is significant, user privacy interests are limited to visits to financial services sites, and enforcement impact is minimal given the small number of companies in the online financial services fraud prevention business.
- Retaining protocol logs for security analysis, similarly constrained by retention periods and internal controls. Commercial interest in identifying and tracing security breaches is substantial. Since protocol logs can be identifying, the privacy interest is not insignificant—but there is a countervailing privacy interest in ensuring compromised

[10] Exceptions to Do Not Track may appear to be violations to automated testing tools, requiring additional enforcement resources.
[11] This standard parallels the means-ends test employed by court for strict scrutiny review.
[12] *E.g.*, *DeviceInsight*, 41ST PARAMETER, http://www.the41.com/land/DeviceID.asp; *Fraud Protection*, BLUE CAVA, http://www.bluecava.com/uses/fraud/fraud-protection/.

third parties do not surreptitiously distribute malware.[13] There is no enforcement impact (see Part IV).

- Third-party analytics, provided third parties use first-party cookies and agree to only disclose data to first parties. Analytics are essential to the operation of many online businesses. Privacy interests with respect to the analytics provider are de minimis since the analytics provider is, by agreement, just outsourced first-party analytics. Enforcement impact will be modest owing to the significant concentration of the analytics market.
- Frequency capping by an advertising network, implemented as a short-term, non-unique, human-interpretable cookie. Advertisers and advertising networks have a sizable interest in ensuring a user doesn't see the same ad many times. The privacy interest is negligible since the cookie is not unique. Enforcement will not be impacted.
- Third-party tracking when a user explicitly opts into the practice. For example, if a user logs into Facebook and explicitly enables social plugin tracking, the Facebook Like button can track the user. In these cases there is no privacy interest since the user has agreed to the practice, and by design the enforcement impact will usually be slight.[14]

D. A rulemaking is the appropriate venue for defining bright-line Do Not Track rules.

We recognize that, in practice, the Do Not Track standards must be distilled into clear rules for online businesses to follow. Rules must also be updated over time as technologies and business models shift. In our view these are quintessential regulatory issues; a Do Not Track rulemaking would provide a forum for balancing privacy and commercial concerns. As discussed in greater detail in Part IX, we believe the FTC is the right agency to conduct this rulemaking.

There is a strong temptation to leap into precise definitions of activities encompassed by Do Not Track; several organizations have attempted just this, and we commend their valuable efforts in mapping out possible approaches.[15] We recognize, however, that stakeholders will disagree on many of the fine distinctions inherent to any definition of Do Not Track. **In our view the next step is to authorize a rulemaking guided by high-level standards, *then* identify specific bright-line rules and technical standards in the rulemaking.**

III. Do Not Track should be implemented as an HTTP header.

A third-party web tracking choice mechanism should possess these characteristics:

- **Comprehensive**: choices can apply to all third-party entities and tracking mechanisms.
- **Persistent**: after configuring privacy choices, a user does not have to reconfigure them.
- **Simple to Use**: the choice mechanism is convenient and requires no special knowledge.

---

[13] *See, e.g.*, Kim Zetter, *Google DoubleClick Caught Serving Malicious Ad*, WIRED THREAT LEVEL (Dec. 10, 2010), http://www.wired.com/threatlevel/2010/12/doubleclick/.

[14] The Do Not Track header supports opting into tracking. Jonathan Mayer, *Minor Updates to the Do Not Track Header*, CENT. FOR INTERNET & SOCIETY (Jan. 27, 2011), http://cyberlaw.stanford.edu/node/6597.

[15] *E.g.*, CENT. FOR DEMOCRACY & TECH., WHAT DOES "DO NOT TRACK" MEAN? (2011), *available at* http://cdt.org/files/pdfs/CDT-DNT-Report.pdf. A number of stakeholders expressed views on defining Do Not Track at a recent event at the University of California, Berkeley. Browser Privacy Mechanisms Roundtable (2011), *available at* http://www.law.berkeley.edu/files/bclt_Browser-Privacy_Transcript.pdf.

- **Verifiable**: compliance with the choice mechanism can be objectively determined.
- **Granular**: a user can make selective choices about third-party tracking.
- **Tailored**: the choice mechanism solely affects third-party tracking; it has no collateral effects.

There are currently three major technology proposals for responding to third-party privacy concerns. The first proposal is the Do Not Track HTTP header, a signal to web services of a user's tracking preferences. Mozilla recently implemented the Do Not Track header in Firefox 4, and a number of Firefox extensions also support the header. We are collaborating with Mozilla and the Center for Democracy and Technology to standardize the header in an IETF Internet-Draft.[16] The second proposal relies on block lists of web resources the browser should not to load. Firefox, Chrome, and Safari support block lists via extensions; Microsoft is adding in-browser block list support with Internet Explorer 9. The third proposal is the Network Advertising Initiative's model of a per-company opt-out cookie. Extensions like TACO and Google's Keep My Opt Outs refine the model by persisting cookies and preventing accidental deletion. The following table reviews each proposal against the design criteria:

Design Criterion	Do Not Track Header	Block List	NAI Opt-out Cookies
Comprehensive	A user can send the header to all third parties, though some may be outside U.S. jurisdiction.	Only covers resources on the block list.	Only covers the ≈70 NAI members.
Persistent	Preferences are browser-specific.	Preferences are browser-specific. A user must ensure their chosen block lists are still actively maintained and have not relocated.	Preferences are browser-specific. Cookies may expire or be accidentally deleted; extensions like TACO and Keep My Opt Outs will retain them.
Simple to Use	The user either activates a browser preference or configures an extension.	The user either activates a browser preference or configures an extension.	The user must either discover the NAI site or install an opt-out cookie extension.
Verifiable	See Part IV.	No verification is needed for blocked resources. Determining what should be blocked requires the same verification as the header and cookies; see Part IV.	See Part IV.

[16] *See* Mayer, *supra* note 14.

Granular	A user interface issue; no protocol limitation.	Preferences apply by domain or URL.	Preferences apply by business.
Tailored	The Do Not Track header is simply an expression of preference.	The block list prevents listed content from being loaded at all.	Opt out cookies are simply an expression of preference.

We believe the choice between the Do Not Track header and the NAI opt-out cookie model is clear: both are mechanisms for expressing a user preference, but Do Not Track is more comprehensive and persistent, simpler to use, and more granular. In our experience hesitation to support the Do Not Track header has been reflective of objections to the Do Not Track policy, not the technology.[17] Existing opt outs are generally far less restrictive; NAI members, for example, must only commit to not using tracking data to target advertisements.[18]

We view block lists as a complementary measure to the Do Not Track header, and we commend Microsoft for building block list support directly into Internet Explorer. Block lists provide privacy protection in the near-term, and in future will provide defense in depth against third parties that refuse to honor the Do Not Track header, especially those outside U.S. jurisdiction. But block lists are far from an ideal solution to third-party privacy choice.

First, users must select a trusted list (or combination of lists) to make the correct blocking decisions. In practice lists vary substantially in quality. EasyPrivacy, for example, covers a fairly large number of third parties (notably not Google Analytics or Facebook). TRUSTe's list, on the other hand, is primarily a *whitelist* of TRUSTe customers, including data aggregators and behavioral advertisers. We have serious doubts that consumers will be able to easily judge the merits of block lists.

Second, lists require updating. If a list ceases to be actively maintained, relocates, or is even just formatted incorrectly, its users may be left in the lurch. Explaining these problems is a user interface challenge; many block list implementations do not even attempt to notify the user of block list problems.

Last, block lists are not a tailored solution; they completely block listed content, even when it could be delivered without third-party tracking. This feature has far-reaching business implications: Opting out of tracking by Facebook's Like button necessitates blocking the Like button—even though it could be trivially delivered without tracking. More concerning yet, opting out of tracking by an advertising network would, in general, block *all* ads from the network—not just require the network to serve non-behavioral ads (and otherwise comply with the Do Not Track policy).

[17] To confirm that reading the Do Not Track header is trivial, we wrote sample code for a number of popular web application platforms. We found the software development for each platform was a matter of minutes. *Do Not Track: Web Application Templates*, http://donottrack.us/application.html.

[18] *See* Rainey Reitman, *Mozilla Leads the Way on Do Not Track*, ELECTRONIC FRONTIER FOUND. (Jan. 24, 2011), https://www.eff.org/deeplinks/2011/01/mozilla-leads-the-way-on-do-not-track.

Tailored choice mechanisms allow for market responses to consumer privacy choices. The Do Not Track header would provide incentives for third parties to both develop privacy-preserving technologies and increase transparency in the interest of encouraging users to allow tracking. For example, advertising networks might standardize a protocol for tracking-free interest-targeted advertising,[19] and social networks might provide clearer privacy statements about their social plugins.

IV. Do Not Track is verifiable.

We envision two technical approaches to verifying Do Not Track compliance. First, most tracking at the application layer[20] can be detected by modifying a browser to report tracking-related activity.[21] If after receiving a Do Not Track header third-party embedded content sets a unique cookie or lists the browser's plug-ins, the third party may be violating Do Not Track. Second, behavioral advertising can be identified by monitoring ads for interest targeting.[22] Data should be sourced using both crawling and crowdsourcing to ensure comprehensive coverage of top websites and a real-world sample of observations. We are beginning development of a Do Not Track verification system with colleagues in the Stanford Security Laboratory, and we look forward to sharing our work in the coming months.

We note that verification does require some measure of human follow-up. If a third party is engaged in tracking covered by an exception, for example, it will likely appear improper to an automated system. While we are confident the degree of human intervention necessary will not be inordinate, we anticipate forming a more precise estimate in the course of our verification work.

V. Do Not Track is unlikely to harm advertising-supported businesses.[23]

A. Do Not Track would only affect a sliver of the online advertising market.

A brief overview of online advertising: Suppose you operate a high-end Napa winery and decide to run an ad. You might place your ad on a specific website ("first-party advertising"), or you might arrange your ad with an advertising network that spans thousands of sites ("third-party advertising"). Here's a sample of how you might target your ad:

---

[19] *See* Vincent Toubiana et al., *Adnostic: Privacy Preserving Targeted Advertising*, PROC. 17TH ANN. NETWORK & DISTRIBUTED SYS. SECURITY SYMP. (2010), *available at* http://crypto.stanford.edu/adnostic/adnostic-ndss.pdf; Matthew Fredrikson & Ben Livshits, *RePriv: Re-Envisioning In-Browser Privacy* (Microsoft Research Technical Report MSR-TR-2010-116, 2010), *available at* http://research.microsoft.com/pubs/137038/tr.pdf.

[20] Collection, retention, and non-advertising use of protocol layer information will be challenging to detect by automated means. Strong internal controls and auditing may be appropriate for the largest third parties.

[21] *See* Dongseok Jang et al., *An Empirical Study of Privacy-Violating Information Flows in JavaScript Web Applications*, PROC. 17TH ACM CONF. ON COMPUTER & COMM. SECURITY 270 (2010), *available at* http://cseweb.ucsd.edu/~hovav/dist/history.pdf.

[22] *See* Saikat Guha et al., *Challenges in Measuring Online Advertising Systems*, PROC. 10TH ANN. CONF. ON INTERNET MEASUREMENT (2010), *available at* http://saikat.guha.cc/pub/imc10-ads.pdf.

[23] The substance of this section is drawn from Jonathan Mayer, *Do Not Track Is No Threat to Ad-Supported Businesses*, CENT. FOR INTERNET & SOCIETY (Jan. 20, 2011), http://cyberlaw.stanford.edu/node/6592.

- Contextual Advertising: Your ad appears on pages about wine.
- Demographic Advertising: Your ad appears on pages whose visitors tend to be wealthy.
- Behavioral Advertising: Your ad appears to users who have viewed a number of pages about wine.[24]
- Search Advertising: Your ad appears on search result pages for the query "wine."
- Placement Advertising: Your ad appears on particular pages.
- Social Network Advertising: Your ad appears to social network users who have listed "wine" as an interest.

Of these myriad modes of advertising, Do Not Track would only affect one: third-party behavioral advertising, because it incorporates third-party tracking. And that accounted for, at most, just 4% (less than $1B) of 2009 U.S. online advertising expenditures.[25] While the use of third-party behavioral advertising is rapidly growing, so is the online advertising market; projections place behavioral advertising at only 7% of the U.S. online advertising market in 2014.[26]

B. Do Not Track would only affect a new segment of the online advertising market.

Not only is third-party behavioral advertising a small piece of the online advertising market, it's also a new piece. Behavioral advertising accounted for a negligible share of online advertising until roughly 2007.[27] Countless ad-supported online businesses launched and thrived before then.

C. Do Not Track would cap—not eliminate—third-party behavioral advertising.

Do Not Track is an opt-out mechanism; uptake is likely to be far from complete. Two helpful benchmarks: After seven years of a permanent opt out, fewer than half of U.S. phone numbers are on the Do Not Call registry.[28] And after four years of availability, fewer than 3% of Firefox users have installed its most popular add-on.[29]

---

[24] For simplicity this section glosses over behavioral retargeting, a small subset of behavioral advertising.
[25] Memorandum from the Democratic Staff of the House Subcomm. on Commerce, Trade, & Consumer Prot. to the Members of the House Subcomm. on Commerce, Trade, & Consumer Prot. 3-4 (Nov. 30, 2010), *available at* http://democrats.energycommerce.house.gov/documents/20101201/Briefing.Memo.12.01.2010.pdf.
[26] eMARKETER, THE GLOBAL MEDIA INTELLIGENCE REPORT NA-6-10 (2010), *available at* http://www.emarketer.com/Reports/All/Emarketer_2000722.aspx; David Hallerman, *Is Behavioral Targeting Outmoded?*, eMARKETER BLOG (Mar. 12, 2010), http://www.emarketer.com/blog/index.php/behavorial-targeting-outmoded/. These figures reflect both first- and third-party behavioral advertising. They should be taken as an upper limit on the market size for third-party behavioral advertising.
[27] Hallerman, *supra* note 26; Press Release, eMarketer, Behavioral Targeting Poised for Growth (June 16, 2008), *available at* http://www.emarketer.com/Article.aspx?R=1006384.
[28] INT'L TELECOMM. UNION, MEASURING THE INFORMATION SOCIETY 104 (2010), *available at* http://www.itu.int/ITU-D/ict/publications/idi/2010/Material/MIS_2010_without_annex_4-e.pdf; Press Release, Federal Trade Comm'n, National Do Not Call Registry Tops 200 Million Phone Numbers (July 27, 2010), *available at* http://www.ftc.gov/opa/2010/07/dnc.shtm.
[29] MOZILLA, THE STATE OF MOZILLA (2010), *available at* http://www.mozilla.org/foundation/annualreport/2009/; *Adblock Plus Statistics*, MOZILLA, https://addons.mozilla.org/en-US/statistics/addon/1865.

D. Advertisers might not reallocate their ad dollars.

Websites that host third-party ads usually receive a fixed share of revenue; they earn more only if advertisers spend more. Do Not Track would thus impact advertising revenue only if it caused advertisers to reallocate online ad dollars. But that would happen only if advertisers have a strong preference for third-party behavioral advertising. There's some evidence that advertisers don't: despite the growing availability of third-party behavioral advertising over the past several years, advertisers haven't rushed to adopt it. In fact, U.S. online advertising revenues grew at an average annual rate of only 3.4% between 2007 (when behavioral advertising first caught on) and 2009.[30]

E. There's a technology fix: interest-targeted advertising without tracking.

Third-party behavioral advertising incorporates tracking to discover a user's interests. But **interest-targeted advertising can be achieved without tracking**. Under one alternative model, the web browser learns a user's interests, and then passes those interests to an advertising network. A number of research and commercial efforts do just this.[31]

F. Advertising-supported businesses could ask—and possibly require—Do Not Track users to allow third-party behavioral advertising.

Do Not Track is not all-or-nothing; users who have opted out can opt back into third-party tracking on specific sites or with specific trackers. So even if third-party behavioral advertising were an important revenue source for ad-supported businesses, even if enough users opted out to have an impact, even if advertisers were inclined to pull their ad dollars, and even if alternative technologies for interest-targeted advertising weren't available, a business would *still* have an easy remedy: ask—or, if allowed, require[32]—visitors to disable Do Not Track on the site. The proposal is about increasing privacy choice and transparency, not restricting online business practices.

VI. Do Not Track should be extended to mobile platforms.[33]

Third-party tracking is proliferating on mobile platforms;[34] such tracking implicates the same privacy concerns as third-party tracking on the web, and likewise warrants a Do Not Track choice mechanism.

---

[30] INTERACTIVE ADVERTISING BUREAU, IAB INTERNET ADVERTISING REVENUE REPORT (2010), *available at* http://www.iab.net/media/file/IAB-Ad-Revenue-Full-Year-2009.pdf. One possible reason: behavioral ads may be only a marginally better deal for advertisers. In Q4 2009, a behavioral ad was 2.1x as effective as the average online ad—but it cost 2x as much. *Behavioral Targeting Doubles Ad Effectiveness*, EMARKETER (Mar. 19, 2010), http://www.emarketer.com/Article.aspx?R=1007599.

[31] *See supra* note 19 and accompanying text. Google Ads Preferences is another example of client-based interest-targeted advertising.

[32] *See infra* Part VII.

[33] These comments equally apply to tablet platforms and other Internet appliances. We focus on mobile platforms as but a convenient example.

[34] Scott Thurm & Yukari Iwatani Kane, *Your Apps Are Watching You*, WALL ST. J., Dec. 17, 2010.

From the technical perspective, the mobile browsing ecosystem differs from desktop browsing in a number of ways that are pertinent to Do Not Track. First, mobile browsers lack a strong identity. They are often named generically (such as "web browser" on Android), and users rarely install a non-default browser. Second, mobile browsers lack sophisticated customizability. Third, much of the third-party tracking on mobile devices happens in apps, outside the context of a traditional browser.

The Do Not Track HTTP header model can be easily adapted to mobile platforms. Instead of a universal browser setting, Do Not Track should be a platform-wide preference that adds a Do Not Track header to all HTTP requests and provides a Do Not Track signal to apps. Much as embedded third-party web trackers would check for the Do Not Track header, embedded third-party mobile app trackers would check for the Do Not Track platform preference. Paralleling the granularity of the header, apps should be able to interact with the platform to request an exception from Do Not Track. As for verification, since third-party mobile tracking is heavily concentrated the problem is much simpler than in the desktop browser context.

Turning to policy, the first vs. third party distinction also seamlessly transitions to the mobile context. An app is a first party; a behavioral advertising network embedded in the app would be a third-party since its presence violates reasonable privacy expectations.

VII. The Commission should adopt a wait-and-see approach to tiering.[35]

A tiered web is one in which sites require users to disable Do Not Track to access certain features or content. We believe that tiering is unlikely, if it does occur it could have positive or negative effects, and there are many possible policy responses. We therefore recommend that the Commission adopt a wait-and-see approach.

We believe widespread tiering is unlikely to occur for two reasons. First, as a comparative matter, ad blocking—which has a far greater per-user impact on advertising revenue than Do Not Track—is generally tolerated. To our knowledge there are no sites that tier service for users of ad blocking technology.[36] Second, there is substantial stigma associated with tiering. Asking a user to wade through ads to reach free content is now par for the course; asking a user to explicitly trade their privacy for free content would be socially unpalatable.

Were tiering to occur, it could be beneficial. For example, it could jumpstart a trend towards readable privacy policies, or lead to innovative micropayment business models. But tiering could also be harmful if overused, especially if consumer confusion led users to simply disable Do Not Track to reach content.

[35] The substance of this section is drawn from Arvind Narayanan, *"Do Not Track" Explained*, 33 BITS OF ENTROPY (Sept. 20, 2010), http://33bits.org/2010/09/20/do-not-track-explained/.

[36] The news site Ars Technica prohibited ad-blocking users for twelve hours. The move was widely criticized and never repeated by the site. Ken Fisher, *Why Ad Blocking Is Devastating to the Sites You Love*, ARS TECHNICA (Mar. 6, 2010), http://arstechnica.com/business/news/2010/03/why-ad-blocking-is-devastating-to-the-sites-you-love.ars.

In response to tiering, the Commission could adopt a variety of policy positions. At the poles, of course, are allow and prohibit tiering. Moderate positions could include requiring payment-based alternatives to opting back in and discouraging, but not prohibiting, tiering.

Given the uncertainty around tiering in practice and the broad range of policy options, the Commission should reserve its position on the issue and return to it in future.

VIII. The Commission should adopt a wait-and-see approach to international third parties.

Several domestic third parties have expressed a concern that Do Not Track would render them unable to compete with international third parties. We are skeptical this issue will arise given the concentration of domestic third parties on top sites. If competitiveness becomes a problem, international third parties doing business with domestic sites could be required to contractually follow Do Not Track.

IX. The FTC should call for legislation authorizing it to define and enforce Do Not Track.

When we initially articulated our vision for Do Not Track, we noted it could be implemented voluntarily or through industry self-regulation. We now believe legislation and FTC involvement are necessary.

In our view, third-party opposition to Do Not Track at a technological level is largely a façade. The HTTP standard is designed to allow flexible signaling with headers; Internet Explorer alone uses at least eight proprietary headers.[37]

The substantive disagreements about Do Not Track arise from policy. A number of third parties oppose a stringent definition of third-party tracking. Given the diversity of online business models and businesses Do Not Track would affect, and given the consensus-based nature of the relevant trade associations, we believe voluntary comprehensive adoption will not occur.

The Federal Trade Commission is the right agency to define and enforce Do Not Track. The Commission's growing technical staff lends it unique domain expertise for defining Do Not Track, and its capacity for and experience with consumer protection actions prime it to enforce Do Not Track.

[37] Eric Lawrence, *Internet Explorer and Custom HTTP Headers*, ERICLAW'S IEINTERNALS (June 30, 2009), http://blogs.msdn.com/b/ieinternals/archive/2009/06/30/internet-explorer-custom-http-headers.aspx.