Comments of Jim Brock, Founder and CEO, PrivacyChoice LLC

The following comments are respectfully submitted to the Federal Trade Commission in response to A Preliminary FTC Staff Report on "Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers."

Background

Founded in early 2009, the PrivacyChoice mission is to make online privacy easier for consumers and websites. PrivacyChoice offers a suite of privacy tools for tracking awareness (TrackerScan), tracking control (TrackerBlock), ad targeting opt-outs (PrivacyMark), automated website privacy disclosure (PrivacyWidget), and social-network privacy protection (PrivacyCheck and Disconnect). Over 400,000 web users have managed their online privacy preferences through PrivacyChoice. The following comments include excerpts from the PrivacyChoice Blog, which is available in full at http://blog.privacychoice.org.

Summary of Comments

- **Consumer choice is best served by combining the "Do Not Track" approaches taken by Mozilla Firefox and Microsoft Internet Explorer. Do-Not-Track should be integrated into the enhanced notice-and-choice framework, so that persistent choices can be activated directly from web-based notices.**
- **Do-Not-Track implementations should shift the burden to tracking companies to specifically identify non-tracking data collection.**
- **Implementing Do-Not-Track through browser headers is useful, but does not provide sufficient verification, context or choice for consumers.**
- **Simple economic modeling indicates that greater availability of do-not-track choices for web users may negatively affect industry ad revenues, depending part on the availability and prominence of choices. However, reduced revenues from do-not-track elections may be offset by increased engagement with online profile management.**
- **While not as persistent as a browser-based do-not-track option, the current opt-out framework can be significantly improved to provide greater assurance to web-users that they will not be tracked after making an opt-out election.**
- **An analysis of actual do-not-track elections made through the PrivacyChoice service indicate that a significant percentage of web users are less likely to block tracking by companies subject to industry oversight processes.**

Reconstructing Do Not Track

Summary: Consumer choice is best served by combining the “Do Not Track” approaches taken by Mozilla Firefox and Microsoft Internet Explorer. Do-Not-Track should be integrated into the enhanced notice-and-choice framework, so that persistent choices can be activated directly from web-based notices.

The major browser makers have now proposed three very different approaches to give users control over online tracking. *Microsoft IE9*'s "Tracking Protection Lists" provide direct blocking of tracking interactions based on lists curated and hosted by independent companies. *Mozilla's Firefox* gives the user an option to transmit a constant browser signal asking not to be tracked. *Google's Chrome* relies on the current opt-out cookie framework, using a browser extension to make them permanent for companies adopting self-regulatory rules.

These approaches can be compared in light of five factors: **Simplicity**, **Findability**, **Certainty**, **Durability** and **Versatility**. Each browser's approach supports these objectives in different ways, and are not technically inconsistent. In hybrid form, features from the Mozilla and Microsoft approaches can support a Do Not Track framework that provides meaningful user choices while still supporting the Web advertising economy.

Findability

Advantage: Microsoft

The Do-Not-Track choice for Firefox appears in the "Advanced" tools menu (not the "Privacy" tab). In Chrome, you need to find their extension on the Chrome Extensions site. Neither of these can be initiated from a web interaction. Unless the option is presented at first install of the browser or startup for a new session, web users otherwise must become aware of them and seek them out.

Microsoft's approach, by contrast, allows a choice to be made from within any webpage where a Tracking Protection List is hosted. List curators can host and promote their own approaches. This also allows the preference setting process to be available as part of the enhanced notice and choice framework being offered by the Digital Advertising Alliance. By making that connection, the tracking-control decision can be available in context of the ads and websites where tracking happens.

Simplicity

Advantage: Firefox

Simplicity matters for both web users and tracking companies.

For web users, simplicity can be measured by the number of clicks required; the number and complexity of choices offered; and the ongoing effort required to keep choices in place. None of the approaches makes the do-not-track the default setting, which would be the simplest for users but with significant disruption to the online ad economy.

Firefox's choice is found is found in browser controls, which are opened with two clicks and requires one more click to make a selection. Because the choice applies to any tracking company (current and future), there is no updating required.

Microsoft's choice is made not in a browser, but from a link or button on a website. It requires two clicks, one to start the process and one to confirm. List updates are handled automatically in the background. The approach is more complex is the sense that consumers will have multiple choices from different providers; these may be offered or endorsed by familiar organizations, perhaps simplifying a user's decision.

Chrome's choice starts on the extension download page, and requires one confirming click. No restart of the browser is required and the choice is immediately effective. However, because the list is not automatically updated in the background, the user must approve an extension update each time a new company is added to the list.

For tracking companies, Microsoft's basic approach is the simplest to implement; because it works in the browser, companies don't need to do anything to effect an opt-out choice. More server-side work is required for the Firefox or Chrome approaches, where data collection and use practices are modified for opted-out users. To allow external verification, companies may need to segregate tracking and non-tracking actions on separate subdomains or paths, with separate cookies designated for each.

To the extent curators of Microsoft lists want to allow non-behavioral interactions, like the serving of contextual ads, there will be a burden on tracking companies to segregate these interactions. That could also require server changes and independent auditing.

Certainty

Advantage: Microsoft

Any Do-Not-Track approach requires a definition of "tracking" activity. Microsoft leaves this up to list curators. Firefox and Chrome depend on the tracking companies themselves to make this determination, presumably with guidance from industry organizations or regulators. In each case, certainty for web users depends on how well the standard is communicated at the point of choice.

However tracking may be defined, the approaches different significantly as to the degree of certainty users have about whether their choices are respected. Microsoft's approach provides the most certainty, by actually blocking browser interactions that can be used for tracking. The

Firefox and Chrome approaches do not necessarily block the collection of behavioral information; they rely on tracking companies to see and honor the preference.

Under any approach, tracking companies may still want to collect ad-serving information (e.g. how many times an ad has been shown), but not behavioral data (e.g. which pages were visited). This information may be still associated with a unique cookie identifier. Tracking companies can label their tracking domains or cookies to provide assurance as to how they are used (e.g. those activities can be conducted only on "no-tracking.adcompany.com" and cookies can be labeled "no-tracking" in their name or text). This has the advantage of creating a more explicit promise from marketers to consumers as to what data may still be collected. Compliance can be tested externally through user panels (only "no-tracking" cookies should be seen when the user is opted-out), as well as independent audits of internal practices.

Versatility

Advantage: Microsoft

"Versatility" considers support for different user choices beyond a blanket preference against tracking. Actual experience on privacychoice.org shows that over 30% of consumers tend to make tracking choices that are more refined than a blanket blocking choice. If consumers can choose selectively to accept more responsible and accountable tracking, this will encourage better privacy practices. Versatile choices also allow websites and ad firms to make the case for more targeted marketing, and to connect it to free content and services.

Microsoft's framework is the only one designed to provide both "Allow" and "Disallow" choices. The Firefox and Chrome approaches could be supplemented with customized choices that override a global do-not-track selection. A user could expressly allow targeting by specific companies or kinds of companies, or potentially only on specific websites, with this selection indicated by an overriding cookie.

Durability

Advantage: Firefox

All three approaches give users "set and forget" choices, which endure even when browser histories are cleared. This corrects a major flaw in the current notice-and-choice framework.

However, durability depends not only on the permanence of settings, but also on how effectively a global choice continues to work as tracking companies come and go. IE9 calls on curators to update Tracking Protection Lists, which means the user doesn't need to do anything; but this creates a dependency on the curator to keep things up to date.

Chrome handles this by pushing extension updates as new tracking companies join the self-regulatory program. Because Firefox enables a global mandate, no updating or curation by the user is necessary.

An Ideal Approach

Each browser's approach to Do-Not-Track has strengths, weaknesses and dependencies. An ideal approach could combine the best attributes of the Microsoft and Firefox approaches:

- **A binary, global do-not-track signal which must be respected as to activities commonly defined as "tracking."** This provides *simplicity* and *durability* for the broadest set of web users, provided that "tracking" can be appropriately defined.
- **Settings to control tracking interactions directly in the browser.** This provides *certainty* that choices are honored, with less dependency on server-side compliance.
- **The ability for any selection to be made in an web interaction, rather than within the browser setting menus.** This makes choices *findable* in a context where users can best understand their purpose and effect.
- **Choices to selectively allow or disallow tracking at the company or website level, as a complement to global settings.** This provides *versatile* choices to afford web users the greatest benefit from their online profile and encourages value exchange with web providers.
- **Independent audits of tracking practices which cannot be externally verified.** This allows marketers to continue to use non-behavioral data without compromising *certainty* for consumers

Tracking companies should identify non-tracking activities

Summary: Do-Not-Track implementations should shift the burden to tracking companies to specifically identify non-tracking data collection.

Following up on our December announcement, we're now offering "tracking protection lists" for the new version of Internet Explorer. These lists activate internal tracking controls in IE9 that are easily installed, verifiably effective and can be activated in the context of any webpage or privacy notice. This new IE9 functionality is an important step toward a tracking-choice framework that empowers consumers while coexisting with the web ad economy. Our experience so far in developing tracking protection lists suggests an important requirement for the success of browser-based tracking control.

Tracking Protection Lists are powerful because they selectively block any interaction with ad delivery companies, even the serving of an image ad. But TPLs were not created for the purpose of blocking ads, but only to block behavioral data collection that may accompany ad delivery. The challenge for curating Tracking Protection Lists is in differentiating those interactions. There is no obvious way for a curator to block tracking but allow the display of contextually (versus behaviorally) targeted advertising, or to allow collection of non-tracking data like the number of times an ad has shown.

Only the tracking company knows whether a particular interaction involves behavioral data collection. They are in the best position to indicate when an interaction does not involve behavioral data. To enable this, we include the following logic in the PrivacyChoice TPLs:

- All interactions with a tracking company are presumed to be behavioral, and therefore *disallowed*, but
- Any interaction is *allowed* if the URL includes the string, "not_tracking"

This approach allows tracking and non-tracking activities to be sorted by the companies themselves, rather than by list curators. It also creates an express affirmation to the web user about how their data will or will not be used. This enhances enforcement, insofar as it would be deceptive to label an action as "not tracking" if it is actually otherwise.

The same self-identification approach can be applied when the Do-Not-Track election is implemented through headers, as in Firefox. Companies that recognize the header may still want to collect data through cookies for non-behavioral purposes. When they do so, a "not tracking" indicator should be part of each interaction.

In either case, this approach makes do-not-track more verifiable. It's simple to test whether opted-out browsers experience any interactions that lack the "not tracking" indicator. When implemented as a subdomain (like "not-tracking.adcompany.com"), the user will see this in their browser cookie list, and can easily spot companies that haven't provided the additional assurance.

Just as no companies currently recognize the Firefox Do-Not-Track header, no companies currently use "not tracking" strings. The power of the IE9 approach is that, unlike headers, by presumptively blocking they truly shift the burden to the ad companies to identify and control behavioral tracking activity, while still accommodating other ad targeting and delivery needs that are properly identified.

Do-Not-Track headers in browsers: Six concerns

Summary: Implementing Do-Not-Track through browser headers is useful, but does not provide sufficient verification, context or choice for consumers.

It's great to see smart minds turned to the question of how to empower consumers when it comes to online tracking, so you have to appreciate the announcement of donottrack.us. This effort from Stanford is giving new life to the notion of modifying browsers to transmit a "do-not-track" preference with each header. When compliant tracking firms see the header, they would be required to recognize the opt-out preference and, presumably, ignore any other information transmitted with that request.

The chief benefit of this approach is that it is universal and potentially more scalable than collecting opt-out cookies on a user's computer. Scalability is an important concern, particularly as the tracking company universe expands from a few hundred ad companies to thousands of brands with their own pools of user cookies.

Here are the issues:

1. *Adoption by Browser Makers.* Like any other browser based solution, it requires adoption by the browser companies. This seems unlikely in the absence of a new law that requires it. The FTC today doesn't have the authority to order it, and browser functionality seems like a difficult thing for Congress to legislate directly.

2. *Opt-out Framework Still Required.* Even if adopted as standard equipment by one or more browser makers, consumers on unsupported browsers still need to be able to opt-out. The system would not become more simple.

3. *What would be the default?* Even if it were adopted as standard equipment by all browser makers, the default settings would largely determine consumer awareness and adoption. It's hard to see the industry accept "off" as the default setting. The worst outcome would be a powerful but buried feature that no one knows about.

4. *No connection or context.* In the current opt-out framework, the consumer's opt-out decision can be made directly and immediately from the notice of tracking. Because it's a browser setting, there's no simple way to connect header selection with the ad and online notice that provide valuable context.

5. *Inferior to blocking.* Compared with actually blocking interactions between the browser and the tracking company, an approach based on headers is less verifiable for the user, since it does not prevent unique identifiers from being written or read. If you're going to modify your browser to control tracking, you should modify it not only for compliant companies, but also those who don't comply. Given that less than a third of tracking companies are enrolled in the

self-regulatory system now, incomplete coverage is likely to be an issue for a long time to come.

6. *Less choice*. The donottrack.us header is elegant because it is universal. But as the primary means to control tracking, that actually restricts choice in important ways. Consumers should have the ability to control which companies to block based on policies, oversight or even whether a tracking company has given them an incentive not to do so. In this way, donottrack.us is at odds with the consumer's opportunity to influence and even have a stake in tracking.

Given these challenges, I'm not sure that the donottrack.us approach would meaningfully enhance the consumer experience compared to the current framework, flawed as it is. The current system, with some simple enhancements and much greater visibility to consumers, still seems like the right starting point. From there, browser enhancements that actually block tracking -- hopefully built in and visible -- provide the best upgrade for privacy-concerned consumers.

What’s the Impact of Do-Not-Track: A Simple Economic Model

Summary: Simple economic modeling indicates that greater availability of do-not-track choices for web users may negatively affect industry ad revenues, depending part on the availability and prominence of choices. However, reduced revenues from do-not-track elections may be offset by increased engagement with online profile management.

The last few weeks were full of news in the debate around online tracking privacy. Not only did the FTC endorse the notion of providing a durable "Do Not Track" option for consumers, Microsoft promptly announced that they are building it into the next version of Internet Explorer, and it includes not only durable opt-out settings, but also opt-in settings as well. Add these developments to the self-regulatory approach already under construction, and a more complete framework for tracking choice is coming into view.

What will be the economic impact of this new framework on the ad business? With my simple model, you can use your own assumptions to approximate how much revenue may be lost -- and how much may be gained -- by providing enhanced tracking choice across the ad ecosystem. The impact depends on factors like what percentage of users click on ad notices, how many of them make a tracking choice, and to what extent they are invited not simply to opt-out, but instead to engage with and improve their own ad profile.

A simple model

The purpose of this simple model is to estimate the potential revenue impact of providing enhanced tracking choices (including Do Not Track). The model is based on published market forecasts for 2014; the impact will likely grow thereafter as behavioral methods become more effective and prevalent.

See the spreadsheet model here (static)
Download the model as an Excel file (to edit)

The model assumes that tracking choices are implemented through integration of simple and durable browser settings (i.e. all browsers implement an approach like Microsoft's Tracking Protection Lists); and that choices are presented through in-ad notices anchored to icons in advertisements and website notices, as provided in the Digital Advertising Alliance framework.

The model assumes there will be both affirmative and negative tracking options: "Do Not Track" elections reduce behavioral ad revenue; but theoretically, "Opt-In" interactions increase behavioral ad revenue through more accurate and extensive targeting profiles. To the extent some value-exchange is required for Opt-In, particularly for lesser known or trusted companies, the model includes a cost factor for inducement.

Significantly, the model assumes that a Do Not Track election by any user does not impair the non-behavioral (i.e. contextual) advertising opportunity as to that user. This means tracking choices are implemented in a refined way, with ad and data companies segregating behavioral and non-behavioral targeting methods. See the spreadsheet for some other caveats.

Key factors

Here are some of the key factors built into in the model including rationales for some initial settings (download the spreadsheet to try with your own assumptions):

- *How many users will notice and then click on a behavioral targeting icon in an ad or on a website?* (Set at 15%, which is a icon-click rate inferred from recent research published by Better Advertising.)
- *What percentage of electing users will choose Do Not Track?* (Arbitrarily set at 20% of users who make it to the choices. Comments from Google indicate that of those who encounter Google's tracking profile manager and opt-out interface, just under 7% elect to opt-out of tracking, 28% edit their profile and the remainder do nothing.)
- *What percentage of electing users will choose to Opt-In?* (Arbitrarily set at 40% of users viewing choices, or two times the number opting-out; with Google's ratio being 4x).
- *How many clicks will it take to make a choice election and what is the conversion rate at each step?* (Arbitrarily the model assumes one interstitial screen with a 50% drop-off, and an 80% completion rate on the browser election process.)
- *Is there a multiplier of targeting value if user Opts-In and how much is it?* (Arbitrarily set at 50% above average value.)
- *What is the cost to induce Opt-In elections?* (Arbitrarily set at zero.)

Observation

With these starting assumptions, it's not hard to imagine that Do Not Track elections could reduce behavioral ad revenue by over $100 million in 2014. However, to the extent that Opt-In choices are compelling to consumers and increase targeting value, they can be modeled to defray or even exceed the cost of Do Not Track.

Keep It Simple: Opt-out cookies should always overwrite tracking cookies

Summary: While not as persistent as a browser-based do-not-track option, the current opt-out framework can be significantly improved to provide greater assurance to web-users that they will not be tracked after making an opt-out election.

The renewed notice-and-choice framework being rolled out for behavioral advertising has an important shortcoming: Opt-out cookies only serve to turn off the delivery of behaviorally targeted ads, they don't promise to turn off data collection in the first place.

The problem is that many companies provide opt-out cookies that are written separately from their tracking cookies. This means a tracking company may be collecting data with one cookie, while being told by another cookie not to apply it. Because the user still remains uniquely identifiable to the tracking system, the risk of technical (or ethical) failure remains. In that sense, the opt-out framework doesn't really answer a central privacy concern.

Perhaps the solution is simple:

When a user opts out, a non-unique opt out cookie should always overwrite each cookie that stores behavioral information.

When this rule is followed, an opted-out consumer can be assured that there's no unique tracking cookie on their computer that could be used for data collection. Several opt-out cookies, notably the Google DoubleClick cookie, already operate this way by replacing the unique user "id" with the non-unique "OPT-OUT".

Compliance with this requirement is simple to verify and monitor. Watchdogs can confirm through external sampling that opt-out cookies are being delivered on request and that they are non-unique. More often than not, opt-out failures will result from simple technical problems that cause the cookie not to be written in the first place; we've seen more than a dozen broken opt-outs in the last year through PrivacyChoice. Through random sampling, you can also confirm that once a company delivers an opt-out cookie, it stays in place and doesn't revert to a unique cookie that could be used for tracking.

If an ad company still needs to use unique cookies for non-targeting purposes (like frequency caps) even on opted-out machines, they should certify and publish which cookies are which, so that behavioral cookies are always identified. Verifying accurate identification should be part of the NAI's annual technical review, as should be confirming that no non-cookie tracking methods are used (which is also easy to monitor externally). These kinds of direct reviews or audits of backend process will be far more simple and effective than trying to determine forensically whether opted-out consumers are only being served untargeted ads.

A cookie overwriting requirement would require some technical changes for companies that currently separate their behavioral cookies from their opt-out cookie, or those that combine behavioral and non-behavioral functions in a single cookie. But this seems like a small price to pay for a significantly more effective and accountable opt-out framework, which is the heart of the self-regulatory effort.

Do people care about tracking oversight?

Summary: An analysis of actual do-not-track elections made through the PrivacyChoice service indicate that a significant percentage of web users are less likely to block tracking by companies subject to industry oversight processes.

It's tempting to think that there are just two kinds of consumer views on behavioral tracking: some people simply don't care about whether they are tracked online, and others care so much they would always choose to to avoid it completely.

Results from real users on PrivacyChoice suggest that a good number of people actually have more refined views, including as to the value of oversight.

The PrivacyChoice do-not-track service, TrackerBlock, gives users a unique set of choices when it comes to tracking:

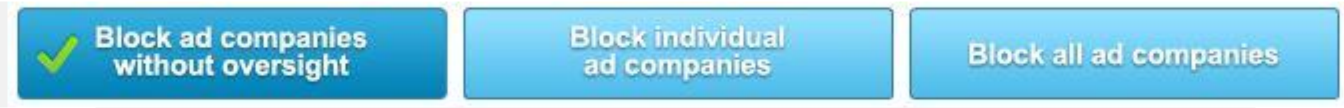
✓ Block ad companies without oversight
Block individual ad companies
Block all ad companies

Click to see TrackerBlock in action

The page explains that "oversight" means that the individual tracking company is a member of the Network Advertising Initiative, and thereby subject to policy requirements and annual compliance reviews. Our explanation goes out of its way to be objective, and arguably "undersells" the idea of oversight as a choice factor.

When given these choices, what do people choose? Here's are the selections made by the last 1,000 unique new TrackerBlock users:

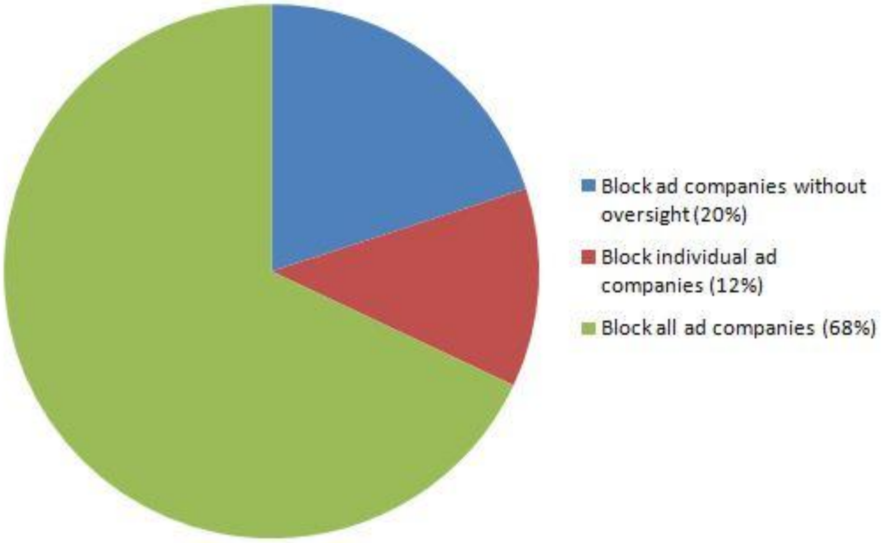

It's no surprise that a substantial majority would choose the "nuclear option"; people who find TrackerBlock tend to be privacy-oriented, and many are attracted to our unique aggregation of all tracking companies (not just NAI companies). But I was surprised that even among a privacy-concerned people, nearly one-third choose a more limited opt-out when it is offered.

A more scientific study would vary the placement of choices and take other steps to ensure representative results. But these results still point in an interesting direction:

- **Many consumers care about the notion of oversight when it comes to tracking.**
- **When offered more granular control over tracking, a significant number of users will use it.**
- **Fortifying the oversight process and explaining it at tracking choice-points is useful to consumers and may temper opt-out rates among companies subject to oversight.**