Browser Privacy Mechanisms Roundtable

Welcome by Robert Barr, BCLT Executive Director

Tutorial on the state of online tracking
Tutorial leader: Ashkan Soltani

Wednesday, February 9, 2011
Bancroft Hotel, Berkeley

ROBERT BARR: We're going to get started; everyone please take their seats.

Good afternoon. I'm Robert Barr. I'm the Executive Director of the Berkeley Center for Law and Technology. On behalf of BCLT, welcome to our conference and to a beautiful day here in Berkeley.

A few announcements, important announcements. The wireless is free. The bathrooms are also free, but they're downstairs. When you go out near the front door, go down the stairs to the bathrooms. We have a Twitter hashtag for the event; it is #BerkeleyDNT. There are three Es in Berkeley, and DNT as in Do Not Track, #BerkeleyDNT.

There will be a break in the afternoon, 3:10 to 3:35, and if you need anything during the day look for me or Louise Lee or David Grady, who I can't see right now.

Our first speaker is Ashkan Soltani. Ashkan is formerly a technologist at the FTC's Privacy Division, and he's currently a researcher and consultant focused on privacy, security and behavioral economics. And he's going to give a tutorial on tracking technologies.

ASHKAN SOLTANI: Thanks everyone for coming. So, there's no need to do introduction. I'm a former graduate, so I come from Berkeley and have great respect for the place.

Quickly to the agenda. I just want to go over Web tracking. A lot of this will be repeat for most of you, so for those that aren't, hopefully it's useful. I'm going to just cover Web tracking, cyber ecosystem a little bit, describe our consumer choice mechanisms a bit and then hopefully lead into Do Not Track.

So, not to dive into the technical details too much, but I want to kind of describe the technical mechanism where most of the Web interactions happen. How many of you guys remember this? A show of hands? (Show of hands.) So, it's about 20 years on now, and it's still very present in our everyday lives. I clicked highlighting because it enumerates the fundamental issues we have with privacy on the Web. Effectively, that's Web pages offer a mosaic of content from all over the Web pulled together for a single page for us to view.

Many of you have read this Huffington Post article on Do Not Track, perhaps. Like most pages on the Web, it contains content served from the first party, Huffington Post, as well as third parties, all the advertising and social widgets. As a result, the data generated from my interacting with this page may go to Huffington Post; it may also go to third-party advertisers or social widgets shown here.

There are varying definitions of what a third party is in this context. Some consider a third party one that does not share common ownership. Others have said it refers to content that is under common branding. The technical definition is often based on what domain appears at the top of the page, like Huffington Post.

Additionally, it's unclear whether when you interact with one of these social widgets, when it's Facebook or Twitter or share widgets, whether this constitutes a first party or third-party interaction because you're kind of interacting directly with them in this view of a first-party interaction.

There are some third-party interactions that are not visible to the user. These are facilitated by what are commonly referred to as Web bugs, or Web beacons. They are effectively elements in the page that enable other entities to track the user's activity on that page. Using a browser program called Ghostery, which looks for these beacons, we see approximately 10 or so on this page that a user may not have been aware of.

I'm not going to bore you with the technical details, but this will clarify every single element on that page that we just saw. It kind of creates an interaction much like this one. So, this is kind of an HTTP GET request. And the core element inside that there's an observer, a host, an identifier, which is this time a random string, which is your cookie, and an event, which is you've viewed this Do Not Track article on The Huffington Post. So, while this is interesting, it is that these elements, multiple observations of observer identifying event, can lead to deep insights about you, i.e., your profile.

This can consist of many touch points. For example, my activity on multiple sites going to the same observer, or conversely my activity on one site going to multiple observers. To clarify, it's the existence of this profile that seems to bother people, not that it's being used for advertising; i.e., there exists a profile that I don't have any knowledge of or control over that's being used to make decisions about me. Now, this has always existed; the difference is the scale with which this occurs.

So, a bit about the ecosystem. I was one of the researchers for the Wall Street Journal's "What They Know" series, and that series identified that if we look at the top 50 Internet properties, and we identified 131 different tracking companies and found on average 54 tracking objects, or beacons installed per site. This survey of kids' sites actually revealed more tracking objects, on average 84.

In our 2009 Privacy Report, we identified a single observer, Google, on 92 of the top 100 websites, and 88 percent of the top 350,000 tied to our sample -- sorry, 88 percent of the

350,000, not the top 350,000. Now, discretion can be of different sorts. I noticed data versus AdCrunch or DoubleClick data, and it's sometimes tied to different identifiers. But the key point is that it's still disclosed to the same entity, which is Google.

Another way to view this, and some of you may have seen this before, is we have the trackers across the top, and we have the top 100 properties on the left. And so, zooming out, each yellow dot identifies a tracker. We see that, for example, on the one domain, Blogspot.com, we encountered 100 different beacons in March of '09. Conversely, on the top 100 domains, we saw Google beacons of some sort on 92 of the top 100. That's 92 different sites in which we've encountered activity from those interactions flowing to one entity.

We recently — October of 2010, Sheryl Sandberg was quoted that "10,000 new websites integrate with Facebook every day," so this is like social widgets and share objects and interpersonal relations. And just in terms of activity, we see for example 650,000 comments per minute. And this isn't even clickstream; this is just people actually volunteering content to one third party.

Now, Facebook and Google aren't the only ones with an ecosystem, obviously. They're just the biggest. And so we see a complete ecosystem with players handling, sometimes they're trimming narrow roles. Some of you may have seen this graph of the ad ecosystem. We see some providers, like data providers, some are identity providers whose sole purpose is to identify you and identify who you are. Finally, some do analytics, tracking, ad serving. And they all work in conjunction, and they kind of skirt some of the privacy concerns because they each have such a narrow piece of the puzzle.

For example, how many people by a show of hands have signed up for a free contract? (Show of hands.) A couple. And how many won?

So, what you've won was the ability to be identified. So often these providers' sole purpose is to help other companies figure out who you are by either your name, your e-mail address. They effectively ping these services, and say, who was this user, say your e-mail address, and then use that to actually provide the query of other data, like offline and online data sources, and then this data can be used to serve you ads, credit reports, whatever, credit card offerings.

So, how do we address this? We often touted CHOICE. And we think the market is OK, and we let CHOICE handle it. We let consumers kind of declare what they want to do and opt out of the tracking they don't want. For example, the NAI is often touted as the primary CHOICE mechanism for consumers. This is how you opt out of at least the advertising space. So, to arrive here, you have to know that this mechanism exists. You also have to realize that this NAI only applies to opting out of being advertised to. So, the data is still collected; it's just not being used for ads. It's not being used to show you ads, so you're not even aware of it.

This page only had 67 out of 130 trackers that I'm aware of, and analyst companies that I'm aware of. Possibly they're going to come online, but, for the most part, opting out on this page by clicking all of those checkboxes still doesn't opt you out of the ones that are not members of NAI.

Lastly, it's actually kind of converse what we tell consumers to do, which is delete their cookies, or go through and pick out the cookies they don't want used. And since the software mechanism is currently enabled by cookies, so this kind of crosses with the best practice of deleting our cookies, or sometimes the cookies themselves are eliminating the opt out.

Another quick mechanism is this dialog box that all the major browsers have, which is the ability to block third-party cookies. So, I don't know if any of you guys use this mechanism. What's interesting about this is that they don't work as expected. With the exception of Firefox blocking third-party cookies, this only blocks the setting of a cookie in a third-party context. What that means is if you visited a site in first-party context, Facebook, even if you have third-party cookies blocking enabled, that site can still track you across the Web and read your cookie. So, again, it's kind of contrary to what we think as a good CHOICE mechanism.

We see much more of this type of thing in the lab. We see technologies that enable companies to circumvent consumer choice. So, cache cookies many of you guys are aware have been used to circumvent the deletion of HTTP cookies. We see some great work on the use of P3P to circumvent browsers' default settings of blocking cookies. CSS history attack, which is the ability to query what sites you've visited, and browser finger-pointing, as well. All of these operate in an area outside of consumer choice. It's very hard to block these attacks or these mechanisms. Part of it is actually an architecture of the way the Web works.

There's a recent attack, which is the ability to detect whether you're logged into services like Gmail. Part of it is at least from the fact that it's recently highlighted. That will likely always exist, like the ability to tell whether you're allowed into a service or not. The question is how do we, knowing that these technical limitations are in place, how do we kind of still enumerate consumer choice.

And just to highlight this process a bit more. A recent company that I don't think has been covered in the mainstream media yet, this company helps address the issue of consumers deleting their cookies, or using different devices, and browsing the Web. Their mechanism completely eliminates the cookie defense; you can still track a user that has deleted their cookies or switched browsers. So, they are in the business of helping you get around those pesky cookies. I'll kind of X-out the name because I don't want to name any names.

That leads us to kind of the purpose of this panel, which is how do we effectively create a consumer choice mechanism that we can then build from? And that's often currently referred to as Do Not Track, right. At a store, my interpretation of what this mechanism

is, is it's effectively a way for consumers to declare their choices. If anything the goal should be to help consumers voice their preferences, and then we can build around that. And we've seen a bunch of different definitions of what this could be, which is Do Not Track, do not collect, which is hard given the way networking equipment works and how everything logs on the Internet — it would be a very hard thing to completely not collect any data.

The other is to, for example, do not use, which is in some ways what we have now, which is do not use for behavioral advertising. Another one is do not identify, which is, for example, what happens when Google, for example, deletes the last octet of your browser's IP, as well as drops your cookies. So, that effectively makes it hard for them to identify who you are. These are all valid kind of approaches that we could talk about.

Lastly, in my opinion the interesting conversations, interesting questions come once we've identified that we have this thing put away. We have to come up with ways to have tight definitions of what tracking is.

If we do that, if we have track mechanisms, are we driving more first-party interactions, or other techniques, like the DNS alias thing, which allows domains to masquerade as other domains? Does consent actually create worse practices, or drive the industry to be able to do more because consumers are now actually making a choice and have opted to allow companies to track them, predicated on access to content? And then the full discussion should also address the model space in embedded devices. As you see more and more kind of appliances kind of on the Internet, with an IP stack, we should be kind of sure to address these, as well, like your smart phones, and your TVs and your cars. Lastly, how would compliance work in this? How would you detect companies that are not kind of adhering to these guidelines? And how do we automatically detect them, and how do we help the FTC enforce bad actors? (Applause.)

ROBERT BARR: Thank you, Ashkan.

The FTC's Privacy Agenda
Commissioner Julie Brill

ROBERT BARR: Our next speaker, and then we're going to have discussion panels the rest of the day, but our next speaker we're honored to have Julie Brill, commissioner with the Federal Trade Commission. She was sworn in with the FTC on April 6th, 2010. Her term expires in 2016. Before she became a commissioner, Julie was the senior deputy attorney general and chief of consumer protection and antitrust for the North Carolina Department of Justice. She also has been a lecturer and chief of consumer protection and antitrust for the North Carolina Department of Justice.

She also has been a lecturer in law at Columbia University School of Law. Prior to her move to North Carolina Department of Justice, she was an assistant attorney general for consumer protection and antitrust for the State of Vermont for over 20 years. Prior to

that, she graduated magna cum laude from Princeton University and from NYU School of Law, where she had a Root-Tilden Scholarship for her commitment to public service.

Please welcome Julie Brill. (Applause.)

COMMISSIONER JULIE BRILL: Good afternoon. Can everybody hear me OK back there? Great. My voice usually carries pretty well, but I just wanted to make sure.

So, first of all, thank you so much to the Berkeley Center for putting on this afternoon's event. I think it's incredibly timely, and I'm really looking forward to the conversation. And thanks especially to Deirdre Mulligan, Chris Hoofnagle, Robert Barr and everybody else who worked so hard to put this together.

I thought I'd start out this afternoon by giving an overview of the FTC's recently released preliminary privacy report. To a great extent, of course, this is the document that has generated so much of the conversation that industry and regulators have been engaged in for the past couple of months. So, it might make some sense to talk about it a little bit, to place in context our upcoming conversation this afternoon about Do Not Track.

So, first in terms of a brief, very brief history of where the FTC has been on privacy, I think it will help to talk a little bit about that in order to understand why we wrote the report, why we made the preliminary recommendations that you see in the report, and why we are recommending a Do Not Track mechanism of some kind.

For decades, the Federal Trade Commission has led the federal government's efforts on privacy and privacy enforcement. First, over 40 years ago, our efforts began with respect to the Fair Credit Reporting Act. Many of us are old enough to have been involved in the early days of enforcement work under the Fair Credit Reporting Act. A little bit later on, we were involved with the Children's Online Privacy Protection Act, the Do Not Call Registry, which Dave Barry calls the most popular federal program since the Elvis stamp. And we're very proud to run it. And also, of course, there's the CAN-SPAM Act.

Enforcement models and the privacy frameworks generally that the Federal Trade Commission, as well as the states have used, has really changed over the past 20, 25 years. In the 1990s, we were using a -- what I call a notice-and-choice model. Even though Congress did not take up the Federal Trade Commission's recommendation to adopt the Fair Information Practices Principles into legislation, there were certain aspects of FIPP, the Fair Information Practices Principles, that were adopted in various pieces of legislation over time. And to a great extent, the enforcement work that the Federal Trade Commission engaged in, as well as the states, and our perspective on privacy was really focused on notice and choice to consumers.

Over time, I believe that the notice-and-choice model really became unworkable for consumers. The privacy policies became incomprehensible to consumers, if they were going to read them, and of course most consumers never took the time to read them, in large part because they were incomprehensible. These things kind of fed upon each

other. Simply put, the notice-and-choice model really creates too great a burden on consumers. And then in the auth, as we said back in the 2000s, the harm model came into prominence.

So, the notice-and-choice model was still there, we were still using it, but we started to focus a lot more on issues like identity theft, and security breaches. And there are a number of reasons for that. The states became very active in adopting laws relating to security breach notifications, and of course there was a lot of work at the federal level in terms of identity theft protection.

So, this model, the harm-based model as I call it, it's certainly necessary in terms of privacy protection. But, from my perspective, it is not sufficient. It's essentially reactive. It only focuses on damage in the broad sense, not monetary sense, but it focuses on problems after they've taken place. So, it's reactive in that sense. And it doesn't place value on nonfinancial damages, or nonfinancial harm, such as embarrassment through disclosure of sensitive information.

Now, at the same time that these two models were sort of working either together or in opposition to each other, that is the notice-and-choice model and the harm model, the landscape, the underlying landscape on the Internet was changing dramatically, both with respect to the amount of data that was being collected and used to which data was being put.

So, there was a growing recognition at the Federal Trade Commission that we needed to grow our privacy framework in order to fit the realities of today. The models that we were using, simply put, needed to be updated and modernized. So, we held a series of roundtables in 2009 and 2010 to gather information from a broad array of stakeholders. I would be -- I was pretty sure that there are quite a few of you here in the audience who participated in those roundtables. How about a show of hands — how many of you did participate in the roundtables that we held? Yes. So, a fair number of you did participate in that. We actually held one here in Berkeley, and there were a couple of others.

There were some themes that emanated from those roundtables. The themes included the following, collection and use of consumer data is ubiquitous. It's happening everywhere, especially online, but also offline. Consumers don't understand and don't have the ability to make informed choices about data collection and use in the current environment.

Another theme that emanated from the roundtable was that privacy is important to consumers. And a fourth theme was that the increasing flow of information, including targeting advertising, provided many of the important benefits that consumers receive online. For instance, much of the free content that is available to consumers is paid for through advertising that arises as a result of the targeting activities that are currently taking place.

On the technological front of interest to many of you here, another theme that emanated from the roundtables is that the depletion between personally identifiable information and

non-personally identifiable information is blurring. There have been huge technological advances that are allowing the tech savvy of you out there, and others who are quite tech savvy, to combine disparate bits of information that are in and of themselves potentially anonymous, or non-identifying, but once these bits of information are all combined together, it's quite possible to re-identify, or associate that information with particular individuals or particular computers or browsers.

More recently since the roundtables were held, we started to hear about even further technological advances that are blurring this distinction between TII and non-TII browser, fingerprinting technology, which I'm sure a lot of you here know much more about than I do, allows websites to gather and combine information about a consumer's Web browser configuration, including the type of operating system used and installed plug-ins and things like that, that will enable the website to uniquely identify — it almost becomes like a fingerprint — to uniquely identify and track that particular consumer.

Of course, in the mobile context there are the unique device identifiers that do the same thing, especially when combined with location information and other information, becomes a very, very rich source of data about a particular consumer. So, we issued our report in December, just a couple of months ago, December of 2010.

The purpose of the report was and is to develop a model of best practices, and it drills down fairly deeply into some of the potential best practices that we are thinking about should be. It is intended, the report, to be used by policy-makers, including Congress, and it is also intended to be used by industry as industry develops potential self-regulatory approaches.

The scope of the report is designed to be quite broad. It applies to both online and offline commercial entities that collect, maintain, share or otherwise use consumer data that can be reasonably linked to a specific consumer's computer or device.

The report had three principle components. Everybody talks about Do Not Track, that's clearly where all the buzz is. But I want to take you back a step because I think it's important, and it's of concern to me that it not be lost that there really are some bigger issues that are mentioned in the report, and that we want industry and others to comment on.

The first component is privacy by design. Just what Ashkan showed you, although I think it's been updated in the last 20 years. I certainly hope it has been. And what we believe is included in the concept of privacy by design are some subconcepts, like privacy and security, should be built into new products not retrofitted after problems arise. The level of security provided should be equal to the sensitivity of the data that is the subject that's being collected and that's being maintained and used. Companies should limit the amount of information collected to what is actually needed. And companies should retain data only so long as they are going to need it, and then it should be discarded and destroyed. So, that's the first component of the report Privacy By Design.

The second component relates to choices, a bunch of issues around choices. And it can be boiled down to the following: Choices for consumers about data practices need to be simplified. The choice mechanisms need to be easy to understand so consumer choice is both informed and meaningful. The scope of the choice — and the report talks about this quite a bit — should cover not just use of the data, that is not just tracking or behavioral advertising, but it should also cover collection of the data; that is the ultimate tracking of the consumer. That's our preliminary proposal. I know there's a lot of discussion about what is Do Not Track, but our proposal is that it should cover both use and collection.

One way to simplify choice is to exempt commonly accepted practices to help remove the clutter in privacy notices and privacy disclosures. So, for instance, the concept that when you go to a particular website that they might hire a contractor to fulfill your purchase, that's not something -- and the fact that your name and address is going to be given to this fulfillment company in order to send you the product that you've ordered, that's not something that really needs to be cluttered up in the first or second layer of notices that are given to consumers.

In the ultimate privacy policy, the full-blown version, yes, perhaps that needs to be disclosed, but let's get rid of the clutter to focus in on the real issues that consumers need and want to know about.

And another way to simplify choice is what I was just alluding to, to use this kind of layered concept where the most important information is given to consumers right up front, then they can find a place — there's a link — where they can get more information. And then there's a further place they can go to get the full-blown privacy notice that discloses all of the practices that the company might be engaged in.

Now, opt-in versus opt-out, there's always been a lot of debate about that in the past — should choices be opt-in or opt-out? The report, and the commission, has not taken a position on whether there ought to be opt-in versus opt-out because we've kind of moved beyond that issue of opt-in versus opt-out. It's really not an issue of which direction the inertia goes, but it's an issue of how robust the notice is. That's what's important at this point in time, at least in terms of the preliminary report — how robust is the information being given to consumers, and how easy is it for them to make the choices that they're being offered?

The indicia of the choice mechanisms that consumers are going to be offered, whether it's through a self-regulatory mechanism, or whether it's through congressional action, there are five indicia that the commission has laid out in its preliminary report in terms of how we're going to judge whatever mechanism we're looking at.

The first is how easy is it to use the choice mechanism? The second is how effective is it, and is it enforceable? So, issues relating to enforcement as a choice mechanism. The third issue, is it universal? That is, is there participation by the vast majority of industry in that choice mechanism? The fourth indicia of how effective the choice mechanism is

is, does it provide for opt-out or opt-in of collection of information, not just use? The last criteria that we've outlined on a preliminary basis is, are the choices persistent? That is, will they last, or instead, as Ashkan was alluding to, are they indicated through cookies that may be delighted or may expire on a relatively short-term basis?

So, those are the first two parts of the report.

The third part is that data practices need to be transparent. Consumers need to understand what kind of data companies are actually collecting. Consumers should have the right to access, and in addition to understanding what kind of information is being collected, we posit and propose that consumers should have access to that information. Now, the right of access that we talk about should be proportional to the sensitivity and intended use of the data.

So, for instance, we can compare to paradigms where data is being used simply for marketing purposes on the one hand versus data that's sold by data brokers to financial institutions who are going to make decisions with respect to a consumer's application for a loan, or for insurance, or perhaps when that information is used for employment purposes. You can see that the sliding scale of the need for access and the need for consumers to see what that information is, and to potentially correct it, is going to change depending upon the kind of paradigm that we're in and what the information is and how it will be used.

And it's probably worth pointing out very briefly that simplified choice, the second component of the report, really works hand in glove with the third component that is transparency. When you make choice easy, and you make it easy to access, and you make it easy for consumers to understand, then that is going to enhance the transparency of a company's practices, so the two really go together.

Let's talk a little bit about Do Not Track, and I know we're going to spend a lot of time this afternoon talking about that. As I said, it certainly has generated the most buzz. So, what is Do Not Track? I pretty much agree with Ashkan's definition of Do Not Track. I think of it as a mechanism for allowing consumers to exercise choice about the collection and use of their online behavior. Fairly simple, although there's obviously a lot that one can dig into in that fairly simple definition.

Now, one mechanism for providing Do Not Track choice for consumers is through a browser-based mechanism, and it can be through a persistent setting, which is known as a header. We can send a signal to consumer choices about being tracked and receiving targeted ads to a particular website or to an ad network. Preliminarily, on a preliminary basis, the commission recommended a browser-based solution in the report. And the reason is because it seems to most easily satisfy the five criteria that I mentioned, especially three of them. It allows the consumer the choice to opt-out of tracking, as well as opting out of receiving ads based on their behavior. It is easy to use. And it promotes universal adoption.

So, that was the reason, among others, why on a preliminary basis, we were saying, “Gee, let's think about going towards a browser-based approach.” There are other proposals, we'll be talking about them this afternoon, and Ashkan outlined them, I thought, very nicely.

There are some enforcement issues with respect to any of these choices. That is both the cookie-based approach that Ashkan talked about, and the header-based approach that might be more in line with the browser-based solution necessarily relies on an understanding that the advertiser or the ad network, that is the entity that is receiving this information about the consumers' wishes with respect to their information, that that entity receiving the information honors the consumer's choice. So, that obviously raises some enforcement issues that I know we'll be talking about as we go forward.

You know, I've been asked a lot what I think about the industry's response to our report, and I know we'll be hearing a lot about that as well this afternoon. But I think it's probably important for me to note right up front that the commission has been calling for industry to develop a self-regulatory response to problems surrounding behavioral advertising — that is one subset of the issues that we talk about in the report — for two years. We first issued our call for an industry response in February of 2009. So, it's my position that industry has been kind of slow to deal with this issue. On the other hand, as I mentioned when I started, we've been very pleased that since we issued our report two months ago, we seem to have gotten industry's attention now, and it does seem as if there is much more of a response occurring, and that there is much more interest in trying to develop some answers to the problems that we point out.

I think one of the things that Deirdre asked me to mention right up front is, in addition to Do Not Track, which we'll be spending a lot of time talking about, what are some of the other concerns that I have with respect to what's happening with consumers' information. And I think that as we talk about Do Not Track some of these other issues shouldn't get lost, because I think they will inform some of what we think about Do Not Track, and some of the parameters, and the perimeter area where Do Not Track ought to move.

For instance, I think it's incredibly important that we think about information that data brokers are providing to financial firms, as I mentioned before. And the Wall Street Journal series that Ashkan was apparently involved in, which I didn't know until he mentioned that, I thought had some terrific pieces about the use of information by insurance companies, what they're doing, for those of you who haven't seen the piece, and assuming that it's true, they talk about insurers purchasing information that has been scraped and sniffed from social networks and using that information to determine whether and at what price to offer insurance to applicants.

This raises some pretty serious issues, issues that one would hope would be addressed through the Fair Credit Reporting Act, but probably aren't because of some limitations given that that law was written in the 1970s; there are some limitations to the definitions of when we're talking about consumer reporting information, what we call in the common parlance credit information. Is it credit information? Is it a credit report? Is the data

broker a credit reporting agency? Because there are limitations in that regulatory framework, the activities that we're now learning about may not easily fit into those series of regulations. And that to me is a huge, huge problem.

OK. So, where do we go from here? The report, the comment period has been extended, I think, as some of you know. It has been extended, I think, until February 18th, which is rapidly approaching. We've received over 200 comments so far, and we expect that on the 18th our server may shut down, hopefully not. We are still committed to filing a final, or creating and issuing, a final report by the end of the year, that is, by the end of 2011. I think there are a couple of big questions that I'd like to pose for us to be thinking about through the afternoon.

First of all, there's the big question about whether self-regulatory responses will be sufficient. And as I said, we're going to judge them by the five criteria I outlined. It's going to be very interesting to see how those proposals develop over the next few months to see if they satisfy those five criteria.

I think the second big question, the one that I'd like to talk about this afternoon if we have an opportunity, is whether the advertising industry and others who receive information about what consumers want to happen with their information online, we'll be willing to honor those choices that consumers are making. That is, will there be a commitment by the advertising industry to honor the choices consumers make through Do Not Track and elsewhere?

In agreement by the advertisers and ad networks that the collection of behavioral information, not just the delivery of targeted ads, but the actual collection of the information in the first place, if they will agree that that's important to honor, when a consumer expresses a choice, with respect to that issue, I think it would move the self-regulation effort a long, long way towards being acceptable. And, of course, I have to close by saying that it's still our position that if the industry does not act quickly and sufficiently, we will ask Congress to take up this issue.

So, thanks very much, and I look forward to the dialogue this afternoon. (Applause.)

02092011 Browser Privacy Roundtable Discussion Topic 2

Discussion topic 2: Technical Approaches

Jim Brock, PrivacyChoice
Brian Carver, UC Berkeley, School of Information
Peter Eckersley, Electronic Frontier Foundation
Edward Felten, Federal Trade Commission
Dean Hachamovitch, Microsoft Corp.
Jonathan Mayer, Stanford University
John Mitchell, Stanford University
Jason Schultz, UC Berkeley, Samuelson Clinic

Christopher Soghoian, Indiana University
Sid Stamm, Mozilla Corp.
Harlan Yu, Princeton University

Moderator: Chris Hoofnagle

Wednesday, February 9, 2011

CHRIS HOOFNAGLE: As Robert Barr mentioned earlier, we are going to do our break after this discussion, and unlike the last discussion, the last session, this is going to be a discussion, so there's going to be no allotted time for the speakers. Instead, we're going to kind of do a Q&A, and we'll do a Q&A with you as audience participants as well.

What I would like to do — and in the back if you could quiet down because it's really hard to hear, yes you — what I would like to do is throw out a question to our two browser makers who are on this discussion, Google is going to join us for the third, but for now we have representatives from Microsoft and from Mozilla, Dean and Sid. And I would like to throw out this first question to you.

Commissioner Brill set forth five indicia for effective mechanisms for consumer control. They were ease of use, can the consumers actually use the mechanism. How effective and enforceable is the mechanism? Is the mechanism universal, meaning is more or less the entire industry using the mechanism, or do you have a bunch of outliers who won't participate? Does it provide for opt-in or opt-out of collection of data, not just use. And so the important back story there is that there is a lot of industry groups that are saying that collection doesn't matter, only use matters. And if that's the case, we can just get rid of confidentiality all together. So, are we talking about collection or use? And then, finally, are the choices persistent?

Do your mechanisms satisfy these five conditions? And, Dean, do you want to start, and then we'll switch to Sid, and then we'll go to our larger group of technical experts.

DEAN HACHAMOVITCH: OK. That's good, that removes a certain quantity of snark from the room.

So, the ease of use, enforceability, universality, persistence and whether enables opt-in or opt-out of information collection, right? I want to make sure I have those down.

I think a good starting point is to say that the two solutions that the two browser vendors here have are less than competitive and more either complementary or superset/subset depending on how you look at it.

From the Microsoft Windows Internet Explorer one, I think ease of use is, yes. You go to a Web page and you click on a link. So, I think we have proved people know how to find Web pages and click on links. So, I feel pretty good about that.

Enforceability. I think the enforceability actually has two aspects, and I think both of them get a yes vote. There's enforceability in terms of indicating to a site a privacy preference. The site can determine, hey, did the user attempt to signal a desire for privacy? Yes, the tracking protection list is on. And there's also enforcement, in a way almost inspired by a security feature, and to be clear I'm not calling this a security feature. You can think about this in the following way. We don't have "do not send me popup window" HVCT headers. We just have popup blockers in browsers. It's not like we decided to signal a website, I really don't want a lot of popup windows, and now I'm going to sit around and wait for legislation to come to bear. So, enforcement I feel good about on two fronts.

Universality, that's actually a great question. We're actually going to bring the design for a tracking protection list, as well as a persistent setting to indicate tracking preferences to the W3C as a proposal for Web standardization. We're doing that because we do want it to be universal, and we think there should be a consistent way that websites, and Web developers can determine the user's preference.

Finally, persistence. Yes, absolutely. And opt-in and opt-out of collection, that is completely up to the people who design the list that people choose.

SID STAMM: All right, great. Thank you.

By the way, we really like the idea of you guys standardizing these CTLs. We're all about these Internet standards any way they can go.

So, let me just run down the same list. It's pretty easy to turn on a Do Not Track header, you check a box, and the header is on and it gets sent. So that addresses the ease part.

Enforceability is going to be a little bit tougher than the rest of them for us because it relies on, as the previous discussions have been, the parties receiving the header doing the right thing. And one of the reasons that we think it's the right idea to express these as preference is because the right thing is going to be different depending on what the site does.

So, it's not exactly clear whether specific behaviors like using third-party cookies are OK in all contexts or not. And so we don't want to turn them off because they might be useful.

So, it is universal. Our Do Not Track header, for those of you who aren't familiar with it, the user turns it on, and it gets sent to every site saying this user doesn't want to be tracked. So, it goes everywhere. I would consider that universal.

Additionally, it is an opt-in feature. If people feel the need to turn this on, they can go seek out the preference in the configuration panel, and opt-in to this expression that they don't want to be tracked.

And, finally, it is persistent. Unlike the stop ads, cookie technology that can be used to opt-out of various advertising networks; this doesn't get erased when you erase cookies. This doesn't expire. When you turn it on, it stays on until you turn it off.

CHRIS HOOFNAGLE: Before we go to the group of technologists, let me ask, what are the primary differences between the two approaches that your companies are putting forth?

DEAN HACHAMOVITCH: Well, in some ways, I'll go back and say, I wish I had noticed this before, that they really are superset/subset relations. I think that the notion of a persistent Do Not Track setting that a browser communicates to a site is present in both models. In the case of one, it's an HTTP header, and more technology. In the case of IE, there is another underlying technology.

I think the primary difference that I see is that the Internet Explorer model does more than allow the user to express a preference. It enables the user to have a mechanism to do something about that preference. Specifically, tracking protection will enable users to block Web content that is associated with tracking, and that's definitely a protection. Instead of just asking a site, "please don't track me," we can just block that content, and pre-empt that tracking.

SID STAMM: All right. So, there is a definite overlap between what we're trying to accomplish here. And for those of you who use Firefox, there are various add-ons that can turn off pieces of the Web if you'd like to. So, you're welcome to use that.

But I think it's really important to give a site an opportunity to explain to users exactly what they're doing without the users making an assumption about whether or not it's acceptable to them without having the full picture. So, one of the ways I see the tracking protection list and the header differing is in that if a person wants to opt out of tracking, they can click on this header, the Do Not Track header in Firefox, and if the site notices this, they can maybe bring the user to an opt-out page that explains what they're opting out of. Hey, we're tracking you for this purpose. You may want to keep it on, is that OK? And they have an opportunity to then create an exception for themselves without having the burden on the user, or whoever is maintaining a list of some sort, to determine whether or not it's OK. Then, the onus is on them to explain to the users what they're doing in a clear way, provide clear choice, so that the right decision can be made on a case-by-case basis where absolutely necessary.

CHRIS HOOFNAGLE: Thank you. And so the rest of our group here, we have a mix of technologists and lawyers, and lawyer/technologists. I would invite you to comment upon the two different approaches. Are they broad enough? Do they do what Commissioner Brill asked for? Are there other approaches that do it better? And whoever can get that red light on first goes first.

CHRISTOPHER SOGHOIAN: So, we live today in a cycle of arms races, and Dean described how there isn't a "do not show me popup signal" that gets sent out. We have a popup blocking feature in all of the major browsers, and for a couple of years that worked really well. And then the ad networks all started innovating around the popup blocking features, right. And so we get pop-unders; we get ads that take up the entire screen, where you cannot close the window for 5 seconds or 10 seconds.

When we have specific technical solutions that address either annoyances or tracking in the market, companies see that as an invitation to innovate around those features, right. And they'll say when they're interacting with their potential clients, "We have a solution that no one else has," and not only that, when they get caught doing what they're doing, when they move from cookies to Flash cookies, they'll say, "Oh, we thought consumers didn't want to be tracked with cookies; we didn't realize they didn't want to be tracked with Flash cookies," or "We thought the consumers didn't want to be tracked with Flash cookies, but we didn't know they had a preference regarding browser fingerprinting, that's a completely different technology."

And because we're in this arms race, because we don't have a way for consumers to say no, we're stuck in this cycle of companies innovating around the technologies to consumers. And the header, which I'm a big fan of and has been part of this discussion, is a way to break out of the arms race. And before I pass it onto everyone else, I want to mention one other thing, which is a key difference between the header and tracking list approaches, is that the header approach pushes the cost of complying and responding to the header completely to the ad networks.

The ad networks have to figure out how they're going to respond to the header. They have to do all the work of implementing it, whereas the blocking list approach, that pushes the cost on to the user, or their agents, the browser vendors. The user, or the browser vendor, has to go and find a list. They have to keep the list updated. When the ad networks innovate around the list they have to try and find the new domain to add to the block list. And, to be honest, if there is an organization or an industry that should have to bear the cost of this, it's the ad networks. Let them pay to respond and respect the user's header. The users shouldn't have to be playing a game of cat and mouse.

JASON SCHULTZ: I feel like we're on a quiz show and it's whoever buzzes in first, right?

CHRIS HOOFNAGLE: We might, in fact, be doing that. In fact, I have a priority button that gives me a focus. So, how about we do Peter, Harlan, and then we'll move to you, Jason?

PETER ECKERSLEY: I think it's important to remember that we are embarking on an incredibly difficult, maybe impossible, but maybe almost impossible task, of retrofitting privacy into this incredibly complicated hypertext system that we've been building for the last 15 or 20 years. Retrofitting it when we didn't do privacy by design, or didn't do it as well as we should have along the way.

This task is so difficult that we are going to need both of these technological approaches. I'm going to abstract away a little bit from the specific implementations in the browsers, and talk about black lists and signaling mechanisms. We're going to need both black lists, but that can stop a particular URL from being loaded in the browser, and signaling mechanisms that have some legal enforcement where a consumer can say "I don't want to be tracked without my permission," and then if companies just ignore that, particularly large third-party companies, then maybe there's some legal mechanism for that.

The reason we're going to need both of those things, as Chris pointed out, we need to break the arms race, if we're just blocking things then some company will find a way to get around the block. We know that. But, we're also definitely going to need to have the blacklist part, because if we're talking about an FTC, or some other agency-enforceable mechanism, there are going to be trackers that just play run and hide. They won't be upstanding corporations that you can find an address for. They'll be out there. They'll be malicious websites doing tracking.

And so, we need to protect ourselves against both of these pieces, and we need both of these technologies working in parallel. I just see one of them as being potentially led by technology, that blacklists are a thing we can do purely with code. The header mechanisms or signaling mechanisms need a policy component as well.

HARLAN YU: I'd like to agree with what Chris and Peter have said already. I see the two different approaches as complementary. The thing that worries me about the Microsoft approach and the reasons why I think it's not sufficient alone is that I'm fairly confident that if that's the only protection that users have, it would be really easy for ad networks and other third parties to innovate around it.

As we've seen with malicious software and botnets, it's really easy for them to get hundreds and hundreds of domains and to rotate them around to generate new ones all the time. And the tracking protection is essentially reactive. You have to know in advance which domains are the tracking ones and which ones are not. Whereas, I see the Do Not Track header being a more proactive approach, where you send the header to everyone and the onus is on the tracking server not to do the tracking.

ED FELTEN: Let me first agree somewhat with Peter about the extent to which these approaches can be complementary, and address different aspects of the issue. I also think it's the case that there are some challenges that you are facing in one way or another with both approaches. Both approaches depend for their impact on the behavior of parties other than the browser vendors and the users. In the case of the tracking list approach, it depends crucially on the list makers, how diligent they are, what criteria they apply, and what means they have to tell which sites might be violating those criteria.

In the case of the header, of course, it relies on the behavior of the site and ad networks that are receiving the header, what do they do and what promises do they make, do they comply with those promises, and again, the question of how someone would know who

might be violating and what can be done about it. So, I think you come back to some of the same issues, and I think it's a very healthy thing that we're seeing different approaches proposed, and we're seeing a pretty rich dialogue about not just approach A, or approach B, but people trying to understand what's the sweet spot in designing a mechanism along with the policy framework.

JASON SCHULTZ: So, I want to add maybe two additional components to the technological ecosystem we're talking about here, because I agree that -- I don't think any one single technology is going to help, or two even is going to help solve this. And I want to focus actually on different people in the ecosystem and the browser maker, and the consumer, per se, which is one of the things that I found in talking a lot about privacy recently, and especially around behavioral targeting and ad networks, is that a lot of websites themselves, company websites that are not in the business of advertising, or making technology, have no idea what's going on.

So, for instance, you go to The New York Times homepage with a clean browser, and say you ask The New York Times, "Who is tracking me when I come to your website?" I don't know that they can answer you. I don't know that The New York Times could tell you the answer to that question. So, actually, in terms of ease of use and enforcement, I actually think that we might need technology for sites that run ad networks, or other technologies embedded within them to know what they're doing. What the hell is going on on my website is sort of the tool I think we might need more of.

There might already be some tools out there. I don't know the full universe of them, but actually the more I talk to people who are installing Drupal or Wordpress, or using cloud-computing sites, they have no idea. And I am worried, actually, about liability for these Web-facing companies who don't make technology or don't run ad networks. I don't want to give them the sort of put your head in the sand excuse and not pay attention, but on the other hand I think the level of sophistication that some of them have is pretty limited, and we want to enable them to maybe get some feedback about what's going on on their website that makes it easy for them to work with some of these other technologies.

The second is for consumers themselves. So, one of the things that's interesting about Do Not Call, right, sort of the comparison, I think I know when I'm getting called on a phone, or at least maybe I used to. I don't know about it with my new phone. But on my old phone, I got a call. This sucks, I want to complain to somebody, and I had to sort of figure out who to complain to. But I sort of know who to complain about. And that is one of the major challenges, I think, with Do Not Track, and the FTC report talked a lot about that in there. There are a lot of conversations going on about that.

But one of the things I want to emphasize is thinking about tools that will help consumers give feedback on this mechanism using the browser mechanisms, or other things; the question is what do consumers actually think about that, are they getting good results? Can they complain? Who should they complain to? How should they complain when it's not working, or is working?

Expecting them to go into the settings of their browser is not a very good way. We have a ton of Web 2.0 feedback mechanisms, right? I know on Facebook how to comment on my friend's phone. That's pretty simple. I know how to tweet about things. I know how to do other things.

And so, actually one of the things I would suggest is that we think a lot about, are there ways for consumers to have their voice, so they're not silent in this debate, and to be able to speak more broadly, and also to maybe even take advantage of some of the social media mechanisms. So, if I've got a problem with a website, and I post it somewhere on Facebook, and my friend has that same problem, and then 50 people have that problem, the like button actually could become a kind of empowerment tool in a different genre. So, I think these kinds of technologies could actually help regulatory agencies understand where the nexus of problems are as we evolve through the technological landscape, and give feedback to the companies that are trying to do the right thing could be useful.

CHRIS HOOFNAGLE: Speaking of companies that are trying to do the right thing, your service tracks 300 tracking companies, and it's quite outstanding. It's PrivacyChoice.org. It's worth checking out. What have you learned from looking at these different trackers, and do you think that given that there's a kind of a consensus here that we do need some policy to enforce at least the header provision of Do Not Track? These 300 companies, are they going to honor Do Not Track? How do you think they're going to react?

JIM BROCK: Well, I've talked to a lot of them. Whenever we index a company, or change their listing, we contact them and let them know that this has happened. And I would say maybe 20 percent of the time we hear back or people clarify. It's been very informative. We've been doing this since early '09. The number one thing to worry about is not necessarily bad intentions because the economic value of ignoring a Do Not Track preference is pretty small. I mean, the number of people who end up doing it, we don't know, but unless that's 20-30 percent, it's not the kind of thing people will build companies around doing.

The thing you have to worry about is not intending to have a working opt-out. We find nonworking opt-outs all the time. They tend to get fixed pretty quick once we let them know, but you'd be surprised how often we find a company that just didn't implement it, isn't testing it, isn't making sure that the consumer preferences are being honored. It happens all the time.

That's one reason why I think a complementary approach that gives the consumer ultimate control, as well as being able to do a signal, I think, is very appealing for that. That's what I would say to that.

The other point I think is important, these two other points about these browser approaches that I would love to make sure are on the table. One is that part of the beauty

of the IE9 approach is that it can be embedded inside the notice-and-choice process that is already rolling out. That certainly is an approach.

But the idea that I would leave a notice-and-choice process that is Web-based without the option to make my settings permanent, that has to be fused together. That means the browsers have to interact at the Web page level to do it. It's not technically difficult to do. It can't be under the advanced tools, and everything that it's been put under. That's not going to work.

CHRIS HOOFNAGLE: Could you explain that? I'm sorry. I did not follow exactly your point there.

JIM BROCK: What IE9 lets you do, and we'll be showcasing this tomorrow with the PrivacyChoice Tracking Protection List, what it allows you to do is embed that action, the act of turning it on, in a Web page. That's super powerful because I don't have to find the advanced tools menu in the browser; I don't have to read an article and then go to my computer. I just come to the ad notice that is already being rolled out, and industry is growing up around these ad notices, and these website notices, and punch that button and I'm done as a consumer. I've sealed the deal. My preferences are now going to be honored whether it's a Do Not Track header or whether it's actually blocking. That's got to be connected, right, and those two things have to be done, in my opinion anyway.

The second key point is, and this goes to Chris' point, too, in terms of pushing the burden onto the ad networks, I've come to believe from studying both approaches that neither of them will work effectively from an audit point of view or otherwise unless tracking networks start to self-identify tracking activity in its own cookies structure, in the subdomain structure they use, in the strings they put in. So, if you look at the tracking protection list that we're going to be showing tomorrow, we have a big list of companies that can't track, right, and that are blocked, but we have a big exception, which is if they state right in the string that they're submitting that this is not a tracking activity, then they're allowed to go through. Now, you could say, “Well, that trusts them too much.” Well, ultimately, I think that's an audit task. I think the FTC is going to have to deal with what an audit looks like to make sure someone is abiding by their promise. But I like that because that means they're telling me; they're deceiving me if they don't honor what is in that string.

And that to me is a very important point not to get lost, and that lets the whole infrastructure self-organize in a way. And if I want bulletproof protection in IE9, that's great. If I want a Firefox solution that is a Do Not Track header, that's fine, too. That's a more auditable action.

In the end, if I audit the company I'll go in and say, “All right, I want to see the activity on the Do Not Track header; when you put the sub-string in, I want to see the activity when you have this identified that you're tracking. “ That's an easier thing to do.

CHRIS HOOFNAGLE: There seems to be some consensus around the need for technology to help with tech violations and enable enforcement, whether it's Jason's model of crowd sourcing or your model of labeling the actual strings. Are there other ways in which technology can help entities like the FTC or self-regulatory organizations determine that tracking is occurring despite the consumer's wishes?

PETER ECKERSLEY: So, I get calls on a weekly basis from class action lawyers who are looking for targets. And if whatever legislation we end up getting, if we get any legislation, if it includes a private right of action, I think we will have class action firms that will be sponsoring their own enforcement. They have the money, they have the will to go out and find firms that are violating things, and eventually once those companies have been sued, then the FTC will find out about them, and can go and investigate them, too.

CHRIS HOOFNAGLE: Well, my question is a little bit different, and that is how can technology identify that this is happening? Now, we know that the class action lawyers are going to sue. So, we know that there's going to be this enforcement hammer out there. But how do we know what we don't know, and can technology help us do that? And maybe we should switch to Jon, and then Dean?

JONATHAN MAYER: I apologize for my voice; I'm just getting over a cold. So, we've been spending a lot of time thinking about the enforceability problem at Stanford. I think broadly speaking, there are two ways you can do it. The first way is kind of look for a suspicious activity. You could crowd source that; you could use a crawler. There has been academic work that shows that's doable. You could also look at ad distribution specifically concerned about online behavioral advertising.

I want to note that the enforceability problem somewhat resolves to the same technical issues, no matter whether you're looking at a header, or whether you're looking at blocking, because you have to figure out who it is you're going to block. And that, from a technical perspective, requires figuring out who is tracking, which in fact is the same as what you need to do for a header.

I guess a couple of other points I wanted to make briefly. First, responding to the UI question, I think it's no real secret that Do Not Track is now in the Firefox 4 beta, a Do Not Track the header approach. And it's in the advanced preferences settings. Everyone knows that's not ideal. That's going to get taken care of. So, I think these are user interface questions, I just want to table them as not really things we have to deal with now. We all know we need to improve user interface.

I do, though, want to pick on two specific points from Commissioner Brill's framework to highlight as a distinction between this blocking approach and this kind of declarative policy approach.

First, this question of universality. So, Dean talked about how from the Microsoft perspective they're going universal because they're developing a W3C standard. And I

want to break universality out maybe into a couple of different components. One being universality as a standard on how someone can do something, that is how a browser can implement a tracking protection list, or how a browser can send a signal. Universal in the sense of you don't have to worry about whether this is specific to individual companies. I think those are both very important. I think as a matter of standards, both approaches are well on their way. The header is moving into an IATF standardization process.

But, in terms of specific companies, I think they don't want it. Chris was talking about this kind of gamesmanship problem around domains and different kinds of sources of tracking, and that's a problem you don't have to worry about really with the header, since it goes to everyone. So, the other place I wanted to add a comment on Commissioner Brill's framework is I'm going to cheat a little bit and add kind of subscripts to the framework, which is I think it's important to keep in mind that there shouldn't be collateral in your mechanism. By expressing a choice, you shouldn't have to worry about the very expression of that choice breaking other things.

In particular, I think something that hasn't been addressed here so far is how blocking actually breaks other things. So, for example, if you were concerned about the Facebook like button, and you wanted to ensure that Facebook wouldn't track you, they could deliver you a like button that doesn't involve tracking. But, on the blocking side, the solution is this kind of hammer of, OK, we're just not going to load the like button.

With that, I think I've kind of exceeded the time for my response, so I'll stop.

DEAN HACHAMOVITCH: Chris, your question was around how can technology help find the tracking that's going on. Is that fair? I think that there -- I think everyone on this panel wishes they could offer you a magic bullet on this, and I think there are basically two issues that I see, and it would be great to hear what other people think. One is that we have no consensus definition of what actually is or isn't tracking. So, you get into this fun, almost Lewis Carroll-like game of, well, you mean analytics, too, or what do you mean by tracking?

The other issue is that technically we are very, very good at identifying what something is in terms of the technology. We can identify an image. We can identify a JavaScript. We can identify a cookie. We can identify objects that way, asking the technology to identify is this tracking or not is a very different exercise. I think a real-world analogy, which will, of course, fail and if nothing else, it's up to everyone else here to offer a better one, is it's very easy for browsers to identify stone, versus drywall, versus wood. It's very hard for browsers to identify kitchen counter versus kitchen floor, versus bathroom counter, versus steps. And those almost semantic uses of material is, frankly, beyond the capacity of plain old code. It's why human intervention, either in the form of a developer on the site, or in the form of a curator of a list, is necessary. I'm curious if people have different opinions about using technology to find tracking.

JONATHAN MAYER: I want to note that I completely agree with Dean that there's going to need to be at some point here some sort of a human component. So, for

example, if we either call certain practices tracking, but allow them. So, for example, certain forms of fraud detection you might want to allow, or on the other hand certain things that might look a lot like tracking but aren't actually. So, for example, you might imagine someone provides analytics for a page and they're a third party, but in fact there is some sort of contract around we promised not to actually give anyone that information. Those could be difficult cases to detect with pure technology.

On the other hand, I don't think the situation is quite so grim. I guess a good fact on our side is that to be good at tracking, to offer useful data to people, you have to be pretty big, you have to be on a lot of sites. It's not that useful to be the tracker on just some obscure website. And the corollary there is it's not so easy to hide. So, the kinds of things you need to do, things like whether it's a unique cookie, or kind of maybe more towards the particularly suspicious end of things. If you start asking a user, please tell me all of the fonts you have and what order you installed them in, chances are good you might want to look into what they're up to.

So, those kinds of approaches can actually give us a pretty good idea of at minimum where there needs to be human follow-on enforcement, and in some cases just a pretty clear indication that these guys are up to no good.

JASON SCHULTZ: I just want to add, so this is maybe a little bit more of where I was going in the sense of what I'm saying is that, I think a one-time choice from a consumer is a hard thing to then assess how it is affecting that consumer's happiness, or expectations, right. So, if it's about -- we want to make it simple and that it's a one-click. But, on the other hand, we don't necessarily know what's happening if we only make it a one-click.

So, this is why I think some feedback mechanism or some crowd-sourcing or some other tools might actually be useful because if we're going to really go with, at least from the FTC's perspective, consumer expectations, consumer expectations change over time, they depend on whether the expectation was met. So, again, I think that we need to assess, and technology will tell us I think even in semantic context if we see consumers respond to it and understand the response.

So, I would just say, I have no doubt that the number of usability engineers out there, including ones who graduated from this university, have already helped solve a lot of these problems of understanding what consumers are doing and what they're experiencing. And I think a little more of that application here in the right context can actually help identify what technologies are creating problems for consumers, and then we can look a little more deeply in them, because yes, I mean, if it's about the kitchen counter versus the kitchen chair, but the consumer is like, well, I don't really care about that difference so much, I figure it all out, versus a hole in the counter that something is falling through, then I think consumers will be quick to complain about it.

ED FELTEN: In talking about enforcement or detection, it's important to sort of zoom out and understand what it is that one would need to achieve in order to deliver a lot of good to consumers. This is not one of those cases where you need 99.99 percent

compliance to generate a benefit. If you get most companies complying most of the time, if you get even that, you're actually reducing the potential harm and trouble for consumers significantly because the number of sites across which they can be tracked would be considerably lower. So, you don't need to get to 100 percent success, and indeed, also depending on the overall policy framework you may not need to -- you may be able to deter tracking, certainly by the larger players who are more important from a consumer benefit point of view, without getting 100 percent certainty of touching them if they violate. You just need enough to deter bad behavior because they want to avoid being shamed or, if the legal and policy rules are set up that way, being punished.

And I think, to put that together with the work that Jon is talking about, about technical methods for finding and detecting tracking, I think it's a more hopeful situation than you might think. As technologists, we like to say, we want to reach 100 percent success all the time, and not be satisfied with less. But, frankly, I don't think we need to get there on this issue to get to the policy goal.

DEAN HACHAMOVITCH: Actually, to what Ed just said, it goes back to Harlan's first statement, which is what is sufficient? And where do we put the bar for sufficiency? I want to be clear that there's a tremendous opportunity and role for policy in this space regardless of which of these two technical paths goes forward. There is a huge role for policy for legislation because it's not clear right now what is compelling about honoring the user's preference. And so there is definitely a role there, again, whether it's header, tracking protection list.

I think that there are some historical examples that show how these things will play out, and we could have a beacon, the Web browsers indicate the user would like to not suffer phishing attacks. Do not pretend to be Citibank when this user visits your site, please. And I think it's a pretty good idea.

In addition to that beacon and that messaging system, it's nice to have an actual protection, a kind of phishing filter. You can imagine the same thing for attack after attack after attack, and I think consumers today are somewhat skeptical, or at least worry for very good reasons, because as they go around the Web they end up seeing the most obvious ones, phishing attacks, malware attacks, cross-site scripting attacks. The list just keeps growing day after day after day.

CHRIS HOOFNAGLE: I actually want to pre-empt you and ask a new question, so that we can move on to some other issues. One of the issues that Ashkan brought up was basically displacement. And, Chris, you commented on this earlier, if you block cookie tracking, then people move to Flash cookies, et cetera. One issue that Ashkan raised is the displacement towards first-party tracking. How realistic is that risk, and what can technology do to kind of detect false first party tracking?

And just as an example, there are some rather large players out there that consider their third-party tracking to be first party tracking. So, you know, you visit some third-party website, you're on the Washington Post, but let's say the analytics provides, and some

others still consider that tracking they're doing on that site first party. Can technology detect, do something about these things, should it?

PETER ECKERSLEY: Sure. I think there will be some important, I don't know whether I would want to say fights over this, but it will be important to get this right. The big nice example in the room is Facebook, and Facebook is interesting because it has this first-party/third-party inversion phenomenon that it engages in. You're familiar with Facebook as a first-party website, and perhaps either by -- we don't expect we can fix the problem that you're tracked when you're on Facebook. It's not great, but Facebook keeps a record of whose profiles you look at, and what messages you send to other people. But that's not the problem that we're trying to deal with at this stage of the policy game. We're trying to deal with the few hundred third parties who see everything we read on every website, virtually. That's the bigger problem. We'll try and solve that one first.

But it's essential to not let Facebook say, "Well, we were a first party when you logged into our social network, and so now we're a first party on every website you go to." That seems to not make sense, and we need to draw the line clearly, and say, "No, Facebook, actually you're a third party when you're on someone else's site, and you have to respect the Do Not Track header if it sits there unless you've gotten extremely clear consent from that user." They absolutely understand, yes, I am permitting Facebook to track me on every website I go to, and there's some reason why I want to say yes to that.

CHRIS HOOFNAGLE: OK, but you're calling for a policy intervention. Is there a technical way of, let's say, sorting out the difference between the cookie you need to log on to Facebook, and a tracking cookie that Facebook might be using elsewhere on the Web?

PETER ECKERSLEY: I think, yes. There is one URL that is in your URL bar, and Facebook can play games to try and get around that, and we've seen a little bit of that. There are some documented instances of advertising service aliasing their servers into a first-party domain. But I think if you write a policy rule, then those weird attempts to game it look pretty obvious.

CHRIS HOOFNAGLE: I'm going to pose one more question, and then turn it over to the audience. So, we do have our great BCLT people out there with microphones. But, until they get to you, I want to ask a question about where opt-out mechanisms should be located? It seems especially with the rise of mobile devices, and the fact that people are tracked across multiple channels, not just the Web, but they're also tracked over other programs. Maybe these privacy mechanisms should be in the OS — any comment on that, instead of the browser? I mean, the tracking is deeper than the browser sometimes.

EDWARD FELTON: Sure. I think there are settings where it might make sense to do that. Mobile settings in some cases you can get broader coverage by putting something into the basic interface of the device. That is where the choice mechanism is located. I certainly see the argument for that.

I want to be cautious here because even though I'm not a spokesman for the FTC, or any commissioner here today, I don't want to give the impression that the FTC has its one favorite approach that we are trying to push on everybody. We really are interested in hearing people talk about different approaches, and debate the pros and cons.

But, having said that, certainly a discussion of whether these things should, in some settings, be offered via the device OS is an interesting one.

CHRIS HOOFNAGLE: Dean?

DEAN HACHAMOVITCH: I think it's a great question. I think the mobile example seems a lot easier for people because you just have it there. Over time, and it might be a very short period of time, I can totally see it making sense to put that in the operating system.

I think this is an interesting question for you, if there were a Windows system setting called Do Not Track Me, would Mozilla on Windows respect that? If there was a setting like that on the Macintosh, would Mozilla on the Mac respect that setting?

SID STAMM: It depends on what you mean. We don't track either.

DEAN HACHAMOVITCH: No, respect that setting in terms of sending indication to sites, et cetera.

SID STAMM: Yes, if there's a clear way in the operating system to turn on and I don't want to be tracked kind of signal, I don't see why we couldn't just hoof it into what we're currently doing.

DEAN HACHAMOVITCH: So, you have at least one example of consensus.

CHRIS HOOFNAGLE: And it's too bad Apple isn't here to have more consensus.

So, Jim, Jason and then Chris, and then we'll move to the audience.

JIM BROCK: And I'll keep it quick just to say, I think it would be great if it was at the OS level. Those companies also are in advertising, so it will take a while, I think, to do that, and I'm not sure there's a mandate is necessarily for it. The prototype we've built tries to do something in the interim, which is simply to have an app that transmits your UDID, the nice thing about mobile is it's more like Do Not Call than like Do Not Track, in the sense that we know what the device is, you can register your ID for that device. And that registration can be available to applications or websites to know not to track you on that site.

Again, that's a lot like where we are with browsers. It's very, very early, but there's no reason not to solve it, and I would hope the FTC would address it in their guidance in terms of how this all plays out in mobile.

EDWARD FELTON: If I could just interject here. Certainly the FTC would welcome comments from everyone on these issues.

JASON SCHULTZ: So, I just want to sort of point out that I think this is another place where Do Not Call and Do Not Track, despite their semantic alignment, aren't necessarily on point, because it does seem to me like Do Not Call was looking sort of at a system where any phone, it doesn't matter what phone you're using, it's like you don't want to be called on any phone. And part of that is because of the consistency of the infrastructure, and the technology.

And one of the things that I think we're trying to struggle with here, and I know this is in the report, too, is that what happens when you have innovation in the infrastructure. So, right now, we're talking about browsers and operating systems and mobile, which is interesting that we're not saying what technology mobile is. So, I do think that there needs to be a way to accommodate that, but I do think that that may also be an ongoing conversation, right.

So, again, just like we're not looking for a single technical solution, the question of where to put the flag, or the header, or what the form of the signal from the consumer is going to be, I think, also needs to have some flexibility so that as consumers gravitate, maybe even dramatically from one infrastructure to another, they can take that signal with them, and it doesn't have to become a renewed conversation each and every time.

CHRISTOPHER SOGHOIAN: So, I just want to quickly talk about the role that the first parties can play in making this happen. So, we saw more than 10 years ago IBM saying, "We won't spend any of our ad dollars with networks that do not have a privacy policy." And almost within a year or so every ad network had a privacy policy, because they all wanted to be eligible for that revenue. Similarly, Wal-Mart has said, "We won't buy any products from suppliers that don't embed RFID tags in the shipping containers because we want to be able to track them."

First parties can move the market in ways that no one else can. And if Procter & Gamble and Wal-Mart and some of these other big companies that actually spend the money said, "We will only send our ad dollars to network that respect the Do Not Track header, or respect some kind of Do Not Track mechanism," every ad network would suddenly get onboard.

We've already seen that there are advertisers that will only agree to spend ad dollars through networks that are members of the NAI because they want to seem like responsible players. There is a way to pressure the first-party sites, and the first-party companies to shift the market even if the ad networks don't want to come along on their own.

Sorry for the derail.

SID STAMM: I'd like to add a little more context to the operating system question. A number of antivirus, or antispyware and anti-fraud software vendors, have rolled out things that detect tracking cookies and delete them. That's based on a blacklist approach. But this is an example of something that's done more or less at the OS level to protect users. And so, there's no reason that we couldn't go down there and do it there. In fact, I think it's great to have multiple layers of defense, and I want to encourage that.

And I also want to mention something to Dean. You mentioned that maybe we should send out a beacon that says, "I don't want to be phished." There's a slight difference between I don't want to be tracked and I don't want to be phished. A large number of the players in this arena for tracking want to play fair, they want to honor consumer's preferences, whether it's voluntary or not, they'd like to do it. They want to have a way to do it. Phishers, I don't know any who are going to respect an opt-out.

CHRIS HOOFNAGLE: Do we have any audience questions? And, you know, it's fine if you don't because I could abuse these kind people forever. Ashkan already got a chance to speak, how about Mary.

QUESTION: Hi, thanks Chris. I'm wondering if instead of Do Not Track any of you guys have considered users can only be the ones who track themselves. So, that could exist in a personal data ecosystem where instead of anybody else tracking me, I can decide never to track myself, or I could decide to track myself through my own data locker. And advertisers, instead of getting, say, 50 percent of the data wrong from the Rapleaf and Intelius companies of the world could come to me and make a deal with me where I'm now a part of the transaction. And I can be involved and say how much I want to share. Do I want to be pseudonymous, do I want to be completely anonymous? How can I be involved, in other words, in my own transactions, and how can I say where my data goes, and how can we create an environment here the technology supports me being in control, and nobody else tracking me but me?

CHRIS HOOFNAGLE: So, this is a consumer data cooperative question. Is anyone working on this? John?

JOHN MITCHELL: So, Stanford is, and so is Microsoft, incidentally. So, I want to note the steps that are being taken in this direction are kind of just baby steps, but there are proposals that start to move there. So, the two I want to highlight at Stanford, there's a project called agnostic, this idea where you build your interest profile into your browser, and then it's up to you to decide whether to share that interest profile with an advertiser, and allow them to then target an ad based on those interests.

At Microsoft Research, there's a project called RePriv, which is an effort to go a little bit beyond that, and in an automated fashion learn the things about the user, build a data store I think very similar to what you're thinking about, and then selectively share information out of that data store.

I want to note that this is a direction we'd love to see things move in, and this is one of the reasons we're really excited about the header approach. If you signal that you would like to opt out, that then gives the opportunity for the market to respond with privacy-preserving ways of still targeting ads, or getting analytics, and possibly even getting more correct and more information you might not otherwise by building trust with that consumer, as opposed to blocking them outright.

JASON SCHULTZ: I think it's an interesting idea, but I still think that from an information transaction cost point of view, right, with every consumer having to think this through, I think there are some challenges to getting adoption. Whereas, I think that we've seen some interesting social models come up with lots of other information transaction costs, even something as straightforward as Groupon, right, for coupons and things, where I think you could get some uptake on that.

So, I think it's an interesting model, but I would -- again, it's like if we try and isolate each consumer and make them assume all the costs of figuring all of this out, I worry a little bit about uptake and adoption. Whereas I think if you do empower them to a certain degree, but then you sort of allow these sort of innovators on top of that to sort of interact with them as "I am your spokesperson" sort of thing, you might actually get a broader adoption.

CHRIS HOOFNAGLE: Nathan Goode (ph) is going to have the last question, because we have two minutes left.

QUESTION: Hi, just to play devil's advocate, do you think that binary choices like Do Not Track can actually lead to more insecurity? So, for example, say I turn on Do Not Track, the next day I can't read my e-mail, I can't log into Facebook, I can't do anything until I turn it off. And then to turn it off, I have to enable all sorts of horrible tracking that I didn't want initially and the ad people then kind of win by default, and they also win by the fact that they can say, now the user has affirmed to all these additional things in order to use our service.

PETER ECKERSLEY: One nifty thing about the informational, like the header style version of Do Not Track is that it doesn't actually break any of the underlying mechanisms. And you can do the opt back in. I think all of them essentially have a way of saying, in general, "I don't want to be tracked, but I'm OK with accepting a particular bargain that involves service in exchange for some tracking that I understand." And those opt-ins can use all the existing Web mechanisms for implementation. You can use a cookie, you can use a local stored object of some sort, you can do those things. And so the ball is in the court of the service providers to say, this user didn't want to be tracked, but maybe this service here is a compelling one that we'll be willing to be tracked for it. Here's a user interface for opting back in, make it non-confusing and everyone is happy.

CHRISTOPHER SOGHOIAN: So, I actually am worried about that scenario and I think the solution for that is strong regulations prohibiting companies from discriminating against users who have said no. I think Jonathan has a lovely blog post that he wrote

recently detailing how small the actual percentage of behavioral advertising is, as compared to the overall advertising market. These sites already provide free e-mail and free news content to us without needing to rely on behavioral advertising, and I think consumers should be allowed to say they don't want to be tracked everywhere on the Web without suffering a loss of e-mail, or a loss of access to the content and services that are a part of our lifestyle, right. You shouldn't be cut off from the communications mechanism that is Facebook because you don't want to be tracked everywhere on the Web.

CHRIS HOOFNAGLE: Harlan, just 30 seconds.

HARLAN YU: I think the browser vendors need to innovate and experiment around the user interface, how to make this easy for consumers to control. I don't think it's an all-or-nothing thing where because I need my e-mail I have to turn Do Not Track off for my whole operating system or for every single interaction with every single party. And I think even -- certainly there's a fear where servers will deny you certain content, or full access to their site if you have the Do Not Track header on, but I think in the end that at least provides transparency for what the bargain is between the user and the server, rather than having that decision made automatically for the user, without them knowing exactly what they're giving up in terms of their privacy.

CHRIS HOOFNAGLE: OK. With that, let's thank the panel.

Thank you for your time today. (Applause.)

Discussion Topic 3: The Implications of DNT

James Dempsey, Center for Democracy and Technology
Pam Dixon, World Privacy Forum
Alex Fowler, Mozilla Corp.
Susan Freiwald, UC San Francisco, School of Law
Beth Givens, Privacy Rights Clearinghouse
Sue Glueck, Microsoft Corp.
Betsy Masiello, Google Inc.
Chris Mejia, Interactive Advertising Bureau
Joe Ridout, Consumer Action
Chuck Teller, Catalog Choice
Lee Tien, Electronic Frontier Foundation
Jen Urban, UC Berkeley, Samuelson Clinic

Moderator: Paul Schwartz, BCLT & UC Berkeley, School of Law

PAUL SCHWARTZ: So on our panel that we'll have now, what I would like to do is it seems to me topic one, we heard about the role of the FTC and different policy approaches. On topic number two, we looked at technical approaches and the interplay between what technical people think they can do and what they think the policy approach should be.

So, the way I'd like to proceed today is give everybody on the panel maybe about five minutes, and this is kind of my question to this panel, and maybe we'll just go around and hear briefly from everybody, and I will interject at times and maybe be Oprah, but just kind of give everybody a chance to talk.

So, this is kind of my opening question, which would be, you know, what does tracking mean? OK, what does tracking mean, and can we make a distinction between, for example, tracking for analytics and other kinds of tracking that get carried out where it's more directed towards customers and shaping the customer's choices?

And then I guess I would also say, the other question to think about is simply what does PII mean? OK, what is personal information and how does that feed into the whole debate about Do Not Track?

So, why don't we start off with Betsy, and let's go around the table, and just everybody I'll try to hold you to that five minutes.

BETSY MASIELLO: Thank you. It's a really good question, what does tracking mean.

I think there's a couple different sort of ends of the spectrum that I hear discussed when I come out and talk to folks about Do Not Track.

On the one end of the spectrum is a really broad sort of discussion around personalization, and whether or not users actually want a personalized Web, and how we can give them options to modify how much personalization they're seeing.

I think most folks in the room who are looking for a really practical, reasonable, short-term sort of "what can we do soon to give users these choices" are focused on behavioral advertising, right, and the use of data collection to compile profiles about a user's interests, and serve ads to them based on those interests, right?

I think for the most part the industry is trying to very hard to solve that problem, and you see a lot of the browser companies putting out new solutions that could be used to solve that much more comprehensively.

Today, the way that this is done is through cookies, and for the most part -- I'm not going to speak for the entire industry, but for the most part this type of behavioral profiling is done through cookies.

Many of the ad companies are members of the NAI, and the NAI has set some standards on what it means to behaviorally target a user, and what sorts of choices you should give them, and mandates that if you're compliant with NAI, you have an opt-out cookie.

So, the only solution that I know of that's on the market today to comprehensively stop Do Not Track or give users a way to stop tracking is what Google announced two weeks

ago with this Keep My Opt-Out extension for Chrome, which persistently stores those opt-out cookies for NAI member companies. It's not perfect, but it works today. You can go install it on Chrome, and it actually will functionally stop that type of behavioral targeting.

I think that's a first step. It's where we can get to today, and what we are able to do right now, and we're going to continue having this discussion over the next several months, if not years, around sort of the broader question of personalization, and how much personalization users want, and how we can give them granular choices in defining that.

I think figuring out that first piece is really important. How much personalization do users actually want, because that's what we're talking about, right? The collection of this data enables advertisers to show you relevant ads, but it also enables Facebook to make your experience on Facebook much more useful, right?

So, someone was mentioning earlier today they didn't really like the idea that Facebook was storing all the clicks that this person was making on their friends' pages.

On the other hand, Facebook storing that enables Facebook to figure out who to show you in the news feed, right, that these are the people that you seem to interact with the most, and these are the people whose content you think is relevant to you, and it's what you want to see in your web experience.

So, I think we have a long way to go as an industry in figuring out what users actually want here, and how much personalization is really ideal for them, and I hope that we will be engaged in that conversation for several months together.

PAUL SCHWARTZ: Excellent.

Pam?

PAM DIXON: I was so hoping you'd start at the other end. (Laughter.)

That being the case, let me go ahead and just say a couple things.

I'd like to zoom back a bit. We've heard a lot of detail on the panels to this point, and I want to roll back three years to when a group of privacy advocates actually met at Berkeley. Chris was kind of enough to host the meeting in a little room, and we all sat there for two days, and we hammered out what our thinking was in advance of the Federal Trade Commission workshop on online behavioral advertising and profiling, targeting, whatnot; the idea of the meeting that everyone was invited, every single member of the privacy community -- and for those of you who are kind of insider privacy geeks, you know that was no small accomplishment. We had CDT and Jeff Chester in the same room, so that was a big deal. There were arguments and there were fights, and it was really great. A lot of the people in that room, I see some of us here on the panel today.

What I want to say is a couple things. Instead of looking at online tracking in an esoteric manner, what we were really thinking about at that table is, what is the consumer experience, what is Joe and Jane doing when they're sitting at home at night watching TV with their remote control and using their iPad or whatever to e-mail or to cruise around looking up the ads for that car that they thought was pretty cool, what's their experience, and do they know what on earth is happening to them while they're having it? And the answer was no, they really don't, they don't know.

The next question is, well, what's the result? And the way I like to put the result is the modern permanent record, and that's really in answer to your question, that's really what tracking is. Tracking is the modern permanent record. It's the new form of what used to be called by schoolmarms, you'd better watch out because that will go on your permanent record. Therefore, every time you didn't get to go out for recess because you had a bad math grade or a bad whatever grade, that was somehow recorded and influenced your college choice or perhaps your career choice.

So, the idea is that tracking should not influence your modern permanent record, and to do that we thought of the Do Not Track list. How could a consumer make an affirmative choice, without having to know all of these details? I don't think any consumer, unless they're an expert, can really know these details.

So, that's really the concept of what is tracking.

And just along those lines, I think in terms of how to distinguish between tracking, I think a lot goes to intent. What is the intent of the anti-fraud mechanisms that so often get a loophole in the current regulatory system? What is the intent of those mechanisms that don't get that loophole? So, I think that that goes far to answer.

But something that's really, really hard to pin down is this: One of the issues with privacy in general is there is a tendency for the privacy community to just say no, and the industry to just say yes, when we both need to be saying both things.

And I'll just only speak from the privacy perspective, which is we really do need to say no to tracking that creates a modern permanent record that can be used to create harm for an individual or to create lots of opportunities, or to provide them an experience that limits something that they may otherwise have seen, noticed or experienced. And there's a lot of ways that that can be seen.

But we also need to provide an affirmative way for consumers to express what they want to have, not just to say no but to also use Facebook, use Google, use these products, and say yes to the products without being damaged by the use of them.

So, those are some big picture things, and I suspect we'll get to the details later.

PAUL SCHWARTZ: OK, before we get to Chris, just to kind of quickly summarize, it's interesting that Betsy pointed us to the broader question of personalization, and really Pam picked up on that theme as well in talking about choice. And so it's interesting to try to think about what's the common ground here, and not have kind of one party say, we want less, defined as I guess less information collected, and one party saying more, you know, or everything. Excellent.

Chris?

CHRIS MEJIA: So, you know, to be honest with you, I'm not sure that I'm qualified to answer the question, what is tracking, based on one thing. As I was walking up to campus today, I realized -- you know, I put on my suit, I'm staying in San Francisco at a hotel -- I realized it's the second time in my life that I've worn a suit to this campus. The first time was my graduation in 1993. So, that puts me at 40 years old for you that can't do the quick math.

And when I think about this, and I take a look around in the room, certainly the folks that are on this panel, I'm not sure that we're age-qualified to make this distinction today with regards to what is tracking.

What I'm leaning toward seeing in the industry certainly from the younger folks that are taking the lead in terms of use of all of these products and Web pages and Internet experience that we're talking about is that tracking is simply the currency of their exchange on the Internet, OK?

When I'm talking about that, I'm saying -- I'm looking at my friend's 16-year old daughter who has no idea what this word tracking means, but freely engages in this Internet experience.

Simply stated, I think what tracking is today is the Internet experience that we've all become accustomed to, so receiving your box scores, your social updates, your stock ticker, the advertising that's relevant to you and to your social network and to your paradigm; I think that's what tracking is today. And it's a currency that's freely traded by the younger generation.

At age 40, yeah, I'm a little bit concerned about my privacy. The folks that are 20 years behind me I don't see the same sorts of concerns, and I think that there's a whole new paradigm for what's going to happen.

So, I would urge that we not just talk about this in the context of months, but that this is going to be an ongoing conversation for a millennia.

PAUL SCHWARTZ: It's kind of great to be reminded how old one is, right, always kind of painful. And as a law professor you become aware of that, because you look at your notes and what you think are topical references to things like Notorious B.I.G. and Nirvana, right, 1991, and you realize that time has marched on.

So, there's a real danger, but I think it's not only the danger that Chris presents, but it's also a regulatory danger that we're essentially fighting the battles of kind of the last technology. So, I think that's part of what's going on here, that we may be worried about websites, and websites are so 1998, which is the year of COPPA, right, the Children's Online Privacy Protection Act, and it may be that the privacy concerns of the 21st century or the early part of the 21st century are things like viral marketing and immersive digital experiences and so on.

Lee?

LEE TIEN: So, I guess I'm going to be a little bit of a contrarian here on a number of points.

Sort of the first point I wanted to make is we've got some evidence right now, done by people sitting very close to me, that strongly suggests that that generation gap in privacy isn't what's the problem that some folks would like it to be, and I think that if we look at any of Dana Boyd's work, you will see a very strong sense that even young people who use social media and seem to be sharing a great deal about themselves are actually trying very hard to cope with a world that has shifted from the one that I grew up, and I happen to be much older than you, where instead of being private by default and public by effort, it's now the other way around, and we're coping with that.

The other place where I want to be a little bit different is on the question of why we care about tracking in the first place, and therefore what it's about.

I mean, we have talked a lot about advertising and tracking, and one of the things about our position at EFF on the kinds of surveillance that's going on here is that it's not just about advertising. In fact, we don't have a negative position towards advertising. We are worried, though, that the kind of ongoing, routinized, widespread, ubiquitous surveillance of people's activities on the Web creates mass quantities of data in private hands that also will flow seamlessly and invisibly into the hands of government decision-makers and policy-makers who also have very, very different kinds of uses for the data that we're often not ever going to know about until it's too late.

There is a -- it is not possible in today's world to really separate the kind of data surveillance threats on the private sector side from those on the public sector side.

So, as an organization that's really trying to champion privacy, as well as expression of freedoms on both sides of the public-private divide, you know, we really don't think it's a good idea to frame the question of tracking purely from a sort of advertising or even a commercial side. It's all one piece.

That said, then therefore when we think about what tracking is, and I'm hoping that someone later down the line will be able to do a better job of defining it, but all I want to say is that if we are going to think about what it means to track, it only makes sense to

think about it in terms of the threat model in the first place. What are you trying to make sure that citizens and consumers are able to do? And that, what should they not be watched about in terms of — I think Peter mentioned one of my favorite articles of all time about the right to be anonymous. This is what we are trying to -- part of what we're trying to protect.

If there is tremendous widespread surveillance of what you read, you don't have the right to read anonymously, and that information will be in the hands of law enforcement and other government officials without you knowing it.

So, to some extent I'm not answering the question, I'm really saying that this is what we have to think about when we want to answer it, and that's also part of why I think that we shouldn't be trying too hard to work out the perfect definition of it. It's going to change, it's going to shift, depending on what kinds of threats there are to civil liberties and privacy. So, we need to have sort of this concept and not worry too much about the details.

PAUL SCHWARTZ: That's a really great point we need to point to what Chris Hoofnagle has called the relationship between Big Brother, namely the government, and then lots of little brothers that also collect information.

Chuck.

CHUCK TELLER: I'm going to try to answer this kind of from two perspectives or two different hats.

I had the opportunity to run a very large portal where consumers come and make opt-out requests to brands, 3,100 brands. So, I talk to direct marketers every day. And we're really focused on name, address from a PII perspective, which I think is key.

So, from the direct marketer perspective, the collection of information is about targeting, right, not necessarily tracking but targeting, so that we can get the relevance, so that they can actually send you an advertisement that you actually buy. So, it's about prospecting so they can give you relevant ads. So, I sound like the DMA for a moment, but that's really what it's like from their perspective, right?

From the consumer's perspective I look at this and I say, well, tracking, if I'm with Safeway, and I give them a Safeway card, it's about a loyalty program, so there's a quid pro quo. I have given you the right to take my information and trade it with manufacturers so that when I'm at the checkout stand I get money off. So, there's a quid pro quo in that perspective.

Tracking, however, without my permission, the collection of information without my knowledge, that perhaps is kind of really what's driving this whole discussion.

So, when we reached out to our consumers and asked them about this issue, about did they -- to be honest with you, first of all, they barely understood it. They barely understand what is going on. And I think what's important to try to understand is what are their basic expectations when they do business on the Internet.

So, we actually recently did this survey to try to get this, because I don't want to make it up from anecdotal evidence. So, I basically asked consumers, "When you make a purchase online, do you expect that that company will then use your e-mail address for marketing, will that company use your mail, name, address for marketing, would that company trade your information with affiliates, will that company trade your information with third parties?"

And what we learned very quickly -- this is a survey of around 2,000 out of our sample -- is that the norm, you know, it's not unanimous but the norm is more towards they don't expect those things to occur, they don't expect those things to occur.

And so what we have an issue here, I think about when it comes to tracking, when it comes to collection of data, when it comes to the use of data outside of the notion of where I just made that purchase, of educating the consumer and giving them proper notice and then actually getting them choice, and giving them choice mechanisms that work and giving them choice mechanisms that are not onerous to execute. And if we can put those things together, the direct marketer and the consumer can actually coexist in an environment where we're collecting data, businesses are expanding in our capitalist environment, and we can prospect, et cetera.

But I think that we've got a long way to go from an education perspective; we've got a long way to go from an execution perspective.

PAUL SCHWARTZ: So, let me ask a quick follow-up question to Chuck, and then we'll go to Susan. What you're describing is kind of what in privacy law, in European privacy law, we call purpose specification, that if the information is collected for one purpose, then the company doesn't use it for an unrelated purpose. And Chuck is suggesting that the polling information shows that consumers may not be happy about that further use and view it as foreign or extreme.

But I guess let's lead it back to analytics. Did you have a sense, did you poll about whether consumers feel that it is appropriate for information that is collected about their Web experience that would be used, for example, to improve a site or to target offers to them?

CHUCK TELLER: So, that's great, I actually left that one out. We did ask -- the last question was, "When I make a purchase at this site, and I return to that site, do I expect that they will use my purchase information to make recommendations to me?" That flipped on the other side, that people do expect.

But there you have the kind of direct relationship between an online purchase and coming back to that online store, right? If the transfer of information is to a different channel, that people don't understand.

We have consumers who contact us on a regular basis to ask us, “How did I get on that mailing list?” and we ask them, “Did you buy from that company online?” Yes. There you go. So, it's fairly straightforward what is happening. But when it's direct -- we didn't ask the last question. The last question I should have asked, and it dawned on me this morning, was, you know, when you go to other websites, do you expect to be tracked there? My thought was not asking that question because it is so esoteric at this moment because most consumers don't even understand it. So, we actually didn't feel like we could actually pose the question properly.

PAUL SCHWARTZ: OK, great. And that was actually Pam's point as well about how actually consumers don't necessarily know too much about what's going on.

SUSAN FREIWALD: Hi. I'm Susan Friewald from University of San Francisco. I've worked on privacy and surveillance issues for the last 20 years.

But having said that, I resent the comment about my age because I just turned 30. (Laughter.) So, I'm quite young, and I can follow this stuff.

So, I have a quick and dirty definition of tracking that I just put together while the other people were talking: gathering detailed information about people, and using it to form profiles that far outstrip people's understanding and the preferences they would have if they understood the risk. So, I'd like to add a little nuance to preference.

And following up on what Lee said from the EFF, the harm from tracking that I see is partly the harm that the commissioner talked about, which is that the information is available to identity thieves and rogues and people who are going to violate your security, but there's also a harm from the government, which is that the more data that's collected, the easier it is for the government to get access to that data, and use that data in a way that consolidates power, persecutes journalists and people who disagree with the government, and lots of things we have a sorry history in this country of doing.

But even if it's not the government, even if it's not the government, the more information that is collected about you, the more there's a chilling effect on your ability to use the Web to engage in political activism, to explore information about your health, to explore different kinds of ways of living.

And so, I think there's a really important harm out there to data mining and to collecting this data that we need to address.

Having said that, and turning to some of the choices that have come up in the report, one of the ideas is privacy by design, which is, of course, a great idea.

And the problem I have with privacy by design is one thing that was already said, which is a little late for a lot of the Web technology, and secondly, I just don't see the incentives for companies to engage in this.

You know, the exceptions are the companies that are very consumer-facing, that get involved in scandals, whether it's Google Buzz or Facebook Beacon, that consumers can easily grasp and exert some pressure on, you know, voting with their feet.

But for a tremendous amount of tracking that goes on, it happens behind the scenes. And when things happen behind the scenes, there's no incentive that the consumer then can bring to change it, and the competitive mechanism falls apart -- or the use of competitive pressure to bring in better practices falls apart.

So then, that brings me to the question of self-regulation versus legislation. And, as you know, someone who's been teaching this for a long time, you know, I feel like we got to the answer in 2000, and the answer then was that self-regulation just doesn't work. I think that's that answer now, and for some reason, for about eight years there in the 2000, there wasn't a lot of legislation that was consumer protecting legislation. I can't think why that was.

But now it seems like maybe there's a chance for legislation, and I hope we don't spend too much time waiting for self-regulation because I think that the information is too complicated, there's too much that goes on behind the scenes, it's too technical, and to the extent self-regulation relies on competition to reinforce it, I just don't think it's going to work.

So, moving to my last point, which is legislation, the problem with legislation -- and this hasn't been discussed yet -- is it's going to run headlong into some other initiatives that Congress is looking at, namely data retention. Congress just had -- and CDT did an excellent job presenting in that -- a hearing about data retention.

You know, one of the best ways to protect privacy, and this has come up, is not to collect data, and some people have recognized that. But there's a tremendous interest on the part of law enforcement to have companies collect data for a long time. So, there's going to be real conflict there. And even the CALEA push is about making sure that government has access to data and information.

So, the legislation to protect against tracking online is going to run headlong into those things, which leads me to my last point, which is I think that in some ways enforcement efforts by the FTC along the unfair and deceptive practices line may be a really fruitful area to pursue because I think legislation is going to have problems.

So, I think some of the suggestions that have been made about expanding the definition of what's deceptive or expanding the definition of what's unfair could be really fruitful avenues to pursue, which I think is a mixed metaphor, and I'll stop there.

PAUL SCHWARTZ: My quick feedback to that is it's an interesting definition, forming profiles that outstrip people's knowledge and expectations. So, it's really a big element of that is transparency.

SUSAN FREIWALD: Yeah, and outstrip the preferences they would have if they understood the risks.

PAUL SCHWARTZ: OK, it's a great definition, really very promising.

And then on the kind of unfair deception --

QUESTION: Can I ask you a question? Is that the definition for all of tracking or comparative tracking?

SUSAN FREIWALD: Tracking.

QUESTION: So, the question was about what is tracking. Yeah, but what is nefarious tracking? (Off mike).

SUSAN FREIWALD: Well, I think if you add in it outstrips people's preferences they would have, then my definition is a negative.

PAUL SCHWARTZ: So, I guess then where Chris may be going is, let's say it's website functionality, and to keep a customer from like, you know, the ill-design department and the big pillar is right in the middle of the living room. So, people walk into it, and then the website says we want to gather information because we want to see if there are any such pillars, as it were, in our website where customers are just hitting the wall.

CHRIS MEJIA: I mean, it's not, but I would take it a step further. I mean, I represent the advertising industry. So, I would argue to say that it's identifying that I'm a male on my Sports Illustrated site, and I probably don't want to be served with an ad about feminine products. I don't think that there's anything nefarious about that or anything bad about that; we certainly expect that in Sports Illustrated that I don't receive those ads, and we expect that in the Wall Street Journal where I'm targeted by content all the time.

So, I just want to make a distinction between tracking for nefarious purposes and tracking for sort of legitimate business purposes or other purposes.

PAUL SCHWARTZ: OK, let's see if we've got a taker. Susan?

SUSAN FREIWALD: So, I mean, you know, the problem you left out of that equation is consumer choice. So, if I'm a consumer, yeah, I want ads that are more designed for me, but on the other hand, I want to know who has my data, and I want to know what they're doing with it.

And by the way, I just checked, and the major browsers think I'm a man, which I think is awesome. (Laughter.)

Anyway, I just looked at it from the footnotes about it, and you can see what Google knows about you and what Yahoo! knows about you, and they think I'm a man. They think I'm a 30-year old man. (Laughter.)

PAUL SCHWARTZ: Jim.

JAMES DEMPSEY: Last week or the end of the week before, January 31st, at CDT we put forth a document that attempts at a definition of tracking. So, you know, as a premise, as Pam said, we were one of the original proponents of the Do Not Track concept. We put it forth with other privacy groups in 2007.

And I do think we're at an important moment here, as Peter Swire was saying this morning, in terms of now an idea coming truly to the point of possible fruition. And in order to have a Do Not Track you have to agree upon what is track or you have to have some understanding of what is track.

The CDT definition, we have a one-sentence version, and we have sort of a small chart, just two columns, it is tracking, is not tracking under our definition. Tracking is the collection and correlation of data about the Internet activities of a particular user of a computer or a device over time and across non-commonly branded websites for any purpose other than fraud prevention or compliance with law enforcement requests.

And then in the chart we have somewhat more caveats. We actually also exempt, for example, analytics, et cetera.

Now, it's not perfect what we put out. We put it out. It's a draft. We put it out as a contribution to this exercise.

Whether you take a list-based approach or a header-based approach, a cookie-based approach I think, you do need to have some understanding of what it is that you're trying to empower consumers to avoid.

PAUL SCHWARTZ: Let's see if we have any buyers. What I'd like to do is very quickly have you repeat the definition, go to Chris, Pam and Betsy, and kind of just give me the thumbs up, thumbs down, like Roger Ebert.

JAMES DEMPSEY: OK, I'm happy to do that. (Laughter.) I'm happy to do that, although by and large I'm not a thumbs up, thumbs down kind of guy. (Laughter.) CDT is sort of not a thumbs up, thumbs down kind of organization.

PAUL SCHWARTZ: Yeah, but I'm the moderator and you're not. (Laughter, applause.)

Chris?

PARTICIPANT: Your preference is not being respected.

PAUL SCHWARTZ: Exactly. So, it's not opt-in. You gave us a definition. Chris, Pam and Betsy, just keep it lively?

PARTICIPANT: Can you repeat the definition?

PAUL SCHWARTZ: Yeah, repeat the definition.

JAMES DEMPSEY: Here we go one more time. Again, recognizing that we -- as I said, we did this a little imperfectly I think in that we have a one-sentence definition, and then we have some exceptions, which are actually not included in the definition.

But the one sentence is, tracking is the collection and correlation of data about the Internet activities of a particular user, computer or device over time and across non-commonly branded websites, excepting fraud prevention or compliance with law enforcement requests. And like I said, we also exempt analytics.

PAUL SCHWARTZ: We'll keep it short.

JAMES DEMPSEY: It's a starting point for discussion. Or is it any point for discussion. Where is it?

PAUL SCHWARTZ: So, Chris, Pam and then Betsy.

CHRIS MEJIA: So, with all due respect to Caesar over there, I'm not going to give a thumbs up or thumbs down, for further settlement of alliance.

I will say that I personally agree, and I'm not sure that this represents my membership, but I personally agree with part of your definition, the collection of non-commonly -- what was it -- branded websites --

JAMES DEMPSEY: Over time across non-commonly branded websites.

CHRIS MEJIA: Yeah, so where you get to "over time" I'd put a period and then full stop. I don't know about the non-commonly branded websites and how that's relevant to the discussion and to the definition, to be honest with you, because the Internet, the last time I checked, has around 7 billion websites, more than people in the world. So, I think we're going to find more than a few Googles and Fords and the common brands out there, and who's to say what the next brand is.

PAUL SCHWARTZ: OK, Pam?

PAM DIXON: Well, I think it's a really good point of discussion. I still like the definition that all of us came up with a couple of years ago. I like that more actually.

But that being said, let me respond in a little more detail because it deserves that it's thoughtful, and I appreciate that.

The first thing, I really have a problem with correlation because then we get into the re-identification thing, and it brings up the shape of HIPAA. That's just a whole can of worms.

We don't have to correlate in order to track, although we do and we don't. I think that needs a lot more nuance before it gets put in a sentence. It's really treacherous territory, identification.

Websites are really objective. I don't think we're talking about websites, I think we're talking about the Smart Grid, I think we're talking about cars, we're talking about mobile devices and the way people live today. We've got to get away from websites.

The exemptions I would just fight with you tooth and nail on. I think that we give anti-fraud tools and techniques far too broad exemptions. We've got to be really careful of that because there are some companies that have sprung up that use this data, that's identical often to Fair Credit Reporting Act data for purposes that I think people would not be thrilled at. So, knowing how some of this exempted data is often used I think we've got to be careful of the exemptions, and put more nuance into that.

And then finally, I guess the question is that if a browser thinks you're a man who's 30 years old, what does that do to your life, opportunities and the way your life looks when you are searching the Web in your house, doing these other things? And that's the ephemeral thing that I'm most concerned about here. How does a person's life change based on the box they're put in, based on their activities on the Web? I'd love to get insight somehow. So, those are just thoughts.

PAUL SCHWARTZ: OK, we'll go to Betsy, let Jim finish, and we'll keep going.

BETSY MASIELLO: Before I respond to the specific question, on this point of thinking that Susan is a 30-year old man, it's the advertising companies, right, it's not the browser. So, it's not affecting more than what ads you are shown, so just to make that point, but there's nothing about her life opportunities affected by that, by specific demographic inference. She might see ads for like, you know, computers that are targeted at 30-year old men, but it's not terribly harmful.

You know, I think I agree with some of what both Pam and Chris have said, Jim, in terms of your definition. I think the point around collecting profiles over time is right. The point around common branding, I agree with Chris that there is some question there. And I also agree with the point around websites.

As one example, and I'm not a retail expert, but Limited brands own The Limited clothing store and Victoria's Secret. So, where do they fit into your definition, and how do we think about data sharing in those types of examples?

PAUL SCHWARTZ: OK, back to you, Jim.

JAMES DEMPSEY: So, these are precisely the kinds of questions that we wanted to surface initially. That's why it was issued as a sort of call for response, and it was issued as a draft, with the intention of trying to be able to reach some kind of common understanding.

And I think what's critical here is that what we're trying to do is to reach some agreement on -- at least if you go with HTTP approach, and also if you're drawing up a list, different people can have different criteria for different lists, obviously, if the list is the other approach or another approach, cookies being a third. But certainly if I were drawing up a list, I would want to have a definition of this is how I draw up my list. And if I'm certainly sending out a header and wanting people to respond to it, I want to have people know what my header is.

Now, here's the point, though. We're coming up with a definition for implementation of Do Not Track. So, in other words, we're not coming up with a definition of all bad things on the Web, we're not coming up with a definition of all forms of data collection on the Web, OK, we're coming up with a definition, however broad or narrow it is, that is supposed to then inform that persistent consumer expression of intent or desire or preference that is intended to inform that such that people on the other end can respond to it or so that it can be built to a list and to the browser.

So, I think the -- well, you know, we want people to nitpick the definition, but the point is not to come up with a perfect definition of all evil on the net, the point is to come up with a definition that can be responded to with some degree of clarity, so that the consumer, who this should be all about, sort of knows what they're getting.

I think this is an important test for both the privacy community and the online community, but for browsers to some extent, the publishers and advertisers to a greater extent because, you know, behavioral advertising is not the worst privacy problem online, A), I mean, from the privacy --

PARTICIPANT: (Off mike). (Laughter.)

JAMES DEMPSEY: Well, you know, I mean, look, I think privacy advocates would generally come to realize this. We sort of skewed the whole debate by allowing it to become such a focus of attention, and it's really not the major problem. If we can't solve this, you know, then what can we ever solve by consensus or by a non-regulatory approach?

I'll say one more thing, Chris, in sort of response to your point about kids don't care about privacy. I agree --

CHRIS MEJIA: Well, I'm not saying that I think that they don't care about privacy. I said they're willing to trade in it.

JAMES DEMPSEY: Well, but hang on a second. Hang on a second, because first of all, I think there's a lot of evidence that that's actually not true, that the kids understand and use privacy controls more than -- the younger you are, the more likely you are to use the privacy controls.

But also to say that the kids don't understand what they're doing or that they have a different -- they have the wrong view of privacy or a different view of privacy leads us down a bad path on COPPA and some other regulatory measures because then the regulators could say, "Well, then we're going to make them care about privacy and we're going to raise the COPPA age for a group like CDT that cares about First Amendment rights of kids." That would actually be the wrong way to drive the policy process of saying, you know, that kids don't know what they're doing or don't care about what they're doing because then the adults come in and want to care for them, and then you get --

CHRIS MEJIA: Just to be clear, though, I didn't say that kids -- first of all, it's not kids but younger people -- don't care about privacy, it's that they care about it in ways that are different than I think an older generation has cared about privacy, and that they use it in different ways, and they exploit it in certain ways.

PAUL SCHWARTZ: OK, Jim, let's go to Alex.

JAMES DEMPSEY: OK, so then to just round it out, yeah, so again it comes back to my main point, which is we should all maybe avoid the generalization and say we've got a specific issue here, and we're trying to implement a specific solution to not the worst problem in the world, but at least let's try to agree upon solving this one problem, and there's a potential for doing that that should be responded.

PAM DIXON: Can I just say something just super-fast, and I'll actually be brief?

I think -- I guess that's really the thing that just keeps nagging at me a little bit when I was speaking at first. I think we've got to be careful because it's not that ads are evil. No one thinks that. And it's not that the evil advertising industry is somehow the worst thing that ever happened.

We do have evidence of harm. We've all read those really scary Wall Street Journal articles. And we've got to be really careful that the -- when we're talking about targeting and Do Not Track that we don't limit it artificially, that we allow it to become broader, and to allow it to stop collections that would cause harm in the way that I was describing,

and I think we've just got to be sure to keep it broad enough. But I do understand your point, and I do think it's a thoughtful aspect.

PAUL SCHWARTZ: We're going to hear from Alex, Beth and Sue.

ALEX FOWLER: Yeah, so thanks for the opportunity to participate this afternoon.

So, from a Mozilla perspective, you know, I think the way we would restate the question is, you know, will people online be able to differentiate different types of tracking. From our perspective, we started with a very broad definition of what the header means. Essentially the signal is tell sites not to track me.

And we've done this intentionally because there are a spectrum of values across our users as it relates to privacy. Some will take, and, in fact, they commingle values of privacy with anti-commercialism and really don't want to see ads and don't want to be tracked in all of its various forms. And from a technology provider, from our perspective, we don't want to say that those values should not be respected or enabled or controllable from within the browser.

At the same time, we also have users that have a much different set of values, that understand some of the value that they get through free services and free content by accepting behavioral advertising, and, in fact, some even enjoy what the ads are about as a result of seeing that connection. So, we want to be able to support those values as well.

So, I think it's important when you're designing solutions that you not project your own personal values into how you construct and enable that, and then more importantly how you communicate to users about what you're trying to solve.

PAUL SCHWARTZ: Great point, and it's about consumer choice.

Beth?

BETH GIVENS: Thank you. I'm Beth Givens, Privacy Rights Clearinghouse.

We work on a daily basis with individuals who contact us. We're kind of a troubleshooter. We've been collecting people's stories of privacy abuses and harm for -- I hate to say it but next year is going to be our 20th year. So, we've seen it go from junk mail, which was our big issue in 1992, '93, to identity theft, which sort of rounded out the rest of the '90s, to now we're dealing with employment background checks and online information brokers as people's big concern.

But I want to talk about perception. You asked, Paul, what does tracking mean. I think just general people who may not be really savvy, technically speaking, they kind of assume they're being tracked. Many of them who are experiencing problems in their lives think that there is maybe a compilation of information about them that's out there

that they just can't control, they can't get to it, they can't control it, and it's somehow causing bad things to happen in their lives.

I can't tell you how many people, who say, you know, "I just can't get a job" -- of course, we all know how the economy is -- "I just can't get a job, and I just think that probably those employers are finding things about me, and that's how they're making those decisions." You know, there may be some truth to that, especially if they're public with the information they put out on social media. But they think it's more than just social media.

So, what I'm getting to is I really like the Do Not Track approach, the header approach especially. When I talk to people about Do Not Track, they get very excited about it. They think this is it, this is really going to help me a lot, and it's going to give me a lot of the control that I really want.

And I think that's good. One of the worries I have is that when this finally does happen, and I think it will, either probably legislatively, I would certainly have a preference for that over what some would call robust self-regulation, that there might be a false sense of wellbeing when that happens. And certainly folks like us are going to have to do a lot of consumer education, the Federal Trade Commission as well, over just what it means and what it doesn't mean.

But anyway, I think that people will need to -- they need to have access to their profile, they need to be able to delete it if they want to, and we haven't talked about that. That's more of an issue say in the European Union, this business of being able to delete your tracks.

I think the two keywords for consumers are transparency and control. And I guess with those words, I will close, just transparency and control are just outrageously important to ordinary individuals.

PAUL SCHWARTZ: That's great, and we'll hear from Sue now.

SUE GLUECK: I'm Sue Glueck from Microsoft, and I apologize that I'm the only thing standing between you and a reception. And with that in mind, I will try to be brief.

Listening to my fellow panelists, I'm sort of concluding -- and not to usurp your Oprah role -- but that defining what tracking means is very difficult, and reasonable people who have a lot of experience in this area are going to disagree about it.

And I'm not sure that I care about that. What I care about is what the ordinary consumer thinks when they see a button that says, Do Not Track. What do they think that means? Do they think, OK, I'm not going to see any more ads? Do they think that it's OK to track me for certain analytics purposes? I honestly don't know, and I think we could spend a lot of time discussing it, and I actually think we should spend a lot of time discussing it. But I don't know that that gets us in the heads of the consumers.

So, instead I sort of fall back to, as the FTC report said, what are the commonly accepted practices, and I believe there's a divergence between the commonly accepted practices by the advertising industry on the one hand where people aren't surprised by what we saw, what Ashkan showed us to begin with, that the diagram of what the ecosystem looks like, and breaking down what a Web page looks like, and how many different entities you're interacting with. I don't have a parallel for that in the physical world. You know, when I go to Cincinnati to visit my brother, I'm not also in Paris, I'm just in Cincinnati. So, I think that it is difficult for people to know and understand what's happening.

So, what's the commonly accepted practice on the one hand by folks in the advertising arena, and on the other hand what do consumers believe?

So, I think the great thing about these discussions and about the tracking protection types of tools that are going to be available to people is that now it's going to actually surface the issue for consumers.

So, I don't happen to believe that once you say yes to tracking, if you're willing to be tracked -- and by the way, Nordstrom can track me all they like -- (laughter) -- and offer me discounts, and I'm probably OK with Safeway doing that as well.

So, it's about letting consumers in on what's behind the curtain. What are the commonly accepted practices behind the curtain; what are they saying yes to?

I don't think that that's going to lead to people saying yes to the really extreme things, like now not only can Nordstrom suggest items I might be interested in and tell me about sales, but they can come to my house and peer in the windows. I think that because you have to tell people what they're agreeing to, and that would be if you don't tell them that they're agreeing to some extreme tracking practice, that that's a deceptive trade practice, and we already have FTC Act Section 5 to deal with that.

So, at the end of the day, I'd like to echo what Beth said, it's all about transparency, and it's also about giving people powerful tools to help prevent you from being tracked.

And I just want to note, because I do work for Microsoft, that, in fact, the tracking protection in Internet Explorer 9 prevents the data from going to those third parties that maybe you didn't know were on that Web page. I'm not sure I know what tracking is, but I know what tracking isn't, and if the data doesn't go there, there's no tracking.

PAUL SCHWARTZ: OK, let me just do this. I want to give Alex a chance to say something very quickly, and then I want to ask Jim a quick question because Jim said -- we didn't follow up on it, but he said this is not the number one problem. So, let's hear from Alex very quickly, and then I want Jim to tell us what the number one problem is, so we can really worry as we have our cocktails and hors d'oeuvres or whatever awaits us. So, quickly to Alex, and then Jim is going to tell us what we should be really scared of.

ALEX FOWLER: So, I just wanted to -- I agree with many of the things that Sue said, but I just want to clarify something around the header because I think there was sort of a veiled reference to enabling something that communicates a preference not being the same as actually protecting you against tracking.

I just want to make it clear that certainly in the short term there is a certain reality to that. But what we're hoping is that this forces the conversation and an opportunity between the sites to make clear what they're doing with the user's information and how they're tracking them, and to give more nuanced choice and control and transparency. And I don't think you can do that if you don't actually have any information coming from that third party. So, this is really about forcing that dialogue and that transparency can only happen when those parties are communicating.

PAUL SCHWARTZ: OK, Jim, what's our number one problem?

JAMES DEMPSEY: Well, obviously, I mean, you know, none of this even makes the top 100 list.

PAUL SCHWARTZ: Just give us one. Just give us one.

JAMES DEMPSEY: Well, no, I mean, if you start with global warming and go on down --

PAUL SCHWARTZ: No, no, privacy.

JAMES DEMPSEY: -- it will be a long time before you get to this problem.

PAUL SCHWARTZ: The top privacy list. What's our top privacy concern?

JAMES DEMPSEY: Well, I do think -- I'll say something that I think is more important, which is the leakage of this data and the uses across sites, compilations of data for purposes other than delivery of ads, as has been referenced particularly to the employment context.

Dean Hachamovitch made the point -- maybe -- at first I thought it was flippant, but which is do not phish me kind of. But it is interesting , we've been thinking about, and our definition is based upon thinking about, I think, Do Not Track for behavioral, targeted ad purposes, and you still get all your -- is that right, Justin, or am I wrong? OK, sort of, says Justin, Justin Brookman, who really knows these issues with certainty.

But I do think it's interesting to think about should there be a -- particularly if you're blocking the collection of the information, maybe the use, Do Not Track me for background check purposes, Do Not Track me for health insurance. No, sort of for the FCRA it would be a big issue. So, I think the FCRA, you know, employment insurance and credit, those are to me the bigger issues.

PAUL SCHWARTZ: OK. I'm going to conclude actually on this point because we will continue over refreshments. It was a great panel, and we could go for another hour, but we'll just stop here. Thank you so much. That was great. (Applause.)

END