FTC PrivacyCon
January 14, 2016
Segment 1
Transcript

[MUSIC PLAYING]

CHRISTINA YOUNG: Excuse me everyone. Can you take your seats please? I think we're starting.

Good morning and welcome to PrivacyCon. I'm Christina Young, a paralegal in FTC's Office of Technology Research, and Investigation, or OTEC. Before we commence I have some brief housekeeping details run through with you. First, if you could please silence any mobile phones and other electronic devices. Second, if you leave the building during the event you will have to come back through security. Please bear this in mind, especially if you're participating on a panel so you don't miss it. Most of you received an FTC lanyard at registration. We reuse these so please return your badge to our event staff when you leave today.

If an emergency occurs they're requires you to leave the conference center but remain in the building, follow the instructions provided over the PA system. If an emergency occurs that requires the evacuation of the building, an alarm will sound. Everyone should leave the building through the main 7th Street exit, turn left, and assemble across E Street. Please remain in the assembly area until further instruction is given. If you notice any suspicious activity please alert building security.

We're almost done. Just a few more items. The building cafeteria is not open to the public today. However box lunches will be available for purchase in the hallway outside the of the auditorium and overflow rooms. You may use the overflow rooms to eat lunch. No food or drink other than water is allowed in the auditorium. The restrooms are in the hallway outside the auditorium. This is a public event which is being webcast and recorded. Welcome to everyone who is watching the live webcast. An archived webcast and the conference materials will be available via FTC.gov after the conference ends. And finally, we would live tweeting today's event under hashtag PrivacyCon.

Thank you and over to Dan Salzburg.

DAN SALZBURG: Thank you Tina. I'm Dan Salzburg. I'm the acting chief in the FTC's Office of Technology Research and Investigation and a member of the PrivacyCon team. We know privacy and data security are important to all of you gathered here today and that many of you are now seeing in person people who you knew had preregistered for this event. We're sorry for sharing that information with you last week and are addressing our bulk distribution set up to avoid such a release from happening again.

I hope you have had a chance to review today's agenda. We have a great and diverse roster of presenters and participants and look forward to an informative day of nonstop, cutting edge presentations covering the latest privacy and data security research. Now let's kick off

PrivacyCon with remarks from FTC Chairwoman Edith Ramirez who has led the agency's efforts to protect consumers from unfair and deceptive privacy and data security practices. Chairwoman Ramirez.

EDITH RAMIREZ: Thank you, Dan. I'm delighted to be here with you. So good morning everybody, and welcome to PrivacyCon, a first of its kind conference at the Federal Trade Commission bringing together leading experts to present original research on privacy and data security. Today companies in almost every sector are eager to scoop up the digital prints that we leave behind when we post, shop, and browse online. The new generation of products we see in the marketplace from smart appliances, to connected medical devices, to semi autonomous cars, all of these mean that consumers must navigate an increasingly complex and dynamic digital ecosystem. In short the interplay between technology and data is radically transforming how we interact with everything around us. These trends will not only continue, they will multiply.

At the FTC we're constantly seeking to expand our understanding of emerging technologies and their impact on consumers as we work to ensure that consumers enjoy the benefits of innovation, confident that their personal information is being handled responsibly. We know that enforcement and policy needs to be guided by research and data. We do a great deal of research and analysis internally but with the increasingly rapid pace of technological change and complexity of the privacy challenges consumer space. More than ever we need to tap into the expertise and insights of the research community to help us fulfill our consumer protection mandate. Today's conference provides a unique opportunity to do you just that.

With PrivacyCon our aim is to bridge the gap between the academic, tech, and policy worlds. Our ambitious agenda is filled with cutting edge and provocative research. Some of the presentations will lend support for current privacy and data security policies. Others may lead us to rethink our assumptions. Either way we hope to spur a richer dialogue about privacy and data security. And we hope that this dialogue will be a two way street. As we seek valuable input from the academic and tech communities we also aim to provide useful feedback to researchers about the type of work that would be most relevant to helping us and other policymakers make informed policy decisions.

So this morning, to set the stage for our program, and to highlight the importance of research at the FTC, I'd like to speak very briefly about the way we've incorporated privacy and data security research into our enforcement and policy work. The FTC was founded on the principle that strong research informs strong policy. Today the agency serves as a research and policy hub on a wide array of front line consumer protection and competition issues. Among them privacy and data security. As you know we've hosted workshops and issued reports on significant and cutting edge issues such as facial recognition, the internet of things, data brokers, mobile device tracking, mobile security, and mobile privacy disclosures.

Our workshops have brought together academics, consumer advocates, industry, technologists, and other key stakeholders to help inform policy discussions. And our reports on emerging technologies provide concrete guidance to businesses on how to protect consumers in today's digital world.

Most recently we held a workshop on cross device tracking. To evaluate the benefits and risks of cross device tracking we need to know what it is and how it works. Our workshop included a session where experts explained how tracking techniques function and discussed whether technical measures such as hashing might be used to protect consumers' privacy.

And just last week we issued our big data report which outlined a number of suggestions for businesses to help ensure that they're use of big data analytics produces benefits for consumers, while avoiding outcomes it may be exclusionary or discriminatory. In this report we highlight possible risks that could result from inaccuracies or biases about certain groups in data sets, including the risk that certain consumers, especially low income or under served consumers, might mistakenly be denied opportunities where the big data analytics might reinforce existing socioeconomic disparities.

On the enforcement front, the work to tech researchers has helped us identify deceptive or unfair practices of companies such as HTC, Snapchat, and Fandango. Last month we announced an action against Oracle in which we alleged that the company's failure to disclose that older insecure versions of Java would not be removed as part of the software update process. We alleged that that was a deceptive practice.

Various researchers had pointed out problems with malware exploits for older versions of Java which led to our investigation of the issue. The consent order that we entered into requires Oracle to make an effective tool for uninstalling older versions of Java available to consumers. In short our enforcement actions have provided important protections for consumers and researchers have often played a critical role in helping us achieve that goal.

In certain areas we've also asked technologists and researchers to help us come up with technological counter measures to address vexing problems. Illegal robocalls are a key example. Voice over IP technology allows callers to spoof identifying information such as the calling parties phone number. Fraudsters can now place millions of cheap automated calls for with the click of a mouse. And they can do so from anywhere in the world that has an internet connection, while hiding their identities in the process. These developments have reduced the effectiveness of the FTC's traditional law enforcement tools.

Recognizing the need to develop new solutions the FTC has held four public contests to spur the creation of technological solutions to the robocall problem. As part of these robocall challenges we solicited technical experts to help select the most innovative submissions. One of the winning solutions in our first challenge Nomorobo is in the marketplace and available to consumers. Nomorobo reports that it has more than 360,000 subscribers and that it has blocked more than 60 million robocalls.

Given the importance of research and technical expertise in so much of the FTC's work we're also continuing to build our internal capacity. Last year we created the Office of Technology Research and Investigation or OTEC as we call it. OTEC, which builds on the work of our former mobile technology unit, identifies and conducts research they can guide the development of enforcement and policy priorities among other important work. The office's interdisciplinary team includes lawyers and technologists who work hand in hand to help us study new

technologies and developments in the marketplace. With OTEC we're embarking on an even broader array of investigative research on technology related issues that will aid us in all facets of the FTC's dual consumer protection and competition mission.

PrivacyCon builds on all of these efforts. Our aim is to deepen our ties to the academic and tech communities and ensure that the FTC and other policymakers have the benefit of the leading thinking in the privacy and data security arenas. Our program today will feature five main topics. As to each, we'll have three or four short research presentations, followed by a period of discussion featuring top experts.

We'll start with sessions addressing the current state of online privacy and consumer expectations about privacy. There's no question that among other issues we need to better understand consumer expectations and the degree to which consumer perceptions of company's data practices align with what is actually happening in the marketplace. Just this morning the Pew Research Center released a study finding that Americans see privacy issues in commercial settings as contingent and context dependent. In certain circumstances a majority of Americans are willing to share their information if they perceive that they're getting value in return and that their information is being protected.

For instance nearly half of those surveyed said the basic bargain offered by retail loyalty cards is acceptable to them. While a third viewed that as unacceptable. But while many consumers may be willing to share personal information in exchange for tangible benefits, the study also found that consumers are often conscious about disclosing their information and frequently unhappy about what happens to that information once companies have collected it. We'll see what our speakers have to say about this in other topics.

Our other sessions will address big data and algorithms, the economics of privacy and data security, and security and usability. Among the issues that will be addressed will be big data and bias, the economic incentives underlying companies' data practices, the cost of cyber incidents, and available options for consumers avoid unwanted tracking. You will also hear from my colleague commissioner Julie Brill and from our new chief technologist Lorrie Cranor. And this is just to give you a flavor of what you will hear today.

We're just now scratching the surface of what is to come as a result of technological advancement. If we want to ensure continued progress we must correct policies that are built on innovative thinking and breakthroughs we make through research. And at the same time, we want to encourage research that will aid the complex and practical questions that policymakers are eagerly seeking to answer. So thank you for being here today. Your presence moves us one step closer to that goal.

Now to close, let me just take this opportunity to express my gratitude to all of the participants in today's conference. We have an incredibly impressive group of the top thinkers in privacy and data security. I'd also like to thank the organizers in OTEC and our privacy division DPIP, and in particular Kristin Andersen and Dan Salzburg for their hard work in putting this event together. So thank you very much.

[APPLAUSE]

JUSTIN BROOKMAN: Good morning everyone. Thank you very much Chairwoman Ramirez. Thank you all for coming out to our first PrivacyCon. I am Justin Brookman. I'm Policy Director of the Office of Technology Research and Investigation. We are co-presenting this workshop along with the Division of Privacy and Identity Protection. And I'm also the chair of our first panel, the current state of online privacy. If my co panelists could make their way to the stage.

So we put out our call for research proposals. We weren't really sure what to expect and we got nearly 90 really fascinating proposals. So originally we're going to try to do 12 or so. We decided to pack the schedule to have at least 19 people presenting and wish we could have done more. So we try to maximize the schedule to let them present their research to you.

They're each going to present for about 15 minutes. We're going to try to keep them aggressively to that. They have a clock right there that shows when they're over time. They'll be a chime in here that plays in over time. So they're know they're over. You'll know they're over. They'll know you know they're over. We we'll try to stay on schedule. After that is going to be a short discussion period-- my discussants, Omer Tene from IAPP and Elana Zeidi from NYU. Give a few thoughts. Ask a few questions and then that'll be it. But this is our first time doing this. Would love your feedback if we're able to do is the future and apparently we have a lot of interest. That's great.

So let me start out I'm going to introduce Ibrahim Altaweel from Berkeley to present on Web Privacy Census 3.0.

IBRAHIM ALTAWEEL: Hello everyone. My name Ibrahim Altaweel. I'm the coauthor of Web Privacy Census. Most people may believe that online activities are tracked more pervasively now than they were in the past. As early as 1999 Beth Givens of the Privacy Rights Clearinghouse suggested that federal agencies create benchmark for online privacy. The census is one such benchmark. And I'll discuss today how the literature shows a dramatic upswing in the use of cookies.

The first attempts at web measurement found relatively little tracking online in 1997. Only 23 of the most popular websites use cookies on their homepages. But within a few years tracking for network advertising appeared on many websites. By 2011 all of the most popular websites employed cookies.

In 2011 we started surveying the online mechanisms used to track people online. We called this our Web Privacy Census. We repeated this study in 2012 and 2015.

The main goal of the census is to collect and analyze key metrics and measures, and monitor the state of online privacy, and use the results to answer the following question. How many entities are tracking users online? What technologies are most popular for tracking users? Is there a shift from one tracking technology to another in tracking practices? Is there a greater concentration of tracking companies online? What entities have the greatest potential for online tracking and why?

I will delve into some detail on the data collection methods. We collected HTTP cookies, HTML5 local storage objects, and flash cookies on QuantCast top 100 top 1,000, and top 25,000 websites using OpenWPM, a web privacy measurement platform developed Princeton University. We ran a shallow crawl and a deeper crawl. A shallow crawl means that we only visited the home page of the popular websites. And a deeper crawl means that we visited the home pages and two links on the same websites.

Data collection methods of course have some limitations. For example, we used a Firefox browser. So we don't have information regarding different browsers. Another example is the crawler did not log into any websites which could potentially result in more cookies to be sent.

Overall these limitations mean that the Web Privacy Census is a conservative measure of the amount of tracking online. So much tracking to is going on. We found that users who merely visit the home pages on the top one the most popular websites would collect over 6,000 HTTP cookies. Twice as many as we detected in 2012. Some popular websites use a lot of cookies. In just visiting the home pages of popular websites we found that 24 website placed over 100 cookies. Six websites placed over 200 cookies. And three websites placed over 300 cookies.

What technologies are most popular for tracking users? One obvious observation is that are significantly more HTML5 local storage objects than flash cookies. HTML5 local storage is a new HTML5 technology that became popular in recent years for its large storage capabilities, roughly 1,000 times of flash cookies. An increase in HTML5 does not directly correlate with an increase in tracking as an HTML5 storage object can hold any information that the developer needs to store locally. However this information can potentially contain information used to track users and it can persist.

Is there a shift from one tracking technology to another in tracking practices? Beside the switch from flash cookies to HTML5 local storage, it is very interesting to see that the total current of cookies has increased and they are more and more third parties that are being used. 83% of HTTP cookies are set by third party hosts. And in just visiting the homepages of popular websites users would have cookies placed by 275 third party hosts. If the user browses just two more links the number of HTTP cookies double.

Is there a greater concentration of tracking companies online? Google's presence on the top 100 websites increased from 74 in 2012 to 92 in 2015. Percentage of cookies set by third party hosts has increased from 84.7% to 93.5%.

So what entities have the greatest potential for online tracking and why? The most prominent one is Google. We find that Google's tracking infrastructure is on 92 of the top 100 most popular websites and 923 of the top 1,000 websites providing Google with a significant surveillance infrastructure online. Google's ability of tracking is unparalleled. Most of the third party cookies are set by Google Analytics and DoubleClick. Facebook had a presence on 57 of the top 100 websites and 548 of the top 1,000 websites. This is important because companies like Google can track users almost as much as an internet service provider such as Verizon or Comcast.

In conclusion the Web Privacy Census is a modest research project that seeks to introduce reliable empirical data on the issue of how much tracking there is on the web. We have found over a series of surveys covering three years, that there's a consistent upward trend in cookie usage, and that a small group companies have tracking cookies almost everywhere on the web. In the future we'll continue to collect and analyze key metrics and measures to monitor the state of online privacy.

Thank you very much. And also like to thank my coauthor Nathan Good. Thank you.

[APPLAUSE]

JUSTIN BROOKMAN: Thank you Ibrahim. Now we're going to hear from Steven Engelhardt from Princeton University on "The Web Never Forgets."

STEVEN ENGELHARDT: Hello everyone. I'm Steven Engelhardt from Princeton University. And today I'm going to be talking to you about how the web privacy problem is a transparency problem, and show you the work that we're doing to improve that.

So when you're browsing the web and you visit a site like, let's say, The New York Times, you're not just visiting that first party site. But you're visiting all of the included third parties on that site. And this might be people you recognize like Facebook provide social buttons, or YouTube provides video. But what about the advertising companies or the analytics companies, and so on. They're not immediately obvious who they are to the consumer.

Well they could be anyone from this graph, right? Users might be able to figure out who they are if they use an extension like Ghostery. But what are their privacy practices? What other tracking practices? Which technologies do they use? That's not really obvious right, because the web lacks transparency. But what I'm going to show you today is how we're changing that. And I'll show you also how we already have.

So throughout this talk I'm going to reference back to our paper from 2014 called "The Web Never Forgets." It's a paper that looked at persistent tracking mechanisms. But in particular I'll focus on canvas fingerprinting.

If you're not familiar with canvas fingerprinting or that type of tracking mechanism, essentially instead of setting some state into the browser or setting cookies on the browser, you can look at the browsers properties and use that to uniquely identify someone across different websites, if you're a tracker. So in 2012 there was a paper called "Pixel Perfect," which talked about canvas fingerprinting. Sometime over the next two years AddThis, Ligatus, and a bunch of other companies, about 20 of them, started using this to track users.

In 2014 we went and did our measurement of this. We wanted to see who was doing it, where they were doing it, how the technology worked, and so on. And then shortly after releasing our paper we saw a bunch of news coverage and this really surprised us. We didn't expect such a response from the news and such a response from users-- things like ProPublica, BBC, and so on. And then just two days after all that news coverage happened, AddThis, who was the largest

provider-- they provided canvas fingerprinting on 95 percent of the sites. They stopped doing it, as well as Ligatus, which was the second largest provider.

So the thing to point out here is that canvas fingerprinting was a known technique for two years. But in just two months following our measurement work people stopped using it. So why was that? What was different about our work than just, say, having canvas fingerprinting be known and people knowing what it was?

And the key point is that our work removing information asymmetry between trackers and really, the rest of the web. So we got a bunch of news coverage from that, from different companies. And then we saw users take to Twitter to complain about it as you can imagine. We saw people say, hey, you should remove AddThis this from your site. This is a way of stalking. The first parties here are violating my privacy. We saw people just complaining about it and then we also saw someone say, I feel gross because I had to use AddThis to share this, but everyone should know about canvas fingerprinting.

So there was definitely a big response on Twitter and it wasn't just on Twitter. We also saw people for example, complain to Mozilla and saying, why doesn't Firefox protect me from this technique. And we even saw that it was beyond just users. It was also between trackers and the sites that they track on. So ProPublica focused on YouPorn which they wanted to point out that this tracking is happening there. And YouPorn responded to them and said, hey, we didn't know AddThis was doing this, could you let your readers know, I guess, that we've removed AddThis from our site.

So we see that transparency is effective at returning control to the users and the publishers of knowing what's going on. The users can see what kind of tracking technology is being used on their site and then they can make decisions, right? They can see what kind of tracking technologies are on the sites that they visit and then they can complain to the first party, to the site they're visiting. They can choose not to go there, right? They can have some control which they didn't have before when they didn't have that knowledge.

And automated large scale measurements like the one we did can help provide this transparency. So at Princeton I'm going to talk about a couple things that we did to make this happen. We developed OpenWPM. So this is really like the first infrastructure that can run a real browser across a large number of sites. And we're using it to run our own monthly million site measurements of this kind of thing. So we'll also build some analysis on top of that to look at who's fingerprinting on these sites. Who's tracking and so on.

So I'm going to walk through a little bit of how OpenWPM is built and how it works, and then I'll go into a case study of how it makes things a lot easier and show you how we can scale this up to all different kinds of technologies. So OpenWPM runs Firefox. And the way that we do it is we run it with something called Selenium. This basically lets us tell the browser, hey, you should go to this website, you should do certain things when you're on the website. And we run everything through a proxy so that lets us record all the traffic or all the communication between the browser and the sites that were visiting. And then we also have a Firefox extension based off of fourth party. So if you're not familiar with that, it's another web measurement framework.

Probably the most well used prior to us building our infrastructure. And we took all the features that had, added some more to it, and built it right into our platform as well.

So we give a researcher access to these different locations in the browser and then we wrapped that up in something called a browser instance. And as you can see here we're basically able to run multiple instances of Firefox, or multiple browser instances, at the same time. So when we do our own crawls. We run it over say 20 browsers, and each one has their own instrumentation. So you can easily scale this up to do measurement on a lot of sites.

And there's a couple things this lets us do. We can keep a profile consistent through crashes or freezes so we can keep the same cookies as we browse the different sites just like a real user would. We can also do things like run this with extensions or privacy features. See how well they work. See if they're actually protecting users or where they're falling short. And if there's any new web technologies being used for tracking like WebRTC or audio and so on, we can take a look at that.

So this is already used by seven research groups and you just heard a great presentation by the Web Privacy Census guy who do it. But it's also used beyond academia from journalists and regulators.

So I'll talk a little bit about the measurements we're doing. We're going on monthly crawls on a million sites and we're collecting things like, all the JavaScript crawls that might be used for fingerprinting, or all the JavaScript files on all of those sites. So we can go and check out what's actually going on later on. And we're also looking at the requests and responses in different storage locations in the browser. And this let's us do a bunch of things like see how effective privacy tools are, like Ghostery or AdBlock Plus, see how effective browser protections are, see how JavaScript might be used for tracking, and also look at tracking practices.

So now I'm going to give you two quick case studies. I'll go through canvas before we built OpenWPM and I'll go through WebRTC after. So canvas, like I said before, is just a site goes and draws text to the HTML5 canvas. And that text looks different on different machines but the same on the same machine. So it's useful if you want to differentiate between different users, but know who the same user is.

As you can see here, the differences can be quite large. This is just a visualization of differences between different machines compared to each other. And I want to give credit to all the co-authors on this study. I was just one part of it. So we work with people like KU Leuven and a bunch of other coauthors at Princeton.

So the way that this works is a website will draw a bunch of text to the canvas and make it overlapping, and try to maximize the chance that its unique. And that's what you see visualized up here. And if we want to measure this we have to do a few things. We first had the right Firefox patch to look for when these methods were called, when write text, or when pulling back the canvas as a string, when that happens. We had to write automation with selenium to go and run this across a bunch of sites, and build that from the ground up. And then of course we had to write some analysis code on top of that.

And now I'm going to show you how things were easier to measure another technique that could potentially be helpful for tracking. If you're not familiar with WebRTC and for using it for local IP discovery, essentially it adds some networking capabilities into the browser that you can access from JavaScript. And basically you're able to get the user's local IP if they're behind a NAT. If you're a home user that might be something like 192.168.1.2. But it can be useful for tracking. You can think of it like that. So I saw a tweet that this was happening and I said, oh, we can measure that. We can take a look at that. This won't be that hard. So I was able to add just a single line of JavaScript into our next crawl to do this.

So this is the same thing, I have a method here, that allows me to look at anytime anyone accesses WebRTC I can see that. I can see what they're setting and what they're doing with it. And this is the same method I used to look who's doing things with Canvas. Right? So it's just one added line of code to run our crawls. I had to write some analysis code on top of that. Very similar to canvas, right? With Canvas I wanted to know who wrote text and who read back from it. Well here I do similar things to see when this technique is being used.

And I found this happening on a bunch of sites beyond The New York Times, right? The New York Times actually stop doing it. So 121 first party sites and 24 of those were unique. Only one of which is blocked by AddBlock Plus or other similar privacy tools. So even if you're using privacy tools this technique may still be able to run on your machine.

I guess the point I want to make here is that web measurement gets much easier with OpenWPM. Instead of writing a Firefox patch we could just write a single line of JavaScript. Instead of writing automation with Selenium we could just use OpenWPM. And of course we still need to write the analysis code. You always need some extra human component in there. But the first two steps got a lot easier.

So where do we want to go with it? We think we can use this to inform the public, right? Let people know, hey, here's what's happening on the sites you are visiting. Here's who's doing canvas fingerprinting. And we think that'll really help people understand what's going on when they're browsing the web. We want to provide data for privacy tools. Disconnect, which is a privacy tool like Adblock Plus or like Ghostery. They actually ended up taking the scripts that we released as part of our canvas study and building into their own tool. So they protect against canvas fingerprinting. We want to provide that same kind of data for other privacy tools with our future studies.

And we also want to make the data accessible to less technical investigators, who may want to dig through it themselves, but don't maybe don't have all the skills necessary to dig through at the same level that save someone who writes the code would do. We'd also love to collaborate with people. So you the infrastructure is open source. You can go on GitHub and I'll have a link on the next slide to use it. You can download it and if you see anything wrong with it, or if you see new features you're welcome to submit back to it. We also envision people using it to run their own measurements like the Web Privacy Census. That's an awesome use case and we really hope that more people start doing that.

And then lastly, in the future, we hope that you'll be able to download our data and build some analysis of your own on top of it. We'll be going further with that in the coming months. So if you want to help us make the web more transparent, you can check out our GitHub repo to collaborate or you could check out our research page. Thank you.

[APPLAUSE]

JUSTIN BROOKMAN: Thank you very much, Steven. Now we're going to hear from Chris Hoofnagle with a critique of Alan Westin's Homo Economicus.

CHRIS JAY HOOFNAGLE: Good morning everyone. I wanted to start by thank the Federal Trade Commission and in particular its staff for putting together this event. The different investigators presenting today are very substantive and I'm proud to be among them. I think you've done a fantastic job and I think you should be proud.

My team at Berkeley over the years has shown different ways that websites and other web services track people. For instance my team published to be the first big paper about flash cookies explaining how flash cookies could be used to override users' cookie deletion. And we also showed how HTML5 paired with JavaScript could be used to do very similar things. And the theme of that work was a conflict between the kind of rhetoric one hears here in Washington about users being in control and users being able to make choices about how they're tracked online, and the technical reality. The technical reality that even mainstream companies could use flash and JavaScript to override deleted cookies. It was an attack that looked somewhat like a computer crime.

My presentation today is in a similar vein. It's about the conflict between theory and rhetoric, and how consumers actually operate in the marketplace. The FTC's notice and choice approach to consumer information privacy is based on the idea that consumers follow a rational choice model of making decisions online. Now the problem with notice and choice then becomes, is that the model of homo economicus, the model of the rational consumer who is making choices in the marketplace, has to be reliable as a model.

So much of my talk today is about the trade off talk. The idea that people are making trade offs in the marketplace on privacy. The theoretical background, of course, is about rational choice theory. And I'm going to skip over a bunch of slides to stay on time today. But the key point of my paper is that Alan Westin's theory was based on rational choice theory. And his main thesis was that public policy should serve the privacy pragmatists. And these are the people who weigh choices in the marketplace and make decisions according to their privacy preferences.

So we're familiar with these different definitions-- the privacy fundamentalists, the pragmatist, and the unconcerned. But let me draw your attention to some of the verbs Westin used to describe the privacy pragmatist. If you look at the verbs that are all highlighted in bold here. These are all active characteristics of consumers. The privacy pragmatists are people who weigh evidence. They are people who examine evidence. They look to see whether fair information practices are being widely observed. This is an active, engaged, consumer. I frankly don't know many people who are like this. I'm not even sure that I'm like this.

But this is the basis for much of US policy on consumer decision making and privacy. And of course Westin famously said, "In the politics of privacy the battle is for the hearts and minds of privacy pragmatists." These are the people you should be paying attention to and these are the people who policy should be designed for.

Well, how did Westin come to the segmentation of Americans? The way he did it was by asking this set of questions. One had to deal with consumer control, one had to do with whether data were treated confidentially, and finally the last question was kind of an attitudinal question about whether law is and self-regulation is sufficient for privacy.

So my first critique focuses on this segmentation text. On the most basic level, the problem with Westin is that he segmented it such so people were pragmatists by default. In this semantically doesn't make sense, because we're not pragmatists by default. Pragmatism requires affirmative action. It requires a certain outlook on life and I would argue the pragmatism is actually quite controversial. There's many Americans who find pragmatism quite distasteful but yet he coded it as the default result.

There are some other problems here. Westin's questions-- the screening questions used really had nothing to do with pragmatism. There's nothing in there asking do you read privacy policies? How much time you spent researching products and the like? It's just not in there. And then finally a significant number of consumers simply won't answer one of the questions. So I'll show you in our studies we found that somewhere between 2% and almost 5% of consumers wouldn't answer one of the three questions. What do you do with people who don't answer the questions. In Westin's methods you make them privacy pragmatists. That's really problematic. And it explains another critique we have in the paper that Westin never academically published his work. In part because I don't think it was publishable. This work, excuse me, this work I don't think was publishable.

So moving on. Another way to look at the data is empirically and this is where I'm standing on the shoulders of people such as Professor Turow. Turow pointed out years ago that when you ask people about the rules of privacy most of them don't get the basic answers right. He shows essentially that consumers think that the privacy policy is a seal. Most consumers think for instance that if a privacy policy is merely present, that website cannot sell personal information to third parties.

And it's for this reason that we should be very skeptical of trade off talk. People don't understand the trade to begin with. I'm going to get to a second reason why we should be skeptical of it. Turow of course, was standing on the shoulders of other people in the privacy field including Oscar Grande and his initial view of Westin's data. He viewed knowledge of privacy as a powerful explanatory factor of why people care about privacy and how they make decisions.

And so this is where a lot of my work picked up. And I wrote a number of studies. The fun covers started when I stopped writing with Joe. The Joe covers are boring, but my covers I think are more exciting. You'll see the Parthenon marbles on all my studies because I think they're quite beautiful. What we did, and starting at looking at Californians, what we did is we asked people about their privacy knowledge and we found a funny thing. The privacy fundamentalists

were always more correct than the other groups about existing law and business practices. And not only that people who shopped online where less knowledgeable of rules and practices than people who didn't shop online. Strange. Right? You'd think those people shopping online would read a privacy policy.

So we did a whole bunch of surveys over the years where we presented people with quizzes, asking them questions that Turow used, and that other investigators used. And we found over and over that the basics-- people failed the basic quizzes. And just as an example in our 2009 survey 75% answered two or fewer questions correctly. 30% got none of them correct. And then people say, well, the digital natives are going to save us. This is a generational problem. The digital natives are going to figure this out. No. They are actually the worst performers in the group. Both online and off, when we ask about offline privacy.

So we replicate the study again in 2012. And we find again that there's are substantial misconceptions about people's rights and about what practices are. And we find over and over again-- and the three stars mean a p-value of 0.001-- that the privacy fundamentalists are more knowledgeable than other groups. The other groups that are so-called who apparently don't care, or who are making trade offs.

So the main point of our paper is that Westin's segmentation has confused pragmatism with ordinary consumer decision making. And that many consumers in the marketplace are simply uninformed when they're viewing privacy policies. Another major part of this paper is the idea about whether Americans are more concerned about government collection of personal information or private sector personal information collection. And what we found over and over in our surveys is that Americans are concerned about both.

And this is not just our findings. If you look at the major literature reviews in public opinion quarterly and these are the political scientists who study privacy and they write these amazing literature reviews looking at all the different studies over decades. They find going back to the '80s Americans say they're just as concerned about the private sector as they are with the government sector. So we argue basically that RCT as a model fails in this field because people are laboring with substantial misconceptions about their rights and that they care about those rights.

Let me say something finally about Westin. Westin was a fantastic academic. And his work, his academic work, was great. And he is truly a progenitor of American information privacy. In his book Privacy and Freedom as you probably have heard, Omer's group has republished it and it's worth a read. Alan Westin was against technology determinism. Which is a philosophy one hears a lot of in DC. And he also saw privacy is a liberal value. So his survey work I critique today is not his academic work, and I have a lot of respect for that academic work.

So what do we do? What are the implications for FTC practice? Among them we could start viewing privacy policies as seals. When you go to the marketplace and you buy the organic vegetable, you don't look for in an organic policy. You look for, you assume, that organic mean certain things. We could start saying that privacy means certain things. Now the FTC has already

started to do this in security. If your privacy policy says anything about security it requires some type of reasonable control over our personal information.

Another approach comes from the history of the Federal Trade Commission. In the 1970s the Federal Trade Commission started recruiting marketing academics to come in house support BCP and this greatly punched up the Federal Trade Commission's understanding of how consumers were misled by false advertising. And if you look at today's commission actions their false advertising theories are much more in line with how consumers really understands adds and how consumers really act. And that has not come over to the privacy side. So we could replicate that.

And then finally I do think that we need to look at unfairness more as a remedy for privacy problems. Now why is this? Notice and choice might work in a world where you're selling physical products. But we are not doing that in this world. These are personal information products. And the transactions are not discrete, the transactions are continuous. That means that lock in, shifting practices, network effects, are all ways in which companies can shape choices and in effect remove choice from the consumer. And I write about this in much greater detail in this paper with Jan Whittington.

Finally let me just say thank you. And I can't avoid making a pitch for my book, which discusses these issues in much greater detail. And I do know that the ad practices division is not in attendance today. So what I'll say about it is if you read this book instead of eating chocolate, and other things, that you are guaranteed to lose weight without exercise.

[LAUGHTER]

Thank you.

[APPLAUSE]

JUSTIN BROOKMAN: Thanks Chris. And finally hear from Professor Joe Turow from the University of Pennsylvania on the trade off fallacy.

JOSEPH TUROW: Hello. Thank you. I'm going to go through this fairly quickly. It's a lot of stuff to talk about. I wanted to give you a sense of the arc.

The idea here, the summary, is that marketers justify their data collection with the notions that Americans want and understand the benefits of data trade offs. We challenged this assertion with results of a national telephone survey. Further we present that what observers interpret as trade off behavior is really widespread resignation among Americans regarding marketing use of their data. So that's the point. It's now what we sometimes interpret as trade offs and can be looked at when people do things, as gee, they're doing trade offs is really reflective of resignation of a large proportion of the population.

Okay. So what's the issue? Polls repeatedly find that consumers are concerned about ways marketers access and use their data online. And they're studies from Annenburg, from Pew, from

Bain & Company, reflecting that. At the same time observers agree that people often release data about themselves that suggest much less concerned about that. Okay? That's called among many people, the privacy paradox. The notion that people say they love privacy but in everyday life it's different. They don't. They give it up. They give up data for anything.

Some marketers read this paradox as evidence that people place other things above privacy, which leads to the notion of trade offs, that Chris was talking about. For example, Yahoo says that online Americans quote, "Demonstrate a willingness to share information as more consumers begin to recognize the value and the benefit of allowing advertisers to use data in the right way." And a president of Mobiquity says, "The average person is more than willing to share their information with companies if the organizations see the overall gain for end users as a goal not just for themselves."

It's just like some of the rational choice thinking that Chris was alluding to. A few corporate voices in papers, white papers, by Accenture, Bain, Brand Bond Loyalty, have put cautions around such generalizations. For example, Bain says customers trust can't be bought by companies offering compensation in exchange for selling or sharing personal data. And others have urged transparency, but really not saying what transparency means. They use the word but it's very difficult to figure out what they mean by it.

Generally though firms argue that consumers understanding of trade offs along with increasing consumer power justifies consumer data collection and use. The big deal today is that consumers have this huge power with the use of the mobile phone, the use of the internet, and other ways, and as a result companies have to push back sometimes in order to maintain some kind of profitable relationship. And marketers increasingly see personalisation resulting from predictive analytics as a savior in an age of hyper competition. So this is a great quote from Yahoo, "This concept of value exchange for personal data is starting to come to life through personalisation but it's a pathway to advertising Nirvana."

Now the trade off justifies 360 degree tracking. We can go into a whole lot of detail about the stuff. I just wanted to cite Gartner, a consulting firm. They talk about four stages through what they call cognizant computing, that will unroll over the next two to five years. It's written, I think, two years ago, with the first two well under way. They call them, sync me, see me, know me, be me. And it's the idea of really getting to know people as much as you can, data wise, in almost an organic way to figure out what's going on and how to make money off of them.

All right. But there are alternative explanations to trade offs. One is the public's lack of knowledge of what marketers are doing with their data behind the computer screen. Chris talked about some of that. A lot of surveys show that lack of knowledge. And Cranor McDonnell, Lorraine Alicia found that people really don't understand privacy policies. Alessandro Acquisti and others talk about the difficulty of understanding the technological and institutional systems.

Essentially this knowledge failure research explains the ease with which data retailers and advertisers retrieve information from individuals, through the proposition hasn't been directly tested. But it might get marketers off the hook too easily. So we say gee people have lack of knowledge it's because the schools don't teach it enough. Or let's figure out an educational

program. Ad choices, those little icons that you're supposed to see. I gave a talk at the Penn law school one day showing a slide and nobody saw it. But they could point to this. And to sound more optimistic about what the public is than people like me or policymakers about this.

So we did a survey to try to look at some hypotheses related to this. A 20 minute, on average, interview taking place in February and March 2015. English speaking or Spanish speaking. 750 land line. Wireless 756. Conducted by Princeton Survey Research Associates. More data about that is in the paper. We look first at people's philosophy of trade offs. Not the particulars but what do they know about, what do they think about the idea of a trade off. And you can see it says, "If companies give me a discount it's a fair exchange for them to collect information about me without my knowing it." 91% said no. It's fair for a physical store to monitor what I'm doing online when I'm there in exchange for letting me use the store's wireless internet and Wi-Fi without choice. 71% said no. It's okay if the store where I shop uses information it has about me to create a picture of me that improves the services they provide about me. 55% said no.

Now oddly if we look at how many people agree with all three propositions, only 4% agreed with all three propositions. We took a broader idea of what agreement was when we gave you numbers to each, like agree strongly, agree, disagree, disagree strongly. And in that broader interpretation of belief in trade offs, we found the still small proportion, 21% believes that common trade offs with marketers amount to a fair deal.

But we wanted to look at the privacy policy in terms of a scenario of real life. So we said, for the next few questions please think about the supermarket you go to most often. Let's say the supermarket says it will give you discounts in exchange for its collecting information about all your grocery purchases. Would you accept the offer or not? 52% said no. 43% said yes. Which is interesting because it's close to the three statements we said, it's okay if a store where I shop uses information it has about me to create a picture. You say, well, that those 43%. Turn out it's not. Because when we looked at it, we found that only 40% of the people who accept that dictum, agreed with the supermarket thing. Notice people are very inconsistent. The lack of correspondence, even when the scenarios appear similar, underscores that a small percentage consistently accepts the idea of trade off.

We wanted to know whether people who say they will accept the supermarket discount will still do it when presented with specific assumptions a supermarket may make. So for example, you might say I'll take the discount, but what if you know what the supermarkets doing with your data. This is knowledge Americans almost never receive directly but may intuit from ads and coupons they think they're targeted toward them.

So we have a variety of things we ask them. The people who said they would accept a discount in the first place, we said, would you accept it the supermarket makes assumptions based on your purchases about whether you buy low fat foods. It went down to 33%. The more we asked particular questions about individuals' lives, the less they said they would do it. So in the end when we asked about social, ethnic background inferences, only 19% said they would accept it.

The table shows the limits of cost benefit analyses as a rationale for marketers' claims that most people will provide personal data in exchange for store deals. The decline in acceptance from

43% to around 20% is inconsistent with marketers' assertions that people are giving up their personal information because of cost benefit analysis. In the supermarket scenario, they're doing just the opposite. Resisting idea of getting data for discounts based on some kind of analysis.

Then we went ahead, our hypothesis about resignation came out of every day realization, when we met people, they would say things like, gee, you know, I have to give up the data. I want to be online, I have to be on Facebook. I know they do this stuff. I don't know. I don't know what's going on. But I have to do it anyway.

So we gave the people two statement separated by many other statements so they weren't right next to each other. I want to have control over what marketers can learn about me. I've come to accept that I have little control over what marketers can learn about me. Okay? It turns out that 58% of people agree with both things, which we say indicates a sense of resignation. Resignation, meaning the acceptance of something undesirable but inevitable. Got that from Google dictionary.

We've found there is a strong positive statistical relationship between believing in trade offs and accepting or rejecting various kinds of supermarket's use of discounts. You'd expect that. By contrast, there's no statistical relationship between being resigned to marketers use of data and accepting or rejecting the supermarket trade off. People who are resigned, sometime they do, sometime they don't. They're trying to navigate a world that they don't understand, are annoyed about possibly, and they sometimes will do it. They may look like they're accepting trade offs, but in their head they're saying, gee, I'm resigned to it.

Put another way people who believe in trade offs give up their data predictably, while people who resigned don't do it in a predictable manner. They do give up their data though. We found 57% of those who took the supermarket deal were resigned. A much smaller 32% were trade off supporters, even using the broader measure of trade off support that I suggested. The larger percentage of people in the population who are resigned compared to those who believe in trade offs, indicate that in the real world people who exchange their data for benefits are more likely to do it while resigned, rather than as a result of cost benefit analysis. Moreover we found the resignation is widespread across the US population, regardless of age, gender, education, or race. There were no statistical differences between age and gender. There were between education and race. But still the large percentage of people resign . Anyway

We found that most Americans don't have basic knowledge to make informed cost benefit choices. This is some of the stuff that Chris was talking about. 51% can't recognize the possibility of phishing. Large percentages believe incorrectly that government and laws protect them from price discrimination and certain forms of data collection. When they don't. These widespread misconception suggest they even when Americans do weigh the cost and benefits of giving up their data, they frequently base those choices on incorrect information.

But we also found, and this really was surprising to me, that those who know more about marketing laws and practices are more likely to be resigned. We found too, that resigned people who accept supermarket discounts even as the supermarket collects increasingly personal data,

have more knowledge than others. So having more knowledge is not protective, as some academics have suggested.

So what do we do about it. The rationale of trade offs is a fig leaf we argue, used by marketers to justify a world of tracking and increasingly personalized profiling that people know is there, don't understand, and say they don't want. We haven't begun to consider the social implications of having a large population that is a resigned about a key aspect of it every day environment. Now this may sound really dark and what do you do about it. But I think it's really important to confront what I see in everyday life when I talk to people. That people do these things online in stores with apps, not because they're thinking in a cost benefit way rationally. But because they feel that they have no other choice if they want to live in this world.

We're only at the beginning of key aspects of this era. This is the beginning of the new era, not even in the middle. And there may be time for concerned parties to guide it. Academics, journalists, and advocates have to translate the key issues for the public. And there are a lot of issues of obfuscation and deception we could talk about. Issues at the FCC might be involved in around public interest, convenience, and necessity. The importance that people alluded to, to praising and naming groups that do the right things and not to right things. Thanks for listening.

[APPLAUSE]

JUSTIN BROOKMAN: Thank you Joe. Thanks to all of our presenters. And now we're going to move into a brief period of discussion, with one caveat. Joe may have to leave early. He's teaching two classes later today at Penn. So if you see him slink off he's not in trouble, we're not angry with him, he's not mad at us,

So I'm going to start with some of the trends I saw some of the presentations. One is the proliferation and growing sophistication and growing complexity of online tracking was reflected in Ibrahim's and Steven's work. There's more cookies. There's more companies who are doing it. I love the revised lumiscape chart with all the hundreds and thousands of companies that can even see them on the big screen. And more technologies too. It's not just cookies, it's HTML5, it's fingerprinting, it's e-tags, it audio beacons, it's who knows? And then logically, perhaps unsurprisingly, then the theory of Joe and Chris's argument is that there's an increasing inability of consumers to really manage or control their privacy. Given all these advances.

So the idea that a consumer goes to a website and reviews the privacy policy, and makes informed choice and I am satisfied with how e-tags are used on the site, and I will now access my content in exchange for that is perhaps flawed. And this builds somewhat on Laurie Cranor's work that if you had read every single privacy policy it would take months of your life. And so, instead of that, it's sounds like there's this resignation. Instead of privacy pragmatism, there is resignation. This is what Joe's work was talking about this.

This hit home with me this weekend. I went skiing with a friend of mine and we were talking and he said, he sent a link to his dad to a news story. And his dad called him and said, I'm not opening that. You know how many cookies are in there? And he's like, yeah, I know. He's not a privacy guy. Neither of them are. He's like, yeah, I know, what are you going to do. You could

have tried to walk through deleting cookies or installing AdBlock, but he had to go pick up his kids, right? We don't have the time to really think of every single one of these questions. We've all talked people who've had similar experiences. We've all probably had similar experiences ourselves. Like, I don't quite know what's going on here, but, I don't have the time to figure it out.

It's not just the web, right? It's the internet of things. We had our cross device tracking workshop where TVs or toasters can collect information about us. It's physical space with the automatic license plate readers. And are we making informed choices when go outside about facial recognition. And so one thing I would like to hear from the folks about and I'm going to turn it over to my other discussants first. So what does the solution look like, right? I mean do we just ride it out. There are a lot of folks who say that. Brandeis was concerned about cameras and we're cool cameras now. Do we want government making rules about how much tracking can happen if consumers can't make the choices themselves. Say 15 cookies and that's it. The part of PrivacyCon is that we can hear from really smart people who are thinking about this to help them influence policy decisions. And so I'd love to hear some thoughts about solutions later. I'll ask a question about that but first I'm going to turn it over to Elana.

ELANA ZEIDE: So one interesting theme I'm noticing is a shift away from the idea about informed notice for individuals but more transparency for the collective populace. Including both consumers and more importantly, sort of experts, , advocates policymakers, academics. And with that seems to also be a shift from the idea of questioning consumer's decision making capabilities to whether that they're actually engaging in a choice at all, or either resigned because they see no agency and no reasonable alternatives to opting out of the mainstream, or because they have trust in the default system.

If you look at the idea of what whether transparency is a means to solve those issues and accomplish that, I think there are several implications based on this research. One is how do you use transparency as a way to galvanize consumers to articulate their preferences, or to engage in privacy self-management, if, in fact, it may lead to them being more resigned because they have a feeling of helplessness. Also, how do you insure, or predict, when companies will actually be prompted by public opinion to make a change and whether those changes will actually occur without regulation or other enforcement mechanisms for the most meaningful potential privacy abuses. Which might also be most likely to be the most profit generating core of many companies businesses.

There's also the question of whether transparency can operate as a mechanism to insure consumer trust in a world where there are unknowable unknowns. Years ago, people would allow their friends to post pictures on Facebook without thinking that their picture would remain in obscurity because they weren't being tagged. In an age of facial recognition that is no longer true. I think these shifts really undermine consumers' sense of what they can predict and how they're choices-- a sense of helplessness, in the sense of the unknown and what may happen in the future.

Finally I'm interested in the idea of whether a move towards transparency or shaming and blaming creates a system where we may be able to get some clarity about consumer norms and

what standards infer. It may also create a situation where sensationalist media stories or small vocal subsets who resist certain practices, end up controlling the conversation and give a false sense of clear consensus. And the last point would be, does this that entail a system where we must wait for harms and abuses to occur before we can then create systems to correct them, and if so, does that imply that along with some transparency mechanics, we also need mechanisms that consumers can see for due process and redress.

OMER TENE: Thank you. So I think all four presentations here drew sort of grim and somber picture of the state of play today with consumers being misled or resigned, and kind of being dragged along for the ride by technology or by business. Given that the stars seemed aligned on this. I feel an urge to play devil's advocate. And in that role I'm going to suggest a couple of different adjectives to describe how consumers are acting, or feeling, or faring. Instead of being resigned I'll suggest that actually thrilled, or maybe even exhilarated or delirious about these new technologies. About the fact that they can hail an Uber and rate the driver. And get the newest iPhone or Android phone and even, yippee, take a selfie, and post it under a SnapChat story. Or use a FitBit and give up their fitness or health information.

And I think we clearly see that in the marketplace. We also see Google and Facebook and Apple, Microsoft, three or four of the strongest brands in terms of brand recognition in the market. And not to mention the number of people flocking to work at these places including people who are now in government and even in regulatory agencies. So the point is that there seems to be something more complex at play here. And you know, I think we see it in another contexts. So I care about health, but I still eat a cheeseburger. I care about the environment, but you know, I drive a four wheel drive. There's a lot of snow in New England. And I think part of your response, your retort, will be yes, but consumers are ignorant. They just don't know. But actually I think Joe's survey and research shows that the more informed, they become more resigned. So maybe it's better to just be blissfully ignorant. So with all that I want to turn back to you and hear your reaction.

JUSTIN BROOKMAN: Do you want to start?

JOSEPH TUROW: I mean these are really important insights. I think that it's a complicated world. It's very hard not to be excited about the ability to walk through a store and compare prices in your hand. There are levels of excitement about being able to show a kid a snippet from the Wizard of Oz on a phone, on a bus, when the kid's starting to get antsy. I mean, there are lots of things that are terrific about this. I couldn't live without Google. But where I'm coming from anyway, is that, I think part of my job, is to say-- I mean there are a lot of companies who are saying all these great things. But underlying it there are some real problems that we have to face. And I think part of being a citizen in a society is to say, yeah, there are terrific things about this, but there are also things that in the long term might, and I really do believe this, might harm our democracy. Might harm our relationship with others.

When you walk through a store now and you're not sure what profile the store has about you. When not too long from now you can get on your phone, in some places this already exists, different prices based upon who you are. That's a scary thing to me in terms of how are people going to understand the public sphere. Their relations to others. How are people going to

understand the political process when they think they're getting information that is developed personally for them, that are personal ads. And so while I agree that there are many terrific things about this. I think that there have to be segments of society, they have to say, stop, we can fix the really difficult things that relate.

CHRIS JAY HOOFNAGLE: Let me unravel some of the issues. What I'd say is that at first but one can look at our work and say it's anti-technology. But I would argue strongly that it is not. I personally love technology and I'm an early adopter of many, many things. I'm also a practitioner. And I do know that much what we call innovation does not depend on personal information. And is fundamentally compatible with what Alan Westin would call modern information privacy law, such as, we're going to de-identify this information after six months. We're going to delete it after a year, et cetera.

So I think one of the rhetorical, it's in a way straw man, that we have to recognize and deal with, is the idea that we can't have privacy and these technologies. We can have Uber. Uber is actually not that innovative. Long before Uber taxicab companies had hail apps and blah blah blah. Don't need personal information for a lot of that. But where you do need personal information you can have rules around it. And I see it from practice all the time. There are situations where we do very interesting forms a personalisation with de-identified data, where we agree that data will disappear after a certain amount of time. Where we agree that certain things won't be the basis of selection and the like. So I think we shouldn't fall under the false dilemma that privacy means we cannot have spectacular convenience in our life.

STEVE ENGELHARDT: So coming at this from I guess from the tracking perspective I wanted to comment on the fear of maybe users becoming resigned by getting more information about what tracking was going on, or the notion that we can't have the services without having the tracking. Because I think there is definitely a chance that if everyone starts fingerprinting and users just see, oh, every site I'm visiting is fingerprinting me, I guess I just have to deal with it. Like that could happen. But I think we could prevent that from happening with the right policies and with the right tools where consumers could protect themselves by releasing that data for not just to consumers but for everyone.

And then the notion that consumers just have to be tracked. I don't think that's really true either. Because a lot of, at least for advertisers, they support opt outs right? You should be able to set an opt out cookie and not be tracked. But we still see that fingerprinting often goes on when those opt out cookies are set. So perhaps there should be some enforcement that if you're going to tell users you've opted out of tracking, you can also guarantee you won't be things like fingerprinting. And the user doesn't just have to trust that that won't happen.

JUSTIN BROOKMAN: I'll ask one more question. The broader policy solution is what's the alternative because I talk to a lot companies and they kind of tell the same story right? Of course consumers can control this stuff, but they argue that really means there should more of an accountability model right? Companies should be responsible stewards of the data right? Consumers can't make control. Companies should make smart, informed decisions about how the information is used? Because what is the alternative to that? That's one option. And then there's

like the FTC or the government could be making prescriptive, paternalistic choices on behalf people. That has its problems as well.

One thread we've heard a few times today is the idea of increased transparency, and then filtered through elites or institutions. And then the name and shame approach that Joe and Steven talked about and Elana talked about in her comments. I guess my question is, is that scalable right? I mean The Wall Street Journal did their "What They Know" series starting in 2010. And yet the reports you guys show is that the tracking that they're concerned about is still increasing. Joe and Kristin have been doing this for even longer. So what is the policy solution? Assuming that this is a problem to be addressed. What is the right approach?

OMER TENE: Can I jump in and say that?

JUSTIN BROOKMAN: Yes.

OMER TENE: Thank you. That I think. And also reacting to what Chris and Joe said, I think there is consensus that we need to deal with data excess, and have the identification and strong data security. But I think to a large extent the industry gets it and certainly industry gets the big impact that privacy fails can have on brand and consumer expectations. And I think one thing that attests to this is the fact that we're having this conference and the existence of a privacy profession that has blossomed so the IPP now has a 25,000 members worldwide, that had less than 10,000 just two and a half years ago. I think the right processes are in place and it's really the excess that we need to deal with. And I think you illustrated this some of this in technological research.

CHRIS JAY HOOFNAGLE: The access and accountability issues have to be dealt with and there are very interesting proposals to focus mainly on use of data. But I think one weakness of those proposals as they don't take into account the attacks on accountability that are occurring, such as the Spokeo case. If you take the Spokeo case seriously and if you read the Amici briefs, a large portion of the technology industry is arguing that they should be able to willfully violate the law. Willfully. That means they know what the law is and they violate it anyway. And that they should be able to be sued. Wyndham was in a way an attack on accountability.

The class action. We don't like class actions. We don't like the FTC doing anything. We don't want congress to do anything. So where exactly does the accountability come from. And I think when you look at use models, the first defense, the first time someone gets caught in a use violation, they're going to make an IMS Health argument. And so I think if we're going to move toward a use model the accountability is going to have to include a contractual waiver of first amendment defenses, and an agreement that there is injury in fact that supports standing. Otherwise you'll never be able to sue, not even you Justin. If you take the position seriously not even the FTC would be able to sue.

OMER TENE: I think some companies have staked radical positions and frankly I think done themselves a disservice. Which is something that I think is prone to occur in litigation. On the whole, the FTC has been successful, and I'm not sure how much traction the First Amendment argument against privacy accountability will have. We'll see.

ELANA ZEIDE: So one question I had following up on that is, so when you talk about use and the assumption of harm, are you looking at-- it seems that use that in this case-- a broader word to really talk about data driven decision making. Is, that I think, where you see the troubles lie?

CHRIS JAY HOOFNAGLE: I'd like to defer to someone else because it's not my area of expertise.

JUSTIN BROOKMAN: So could you repeat that?

ELANA ZEIDE: Sorry. I was just saying. In this case when we're talking about what abuses are and what sort of harm, is it really about the uses is in terms of the tracking and what people are theoretically doing with information, or abstractly, or does it really become an issue when there's data driven decision making?

STEVE ENGELHARDT: So I guess the fear. I would say it's more of the data use. It's the fear that if this data is being collected, how is it being used? And that the consumer has no ability to go and prevent that collection or no ability to control that collection beyond preventing it from happening. So once the data gets put into the company's databases that kind of this up to trust.

JUSTIN BROOKMAN: And so I guess that kind of goes to the point. So should we be concerned about the collection itself. You guys both had your studies. There's a lot more collection going on. I'm sure there are lot of people room are like, yeah, but it's not bad collection. It's not malicious collection. It's being done to support of the ad ecosystem. Which there's nothing inherently wrong with that. And then there's definitely been folks who said the FTC should be focused on the cases where there is the harm down the road right. Commissioner Ohlhausen has written about this. There's a lot of focus on the use of data for discrimination right. We have a panel and on that later today. So should we be focused-- should FTC policy in general be concerned about the raw collection in the first place or is it just the fact that we should be worried about how it could be abused down the road?

CHRIS JAY HOOFNAGLE: I've written pretty extensively about the need to focus on collection because of the inability to police uses. And I think to get to a point where we come police use we need to have sea change and a form of accountability that doesn't currently exist. What my team is found over and over when we discover things like HTML5 or flash cookie re spawning. We go to the companies and we say, we think you're doing this and they say, no, we're not doing it. And they actually don't know what they're doing. Any other closing thoughts. With that we are over time. Thank you all so much. We're going to have a quick 10 minute break and then we'll come back for the second session at 10:45.

[APPLAUSE]