In the Matter of:

PrivacyCon Workshop

January 14, 2016
Final Version

Condensed Transcript with Word Index

For The Record, Inc.
(301) 870-8025 - www.ftrinc.net - (800) 921-5555

UNITED STATES OF AMERICA
FEDERAL TRADE COMMISSION

PRIVACYCON WORKSHOP

THURSDAY, JANUARY 14, 2016

FEDERAL TRADE COMMISSION
Constitution Center
400 Seventh Street, S.W.
Washington, DC

I N D E X

P R O C E E D I N G S

\- - - - -

MS. YEUNG: Good morning and welcome to PrivacyCon. I am Tina Yeung, a paralegal in the FTC's Office of Technology, Research, and Investigation, or OTech. Before we commence, I have some brief housekeeping details to 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 that 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, boxed lunches will be available for purchase in the hallway outside 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 conference materials will be available via FTC.gov after the conference ends.

And, finally, we are live-tweeting today's event under #PrivacyCon. Thank you, and over to Dan Salsburg.

MR. SALSBURG: Thank you, Tina. I'm Dan Salsburg, 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

5

had preregistered for this event. We're sorry for sharing that information with you last week, and are addressing our bulk distribution setup 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?

(Applause.)

CHAIRWOMAN 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

6

smart appliances to connected medical devices to semi-autonomous cars. All of these mean the consumers must navigate an increasingly complex and dynamic 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 are 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 need 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 consumers face, 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 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

7

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 would like to speak very briefly about the way that 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

8

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 the 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 their use of big data analytics produces benefits for consumers while avoiding outcomes that may be exclusionary or discriminatory.

In this report, we highlight possible risks that could result from inaccuracies or biases about certain groups and data sets, including the risk that certain consumers, especially low-income or underserved consumers, might mistakenly be denied opportunities or that big data analytics might reinforce existing socioeconomic disparities.

On the enforcement front, the work of 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 have also asked technologists and researchers to help us come up with technological countermeasures to address vexing problems. Illegal robocalls are a key example. Voice Over IP technology allows callers to spoof identifying

information such as the calling party's phone number. Fraudsters can now place millions of cheap, automated calls 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 are also continuing to build our internal capacity. Last year, we created the Office of Technology, Research, and Investigation, or OTech, as we call it. OTech, which builds on the work of our former mobile technology unit, identifies and conducts research that 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 OTech, 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 companies' 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 that 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 cautious 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 and 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 to 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 are 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 craft 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 would also like to thank the organizers in OTech and our privacy division, DPIP, and in particular, Kristen Anderson and Dan Salsburg for their hard work in putting this event together. So, thank you very much.

(Applause.)

SESSION 1

THE CURRENT STATE OF PRIVACY

MR. 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 am a 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. We got nearly 90 really fascinating proposals, so originally we were going to just try to do 12 or so, which was tight. But we tried tohe pack the schedule so we could have at least 19 people presenting, and we honestly wish we could have done more. So, we have tried to maximize the schedule to let them present their research to you. They're each going to present for about 15 minutes. We are going to try to keep them aggressively to that. They have a clock right there where it shows they're over time. They have a chime that plays they're over time. So, they will know

they're over, you will know they're over, they will know that you know they're over, so we will try to stay on schedule.

After that there will be a short discussion period. My co-discussants are Omer Tene from IAPP and Elana Zeide from NYU, we will give a few thoughts, ask a few questions and then that will be it. This is our first time doing this, we would love your feedback, if we're able to do this in the future and you apparently have a lot of interest and 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.

MR. ALTAWEEL: Hello, everyone. My name is Ibrahim Altaweel, I am the co-author of 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 showed relatively little tracking online in 1997. Only 23 of

the most popular websites used cookies on their home pages. But within a few years, tracking for commercial 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 that study in 2012 and in 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 questions: 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 the Quantcast top 100, top 1,000 and top 25,000 websites using OpenWPA, a web privacy measurement platform developed by Princeton University.

We ran a shallow crawl and deep crawl. A shallow crawl means that we only visited the home page

17

of the popular websites, and a deep 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 to have 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 being set. Overall, these limitations mean that web privacy census is a conservative measure of the amount of tracking online.

So, how much tracking is going on? We found that users who merely visited the home pages of the top 100 most popular websites would collect over 6,000 HTTP cookies, twice as many as detected in 2012. Some popular websites use a lot of cookies. In just visiting the home pages of popular websites, we found that 24 websites placed over 100 cookies, six websites that placed over 200 cookies, and three websites placed over 300 cookies.

What technologies are most popular for tracking users? One obvious observation is that there are significantly more HTML5 local storage objects than Flash cookies. HTML5 local storage is a new technology that became popular in recent years for its large

18

storage capabilities, roughly a thousand 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 browser needs it 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? We show the percent from Flash cookies to HTML5 local storage, it is very interesting to see that the total count of cookies has increased, and there are more and more third party cookies being used. 83 percent of HTTP cookies are set by third-party hosts, and in just visiting the home pages of popular websites, users would have cookies collected by -- I mean cookies placed by 275 third-party hosts. If the user browsed to just two more links, the number of HTTP cookies would double.

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

So, what entities have the greatest potential

19

for online tracking and why? The most prominent one is Google. We found that Google's tracking infrastructure is number 92 of the top 100 most popular websites. And on 923 of the top 1,000 websites, providing Google with a significant surveillance infrastructure online.

Google's ability of tracking is unparalleled. Most of 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 is a consistent upward trend in cookie usage and that a small group of companies have been tracking cookies almost everywhere on the web.

In the future, we will continue to collect and analyze key metrics and measures to monitor the state of online privacy. Thank you very much. And I also would like to thank my co-author, Nathan Good. Thank you.

(Applause.)

20

MR. BROOKMAN: Now we're going to hear from Steven Englehardt of Princeton University on The Web Never Forgets.

MR. ENGLEHARDT: Hello, everyone. I'm Steven Englehardt 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, right, but you're visiting all of the included third parties on that site. And this might be people you recognize, like Facebook provides social buttons or YouTube provides video, but what about the advertising companies and the analytics companies and so on that are not immediately obvious who they are to the consumer?

Well, they could be, you know, anyone from this graph, right? It could be, you know, users might be able to figure out who they are if they use an extension like Ghostery, but what are their privacy practices, what are their 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 talk about and 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 instead of setting cookies on the browser, you can look at the browser's 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. Some time over the next two years, AddThis, Ligatus and a bunch of other companies, about 20 of them, starting using this to track users.

In 2014, we went and did our own 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 of that news coverage happened, AddThis, who was the largest provider, they provided canvas fingerprinting on 95 percent of their sites, they ended up -- 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? You know, what was different about our work than just, say, having canvas fingerprinting being known and having people know what it was. And the key point is that our work removed the information asymmetry between trackers and really the rest of the web.

So, like I said, 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 from your site, this is a way of stalking, this is -- you know, the first parties here are violating my privacy.

We saw people -- we saw people just

complaining about it, and then we also saw someone say, you know, I feel gross because I had to use AddThis to show 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, complaining 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, of course, 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 -- you know, 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 complain -- or you can see what kind of tracking technologies are on the site that they visit, and then they can complain to the first party or to the site that 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 of things that we did to make this happen. We developed OpenWPM. So, this is really like the first infrastructure that can run on a bunch of sites or 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 will 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 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 let's 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 and all

25

the communication between the browser and the sites that we're 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 that had, added some more to it and built it right into our platform as well.

So, we give our researcher access to these different locations in the browser and then we wrap 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 of things this lets us do, right? We can keep a profile consistent through crashes or freezes, so we can keep the same cookies as we browse through 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.

26

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 guys 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 of a million sites and we're collecting things like all the JavaScript calls 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, you know, the requests and responses, and different storage locations in the browser.

And this lets 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, I'll go through WebRTC after. So, canvas, like

27

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 keep -- you know, know who the same user is.

As you can see here, the differences can be quite large. This is just a visualization of the differences between different machines compared to each other.

And I want to give credit to all the co-authors of this study. I was just one part of it. So, we worked with people at KU Leuven and a bunch of other co-authors at Princeton.

So, the way that this works is a website will draw a bunch of canvas -- a bunch of text to the canvas and make it overlapping and try to maximize the chance that it's 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 to write a Firefox patch to look for when these methods were called, when right text or when pulling back the canvas has a string when that happens.

We had to write -- sorry, 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

28

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, it's a -- or WebRTC 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 in that. 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, you know, 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 -- you know, I have a method here that allows you to look at any time anyone accesses WebRTC and I can see that, right? I can see what they're setting and what they're doing with it.

And it's the same method I use to look at 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 want 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 stopped doing it. So, 121 first-party sites, and 24 of those were unique, only one of which is blocked by, say, Adblock Plus or other similar privacy tools. So, if you're going to be using privacy tools, this technique may still be able to run on your machine.

And 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. And 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 -- 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're visiting. Here's who's doing canvas fingerprinting, and

we think that will 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 it into their own tool, so they protected 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 maybe don't have all the skills necessary to dig through it at the same level that, say, someone who writes the code would do.

And we would also love to collaborate with people. So, you can -- you know, the infrastructure is open source, you can go and 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 needed features, you are 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 will be able to download our data and build some analysis of your own on top of it. We will be, you know, 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 can check out our research page.

Thank you.

(Applause.)

MR. BROOKMAN: Thank you very much, Steven.

Now we're going to hear from Chris Hoofnagle with a critique of Alan Westin's Homo Economicus.

MR. HOOFNAGLE: Good morning, everyone. I wanted to start by thanking the Federal Trade Commission and, in particular, its staff for putting together this event. The different researchers presenting today are very substantive and I am proud to be among them. I think you have done a fantastic job, and thank you, 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 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 are 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 and 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 that the model of a 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 tradeoff talk. The idea that people are making tradeoffs in the marketplace on privacy.

The theoretical background, of course, is about rational choice theory, and I am 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 in rational choice theory, and his main thesis was that public policy should serve the privacy pragmatists, so 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 pragmatists and the unconcerned, but let me draw your attention to some of the verbs Westin used to describe the privacy pragmatists. If you look at the verbs, they're 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 U.S. policy on consumer decisionmaking and privacy. And, of course, Westin famously said, in the politics of privacy, the battle is for the hearts and minds of the privacy pragmatists. These are the people we 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 is kind of an attitudinal question about whether law 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 -- he segmented it such so people were pragmatists by default, and 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 that pragmatism is actually quite controversial. There are 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, you know, do you read privacy policies, how much time do you spend 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 two and almost five percent 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 tradeoff talk. People don't understand the tradeoff to begin with, and I'm going to

get to a second reason of why we should be skeptical of it.

Turow, of course, was standing on the shoulders of other people in the privacy field, including Oscar Gandy, in his initial view of Westin's data, he viewed knowledge of privacy as a powerful explanatory factor of why people care about privacy in how they make decisions.

So, this is where a lot of my work picked up, and I wrote a number of studies with fun -- well, actually, 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 in all my studies because I think they're quite beautiful.

You know, what we did, and starting in 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 traditional practices. And not only that, people who shopped online were less knowledgeable of rules and practices than people who didn't shop online. Strange. Right? You would think those people shopping online would read the 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 on the basic quizzes. Just as an example, in our 2009 survey, 75 percent answered two or fewer questions correctly, 30 percent got none of them correctly, 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 asked about offline privacy.

So, we replicate the study again in 2012, and we find, again, that there's a -- that there 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 tradeoffs.

So, the main point of our paper is that Westin's segmentation has confused pragmatism with ordinary consumer decisionmaking, and that most -- many consumers in the marketplace are something uninformed.

They're viewing privacy policies essentially.

Another major part of this paper is the idea about whether people -- 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 -- you know, these are the political scientists who study privacy and they write these amazing literature reviews looking at all of the different studies over decades. They find, going back to the 1980s, 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 they do care about those rights.

Let me say something, finally, about Westin. Westin was a fantastic academic, and his work, his academic work, was great. He is truly a generator of American information on 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 D.C., and he also saw privacy as 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 can start viewing privacy policies as seals. When you go to the marketplace and you buy the organic vegetable, you don't look for an organic policy. You look for -- you assume that organic means 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 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 at the PCP, 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 understand ads, 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 that is, if you read this book instead of eating chocolate and other things, you are guaranteed to lose weight, without exercise.

(Laughter.)

(Applause.)

MR. BROOKMAN: Thanks, Chris. And, finally, we are going to hear from Professor Joe Turow from the University of Pennsylvania on The Tradeoff Fallacy.

MR. TUROW: Hello. Thank you. I am going to go through this fairly quickly. It's a lot of stuff to talk about, but I wanted you to get a sense of the arc. The idea here, a summary, is that marketers justify their data collection with the notions that Americans want and understand the benefits of data tradeoffs. We challenge this assertion with the results of a national telephone survey. Further, we present evidence that what observers interpret as tradeoff behavior is really widespread resignation among Americans regarding marketers' use of their data. So, that's the point. It's not what we sometimes interpret as tradeoffs and can be looked at when people do things as, gee, they're doing tradeoffs, 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 there are studies from Annenberg, from Pew, from Bain & Company, reflecting that. At the same time, observers agree that people often release data about themselves that suggests much less concern about that, okay? That's called by

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 tradeoffs that Chris was talking about. For example, Yahoo says that online Americans "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 the president of Mobiquity says, "The average person is more than willing to share their information with companies if these organizations see the overall gain for end-users as a goal, not just for themselves." This reflects 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 generalization. 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 tradeoffs 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 personalization 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 personalization, that it's a pathway to advertising nirvana."

Now, the tradeoff justifies 360-degree tracking. We can go into a whole lot of detail about this 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 was 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 tradeoffs. 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 and McDonald, Lorrie and Aleecia, 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. So, the proposition hasn't been directly tested, but it might get marketers off the hook too easily, so we say, gee, people have a lack of knowledge. It's because the schools don't teach them enough. Or let's figure out an educational program.

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

45

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/March 2015, English-speaking or Spanish-speaking, 750 landline, wireless 756, conducted by Princeton Survey Research Associates. More data about that is in the paper.

We looked first at people's philosophy of tradeoffs, not the particulars, but what do they know about, what do they think about the idea of tradeoffs? And you can see it says, if companies give you a discount, it's a fair exchange for them to collect information about me without my knowing it; 91 percent said no. Is it fair for an online or physical store to monitor what I'm doing online when I'm there in exchange for letting me use the store's wireless Internet or WiFi without charge; 71 percent said no. Is it okay if a 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 percent said no.

Now, oddly, if we look at all of how many people agreed with all three propositions, only four percent agreed with all three propositions. We took a broader idea of what agreement was when we gave numbers to each, like agree strongly, agree, disagree, disagree strongly, and in that broader interpretation of belief

46

in tradeoffs, we found there's still a small proportion, 21 percent believes that common tradeoffs 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 this 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? Fifty-two percent said no, 43 percent said yes, which is interesting, because it's closer to that other of 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's those 43 percent. It turns out it's not, because when we looked at it, we found that only 40 percent of the people who accept that dictum agreed with the supermarket thing.

Those people are very inconsistent. The lack of correspondence, even when the scenarios appear similar, underscores that a small percentage consistently accepts the idea of tradeoff.

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

47

may make. So, for example, you might say I'll take the discount, but what if you know what the supermarket is doing with your data? This is knowledge Americans almost never receive directly, but may intuit from ads and coupons they think are targeted toward them.

So, we have a variety of things we asked them that said, will you accept it, to the people who said they would accept the discount in the first place, we said, would you accept it if they -- accept it if the supermarket makes assumptions based on your purchases about whether you buy low-fat foods, it went down to 33 percent. 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 percent 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 percent to around 20 percent isn't consistent 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 the idea of giving data for

48

discounts based on some kind of analysis. Then we went ahead and our hypothesis about resignation came out of our everyday realization when we met people, they would say things like, gee, you know, I have to give up the data, I have to be online, I have to be on Facebook, I know they do this stuff or I don't know what's going on, but I have to do it anyway.

So, we gave the people two statements 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 percent of people agree with both statements, which we say indicates a sense of resignation. Resignation meaning the acceptance of something undesirable but inevitable. I got that from Google, Google Dictionary.

We found there's a strong positive statistical relationship between believing in tradeoffs and accepting or rejecting various kinds of supermarkets' use of discounts. You would expect that. By contrast, there's no statistical relationship between being resigned to marketers' use of data and accepting or rejecting the supermarket tradeoff. People who are resigned, sometimes they do, sometimes they don't. They try and 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 tradeoffs, but in their heads, they're saying, gee, I'm resigned to it.

Put another way, people who believe in tradeoffs give up their data predictably, while people who are resigned don't do it in a predictable manner. They do give up their data, though. We found 57 percent of those who took the supermarket deal were resigned. A much smaller 32 percent were tradeoff supporters, even using the broader measure of tradeoff support that I suggested.

The larger percentage of people in the population who are resigned compared to those who believe in tradeoffs 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 that resignation is widespread across the U.S. 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 were resigned 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 percent 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 misconceptions suggest that even when Americans do weigh the costs 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 future -- as a protective feature as some academics have suggested.

So, what do we do about it? The rationale of tradeoffs 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 resigned about a key aspect of its everyday environment.

Now, this may sound really dark, and, you know, 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 a new era, not even 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 that we can talk about. Issues that the FCC might be involved in around public interest, convenience and necessity. The importance that people have alluded to in praising and naming groups that do right things and not so right things.

Thanks for listening.

(Applause.)

MR. BROOKMAN: Thank you, Joe. Thanks to all of our presenters, and now we're going to move into a

brief period of discussion. There's 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 are not angry with him, he's not mad at us.

So, I'm go to start with some of the trends I saw from the presentations. There's one, the proliferation and growing sophistication and growing complexity of online tracking is reflected in Abe’s and Steven's work. There are more cookies, there are more companies who are doing it. I love the revised Lunascape chart, with all the hundreds of thousands of companies that you couldn't even see them on this big screen. And more technologies, too, right? It's not just cookies, it's HTML5, it's fingerprinting, it's ETags, it's audio beacons, it's endless.

And then logically, perhaps unsurprisingly, the theory of Joe and Chris' 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 an informed choice that I am satisfied with how ETags are used on the site and I will now access my content in exchange for that is flawed.

And this builds somewhat on Lorrie Cranor's

53

work that if you had to read every single privacy policy, it would take like months of your life. And, so, instead of that, it sounds like that there's this resignation, right? Instead of privacy pragmatism, there is resignation. This is what Joe's work was talking about.

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, do you know how many cookies are in there?” And he was like, yeah, I know. I mean, he's not a privacy guy, neither of them, and he's like, yeah, I know, but what are you going to do, right? You could have tried to walk through deleting cookies or installing Adblock, but he had to pick up his kids, right, he didn't have time to really think about everything in his question. And we've all talked to people who have had similar experiences, we have all probably had similar experiences ourselves. I don't really know what's going on here, but I just don't have the time to figure it out.

And it's not just the web, right, I mean, it's the Internet of Things, we had our cross-device tracking workshop, we're letting TVs or toasters collect

54

information about us. It's physical space with automatic license plate readers, and are we making an informed choice when we go outside with facial recognition.

And, so, one thing I would like to hear from the folks about, and I am going to turn it over to my other co-discussants first, is, 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 with cameras now. Do we want government making, you know, rules about how much tracking can happen if consumers can't make the choices themselves? Say 15 cookies and that's it.

And, so, the point of PrivacyCon is we can hear from really smart people who are thinking about this to help them influence, you know, policy decisions. And, so, I would love to hear some of their thoughts or solutions later.

And I will ask you a question about that, but first I'm going to turn it over to Elana.

MS. 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 populus, including both consumers and, more

55

importantly, the sort of experts, advocates, policymakers, academics. And with that seems to also be a shift from the idea of questioning consumers' decision-making capabilities to whether they're, in fact, actually engaging in a choice at all, or you're resigned because they see no agency and no reasonable alternatives to opting out of the mainstream, or because they have trusted the default system.

If you look at the idea of what transparency is as 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 their being more resigned because they have a feeling of helplessness?

Also, how do you ensure 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 private potential privacy abuses, and which might also be most likely to be the most profit-generating core of many companies' businesses.

There's also the question of whether

56

transparency can operate as a mechanism to ensure 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 their choices -- a sense of helplessness in the sense of the unknown and what may happen in the future.

Finally, I am 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 then 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 mechanisms we also need mechanisms that consumers can see for due process and redress?

MR. BROOKMAN: All right.

MR. TENE: Thank you. So, I think all four presentations here drew sort of a 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 seem aligned like this, I feel an urge to play devil's advocate, and in that role, I am 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 they're actually thrilled or maybe even exhilarated or delirious about these new technologies, about the fact that, you know, they can hail an Uber and rate the driver, and get like the newest iPhone or Android phone, and, you know, even, yippy, take like a selfie and post it on their Snapchat story, or use a Fitbit and sort of 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 as three of the -- 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 in these places, including people who are now in government and even 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 other contexts. So, I care about health, but I still eat a cheeseburger. I care about the environment but I drive a four-wheel drive. There's a lot of snow in New England. And I think part of your response, your report 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 actually 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 a reaction.

MR. 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 is starting to get antsy. I mean, there are lot is of things that are terrific about this.

I couldn't live without Google. But what I'm saying, 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 this 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, and in some places it 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's fear, 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?

So, while I agree that there are many terrific things about this, I think that there have to be segments of society that have to say, stop, we can fix the really difficult things that relate.

MR. HOOFNAGLE: But let me unravel some of the issues that -- and there are -- what I'd say is that, first, that one can look at our work and say it's anti-technology, but I would argue strongly that it is not, and 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 of 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 the -- one of the rhetorical -- it's in the way a strawman 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. You know, long before Uber, taxi cab companies had hail apps in mobile. You don't need personal information for a lot of that.

But when you do need personal information, you have rules around it. And I see it from practice all the time. There are situations where we do very interesting forms of personalization, 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 into the false dilemma that privacy means we cannot have a spectacular convenience in our life.

MR. ENGLEHARDT: So, coming at this from I guess the tracking perspective, I wanted to comment on the fear of, you know, 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 the consumers, but for everyone.

And then, the notion that consumers might, you know, see -- or consumers just have to be tracked. I don't think that's really true either, because a lot of, at least for advertisers that support opt-outs, right, you should be able to set up 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 do things like fingerprinting, and the user doesn't just have to

trust that that won't happen. So, thanks.

MR. BROOKMAN: I'll ask one more question. And if we come to the broader policy question of what's the alternative, because I've talked to a lot of companies, and they kind of tell the same story, right? Oh, of course consumers can't control all this stuff, but they argue that really means there should be more of an accountability model, right? Companies should be responsible stewards of the data, right? Consumers can be in control, the company should make smart, informed decisions about how the information is used.

Because what is the alternative to that? That's one option, right? And then there is the FTC or the government could be making prescriptive policy choices on behalf of people. You know, that has its problems as well. One threat we've heard a few times today is the idea of increased transparency and then filtered through elites or institutions, and the name-and-shame approach that Joe and Steven talked about, and Elana talked about in her comments.

And I guess my question is, is that scalable, right? I mean, you know, 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 were concerned about is still increasing. Joe and Chris

have been doing this for even longer.

You know, so what is the policy solution, assuming that, you know, there is a problem to be addressed, you know, what is the right approach?

MR. TENE: Can I jump in and say that --

MR. BROOKMAN: Yes.

MR. TENE: Thank you. That I think, you know, and also reacting to what Chris and Joe said, I think there is consensus that we need to deal with data excess, and have, like, de- identification, and clearly we need strong data security, but I think to a large extent, 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 are having this conference and the existence of the privacy profession that has blossomed so the IPP now has 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 some of this in the technological research.

MR. HOOFNAGLE: Well, the access and accountability issues have to be dealt with, and there

were very interesting proposals to focus mainly on use of data, but I think one weakness of those proposals is they don't take into account the attacks on accountability that are occurring, such as the Spokeo case. You know, 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 shouldn't be able to be sued.

Wyndham was, in a way, an attack on accountability. You know, 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 will never be able to sue. Not even you, Justin. If you take the position seriously, not even the FTC would be able to sue.

MR. TENE: I think, you know, 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, you know, 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.

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

MR. HOOFNAGLE: I would like to defer to someone else, because it is not my area of expertise.

MR. ENGLEHARDT: So, can you repeat that question?

MS. ZEIDE: So, I'm just saying, in this case, when we talk about what the abuses are and the sort of harm, is it really about the uses 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?

MR. ENGLEHARDT: So, I think -- I guess the

fear, like I would say it's more of the data use, right? 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, like, preventing it from happening, right? So, once the data gets put into the company's databases, that kind of is up to trust.

MR. BROOKMAN: I mean, so I guess it kind of goes to the point, I mean, so should we be concerned about the collection itself, right? You guys both had in your studies, like, there's a lot more collection going on. And, like, I'm sure there a lot of people in the room are like, yeah, but it's not bad collection, right? It's not malicious collection. It's being done to support the ad ecosystem, which there's nothing inherently wrong with that, and then there has definitely been folks who have said that FTC should be focused on the cases where there is the harm down the road, right?

Commissioner Ohlhausen has written about this, there has been a lot of focus on the use of data for discrimination, right? We have a panel on that later today. I mean, so should we be focused at the FTC or 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?

MR. 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 can police use, we need to really see change and a form of accountability that doesn't really exist.

What my team has found over and over, when we discover things like HTML5 or Flash cookie responding, 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 that they're doing it.

MR. BROOKMAN: Any other closing thoughts?

(No response.)

MR. BROOKMAN: And with that, we are over time. Thank you all so much. We are going to have a quick 10-minute break and then we will come back with our next session.

(Applause.)

(Whereupon, there was a recess in the proceedings.)

SESSION 2

CONSUMERS' PRIVACY EXPECTATIONS

MS. ANDERSON: Please take your seats. We're about to start with the next session. Good morning, I'm Kristen Anderson, and I'm an attorney with the Division of Privacy and Identity Protection within the FTC's Bureau of Consumer Protection. I'm here to do the second session of the day, which is on consumers' privacy expectations.

We will hear from six researchers in four 15-minute presentations, and then we will conclude with about 20 minutes of discussion where we will identify common themes and ask the presenters about their work and its implications.

Without further ado, I will introduce our first presenter. We have Serge Egelman of the International Computer Science Institute at the University of California at Berkeley. Serge will start us off with his presentation on Android permissions.

Serge?

MR. EGELMAN: Thank you for that introduction.

So, this is work that I have been doing with several students recently where we've been looking at privacy and how private information is regulated on mobile platforms. So, to give you I guess a brief overview, most of this work is on Android, and that's only because Android actually has a pretty intricate permission system to try and implement notice and choice.

So, whenever an application requests access to certain sensitive data, it's regulated by this permission system, and so when users install an application, they see a screen that informs them of all of the possible types of sensitive data that that application might be requesting in the future.

And, so, the question was, does this actually implement effective notice and choice? So, do users understand these messages about how applications could be using their data in the future?

So, we started this project a couple of years ago by doing an online survey. We had over 300 Android users, and we just showed them screen shots of these permission screens, and simply asked them if an application was granted these abilities, what might that allow the application to do.

We then followed that up with a qualitative study where we had 24 people come to our laboratory, and we interviewed them about similar concepts. And what we concluded from this was that many people were simply habituated, since these appear every time people install

applications, not only does it list what abilities and, you know, types of sensitive data that application is requesting in the future, but all the possible types that it could request, even if the application never takes advantage of that.

And, so, people become habituated. They see lots of these requests that have lots of different data types, some of which they don't understand, and therefore, they learn to ignore these, because there's just so much information there. Another problem was that people were simply unaware. Since this occurs whenever you install an application, a lot of people said that, oh, this is just part of the license agreement, and we know that we need to click through that in order to continue installing the application. So, maybe this occurs at the wrong time in the process.

And since it happens after the user clicks install, it could be that they are already committed to installing the application; there are various cognitive biases that relate to this, and so therefore it's unlikely that they are actually comparison shopping based on privacy, even if they wanted to.

Another issue is that understanding of whether a particular application is going to access a particular type of data really requires a good understanding of

this whole permission system, and what are the different types of data that are regulated by the permission systems.

So, you know, understanding whether an application is requesting a data type requires understanding the whole universe of data types that are governed here.

And, so, we made these recommendations, and what we concluded was that a lot of this could be taken away. So, transparency is great. Notice and choice is good, but the problem is, when people are overwhelmed by the notice, which is what we see with privacy policies on websites, they eventually just ignore it all, because there's so much information.

So, you know, what we found was that a majority of these permissions could probably just be granted automatically without showing the user lots of information, because either the dangers are very low risk. For instance, you know, changing the time, for instance, or causing the device to vibrate, or is simply reversible. So, you know, if an application does abuse one of these abilities, chances are the user can find out about it and simply undo it and there's no lasting harm in that.

At the same time, there are a few very

sensitive things, which because of doing this at install time, that's probably the wrong time during the process, the user has no context about how the data might be used in the future, these could probably be replaced with runtime dialogues. But another open question is, this is just looking at all of the different abilities and data types that could be requested by an application. We didn't look at how frequently these data types and abilities are actually used in reality.

And, so, things actually improved. So, we did this study two or three years ago in the most recent versions of both Android and iOS. They now have a few runtime dialogues that prompt the user at the time that an application is going to first request access to certain sensitive data types. But the problem with this is it also -- well, so it adds some contextual information. The user is doing something, this dialogue appears, and then they could probably use information about what they were doing to make a decision about whether this request is reasonable or not. So, maybe clicking a button to find things near you, it then would be expected that an application would request access to GPS data.

The problem with this is, it only appears the first time that data type is requested. Once this is granted, the user never sees one of these dialogues again. And, so, future access to that type of data might be under completely different circumstances that might actually surprise the user or be really concerning.

And, so, another question we have is, how often are these types of data on mobile platforms really accessed in practice? And, so, we performed another study, last summer, where we looked at real applications in the wild, and we instrumented the Android operating system so that every time one of these data types is requested by a third-party application, we made a log of it. And then we gave these instrumented phones to 40 people, 36 of them returned said phones, and then we ended up with a pretty robust data set.

So, each time one of these sensitive data types was requested, and I'm talking about things like access to the contact list, GPS data, things like that. We also requested -- we also collected things about what the user was actually doing on the phone. So, contextual data. Things like the time stamp, whether the application that was requesting this data was even visible to the user, so whether the application was running in the background, maybe the screen was off. Most people don't realize that, you know, applications

might not be visible to the user and are still accessing, you know, data on the phone.

Connectivity, location, what part of the application they are currently viewing. So, what UI elements were exposed, that might yield some information about whether or not this access to sensitive data was expected or not. And then also the history of other applications that were run.

So, we let people use these phones for about a week. We transferred their actual real data on them, so they were using them as they would their normal phones, they popped their SIM cards into them. And then afterwards, at the end of that week, they came back to our lab, and we gave them some questionnaires. We randomly showed them some screen shots that occurred during the course of that week and then asked them questions.

So, these screen shots were taken randomly whenever one of these sensitive data types was accessed, so that we can ask them, as a prompt, you were doing something, this is what you were viewing on the screen of your phone, it was requesting this particular type of data, how -- you know, was that expected? Did you expect that application to be requesting that particular data type at this moment in time? And also, if you were

given the ability to, would you have prevented that from happening?

And, so, we used that as ground truth to see whether we could actually predict whether a user would have wanted that data to be accessed by the application or not. And, so, this resulted in, you know, we had 36 people participate, we had over 6,000 hours of real-time usage, and during that one-week period with 36 people, we found 27 million requests for sensitive data that was protected by this permission system.

So, some of the problems that we found were due to incorrect mental models. So, again, you know, the goal of this is transparency, show the user all the possible ways that an application might be accessing sensitive data. Is that really working?

Well, we found that in 75 percent of cases, the application that was requesting one of these data types was completely invisible to the user. So, this was mainly due to the screen being off, in 60 percent of the cases. So applications running, you know, the user wasn't actually using their phone. Or background services.

Another thing that we found was, despite the fact that there are some privacy indicators built into the operating system, so both Android and iOS have

indicators for when GPS is accessed. There's a -- this is an example of one of those indicators, it appears in the top status bar and most people assume that, you know, the only time that GPS information is collected, this icon will appear.

It turns out that's not true at all. And, in fact, the icon only appears in 0.04 percent of the cases where location data was accessed. And that's because every time an application requests location data, the operating system caches that for performance reasons and also to preserve battery life, but then when another application accesses just the cached location data as opposed to querying the GPS hardware directly, this icon never appears.

Similarly, applications can infer your location based on cellular network data, nearby WiFi hotspots, and it turns out most applications are using those methods to infer location, rather than the GPS hardware. And therefore, most of the time, when location data is collected, people have no indication that that's occurring.

So, you know, what if -- so, having this, the notice and choice at the beginning when users install the application obviously doesn't work. We've tested that. The ask on first use that's currently happening isn't really working because of the different contexts in which users might be interacting with applications.

So, maybe we could have runtime requests all the time, so every time applications request data, we can have a little notice appear. Well, obviously that's really impractical, too. So, you know, the 27 million data points that we collected, that would result in, per person, about 200 popups per hour, most of which is due to requests for location data, but you could see that there are other data types that were pretty frequently requested. And, so, having lots of popups appear on the phone is not really a good way of going forward either, because that's also going to lead to habituation.

But at the same time, in our exit survey, what we found was that the vast majority of participants said that given the opportunity, they would have denied at least one of these requests, and on average, you know, they would have denied a third of the requests.

So, how do we do this? How do we give users control over the things that they actually care about without overwhelming them? So, we're doing some work now to try and predict the cases where applications access data where people would want to know that this is occurring, whereas the other ones where applications access data that might be expected, well, obviously we

shouldn't prompt the user in those cases.

And, so, what we found was that expectations really did predict behavior in this case. So, we asked people if this access to personal data was expected or not, and then whether they would have blocked it, there was a pretty strong correlation there.

We also found that using the current model on ask on first use, so if you look at for each unique application, in each unique data type, if you ask users the first time that application requests that data, we're going to get it right about 50 percent of the time, which is what's currently happening. So, that's a coin flip.

But we also found that looking at, you know, the visibility of the application was a pretty strong predictor of user expectation. So, applications running in the background requesting data were pretty often -- those were unexpected. And, so, if we add that to the equation, we can get this right about 85 percent of the time. So, instead of just asking on the first use, we could ask the first time that the application requests the data in the foreground and then ask the first time the application requests the data in the background, and then we're going to get it right about 85 percent of the time.

But one of the main things we also observed was that the data was really nuanced. So looking at one user's preference and comparing that to another didn't really work among our 36 participants, because there was just so much variance in the data with regard to what people wanted and what their expectations were. Which suggests that, you know, having a one-size-fits-all solution about what people care about and what should they be shown is unlikely to work either. And, so, maybe we need more intelligent systems that can predict user preferences on a per-user basis.

So, going forward, we're actually trying to implement these systems right now that can try and predict a given user's preferences based on their previous behaviors. And this is part of a pretty complex ecosystem. So, we have what we're calling hard policy, which is preferences that people have explicitly stated. So, I don't want applications to be using data for X reason, and then trying to augment that with soft policy. So, inferred preferences that systems can make about users, such as maybe looking at, you know, hundreds of thousands or millions of users, we can infer one user's preferences based on other users who are like them. Like recommender systems.

And also, you know, based on the feedback from

prompts. So, if we can design more efficient prompts that cater to individual user expectations, we can then use the output of those, so what did the user actually decide to ensure that they see fewer prompts in the future. And that's it. I'll leave it at that.

So, well, the conclusion is, you know, notice and choice is great, the problem is figuring out what notice to give people, since attention is a finite resource. So, I'll leave it at that.

(Applause.)

MS. ANDERSON: Thank you, Serge.

Next we will hear from Ashwini Rao of Carnegie Mellon University about mismatched privacy expectations online.

MS. RAO: Hello, thank you. So, yeah, my talk is about expecting the unexpected, understanding mismatched privacy expectations online.

So, I'll start with the motivation. So, many of us on a daily basis interact with online websites, and as we interact with online websites, we may have questions, such as what types of data does this website collect about me, how does it share this data, and does it allow deletion of this data?

And to answer these questions, a user could read the website's privacy policy, which is usually a

81

textual document in English, and it discloses the data practices of the website, such as collection, sharing, and deletion; however, these policies in their current form are long and difficult to read. So, users usually ignore them.

So, the main motivation is, how can we help users understand online data practices? And our approach is to focus on user expectations. So, we assume here that users expect websites to engage in certain data practices. For example, users may expect banking websites to collect financial information and health websites to collect health information.

And these expectations may vary based on context; for example the type of website or user characteristics: Age, their privacy knowledge, their privacy concern. However, user expectations may not match what websites actually do. For example, users may not expect banking websites to collect health information.

Now, the question here is could we generate effective privacy notices by extracting and highlighting these data practices that do not match user expectations?

So, the concept is simple. A privacy notice does not have to inform you about things that you

82

already expect or know. A privacy notice has to inform you about things that you do not expect or do not know.

So, I want to make a distinction between policy and notice. A policy is usually a textual document, but a notice, which is based on the policy, is usually more -- is usually shorter and more usable.

So, here I'm showing you the privacy nutrition label, which focuses on visual format, and so far, notices that make -- that are more effective, our research has focused on visual formats. And our approach of extracting and highlighting mismatched expectations is complementary to this approach.

Once we identify and extract these mismatched expectations, we could present them to the user in any visual format that is effective. I also want to say here that these privacy notices do not have to be generated or provided by the website operators themselves. These could be provided by a third party, for example through a browser plugin. And this is something important to note.

So, the main research questions are how do we define expectation, and how do we measure expectations, and identify mismatches in these expectations? So, research in nonprivacy domains shows that users can have different types or multiple types of expectations. And

83

privacy research has predominantly not focused on multiple types of expectations.

So, in our research, we make a distinction between two types of expectations. The first: Expectation in the likelihood sense. What does the user expect that the website will do versus what does the user expect the website should do? And this is in the desired sense. And then we compared that with practices, data practices of websites.

To measure expectations, we conducted user studies. So, one of the user studies that we conducted focused on the expectation in the likelihood sense. And in future, we also plan to measure expectation in the desired sense. So, we presented users with different types of websites, and then after the users interacted with these websites, we asked them, what do you assume that the website will do?

And once we elicited user expectations, we next extracted the data practices from privacy policies, and then we compared these two to identify mismatches. So, in our study, we used -- we varied the website characteristics and user characteristics. So, as I mentioned earlier, user expectations can vary based on these website and user characteristics.

We looked at 17 different data practices,

84

which were split among collection, sharing, and deletion. And for collection and sharing, we looked at four different types of data: Contact information, financial, health, and current location information.

So, here's an example scenario. So, here the scenario is describing the collection of different types of data when the user does not have an account on the website. So, you can see that we are asking the user, what is the likelihood that this website will collect your contact information?

So, in future if we wanted to also measure desired expectations, we could also ask them, do you think the website should be or should not be allowed to collect this information, in addition to do you think it's likely that the website would or would not collect this information.

So, we deployed the study as an online survey, and we studied in total 16 websites. We had 240 participants that we recruited from Mechanical Turk crowdsourcing platform.

So, this was to elicit user expectations, the latter part is to actually extract data practices from privacy policies. And to do this, we used two annotators, two experts, one in the privacy domain and another in the legal domain, and they manually read these policies and answered questions such as does this policy disclose that the website collects health information?

Now, to scale up, we are also developing techniques that are semi-automated and that use natural language processing and machine learning that can go and extract answers to these questions. So, the annotations say whether a website is clear, whether it engages in the sort of practice, it does not engage; whether it's unclear or the policy does not contain any statements that addresses this data practice.

Now, it's important to note that there can be different types of mismatches. Here I'm showing you two, the yes/no mismatch and a no/yes mismatch. And this is important because the type of mismatch can impact users' privacy differently.

So, consider the yes/no mismatch. The website states that, yes, we collect your information, but the user thinks, no, the website is not collecting my information.

So, in this case, the user may go ahead and actually use the website, and unknowingly give up data. And lose privacy. Whereas in the no/yes mismatch, the website is saying, no, we do not collect your information, but the user thinks incorrectly that,

indeed, the website is collecting their information. So, in this case, the user may decide not to use the website, in which case the user may lose the utility but not privacy.

So, some results. So, we have looked at different types of website characteristics, and we found that only website type had a statistically significant impact, and the type impacted users' expectations only for financial and health information, but not for contact or current location information.

Several user characteristics also had a significant impact on what users expected. So, for example, users' age impacted whether they expect websites to allow deletion of data.

So, now, here I present two examples of mismatches that we found. This one is a mismatch in collection data practice, and this is an example of a yes/no mismatch. So, websites can collect users' information, even when users do not have an account on the website. However, users do not think that happens, or they do not expect that data practice.

Now, compare this with a no/yes mismatch, and this is a mismatch in sharing data practice. Users expect that websites will share their contact information for any purpose; however, websites do not do

so. They only share contact information for specified and very narrow purposes.

So, as regards to deletion, users predominantly expect their websites to allow deletion of the collected data, but websites generally do not allow that.

So, there can be other types of mismatches as well. One example is a website-specific mismatch. For example, users do not expect banking websites to collect health information, and most of the banking websites we looked at do not do so; however, there can be specific websites, for example Bank of America, which was one of the websites we looked at, that indeed collect health information. So, you can see this is a mismatch that is specific to a certain website.

So, based on the results of our study, we could come up with notices that have less amount of information than a full notice. For example, we looked at 17 data practices. A notice could show information about all 17 data practices, or we could show information about data practices where there's a mismatch between what users expect and what websites do or actual data practices of websites.

So, for example here, for the Bank of America privacy notice, there were mismatches for 11 data

practices out of the 17. So, that's -- if you show only 11, that would be about 35 percent reduction in the amount of information that the user has to read and process.

We could also just show information about mismatches that are more privacy invasive from a user standpoint. For example, I talked about the yes/no mismatch versus the no/yes. If we find that the yes/no mismatch is more invasive, we could only show information about those mismatches, and in the case of Bank of America, it's only five data practices for which there's a yes/no mismatch. So, that would be 70 percent reduction in the amount of information shown in the notice.

However, the caveat here is that we do have to go ahead and test with users how effective the shorter notices will be, and -- yeah.

So, as part of future work, we are planning to also study expectations in the desired sense, and compare that with expectations in the likelihood sense, and may also compare both of them to actual data practices of websites.

We are also, as I mentioned, we will test the effectiveness of notices that highlight mismatched expectations and see whether they actually reduce user burden and see whether users can make better privacy decisions.

Yeah, that was all. Thank you.

(Applause.)

MS. ANDERSON: Thank you, Ashwini.

Next we'll hear from co-presenters Heather Shoenberger of the University of Oregon and Jasmine McNealy of the University of Florida. They will be presenting on reasonable consumer standards in the digital context.

MS. McNEALY: So, good morning and thank you for having us. Our project is online or -- offline versus online, re-examining the reasonable consumer standard in the digital context. The impetus for this project is really trying to get a deeper understanding of how consumers act when online.

So, we know from prior literature that people, individuals, act differently, supposedly, offline than they do online. So, we wanted to take this into a further exploration of consumers. And we know that the reasonableness standard is a standard that is used for regulators, for example, in assessing complaints related to deception. So, we wanted to find out more and explore this a bit more.

So, we came up with an umbrella project that

used mixed methods to examine this question. One of the first things we did was start to interview. We did qualitative interviews, and just to skip forward a little bit, so we asked our interviewees questions related to how they behave both online and offline, and we have this quote from an interviewee who we asked questions related to their expectations related to privacy or how their information would be used and how they attempt to control their information.

So, when we asked about whether or not they showed photos offline, if they just met a person, so it's a stranger, they invite them into their home, and they break out their family photo album. We asked about that. And the interviewee said, you know, I would wait for a friendship to develop offline before showing any photos to someone in person, but this seems almost diametrically opposed to what they do when they participate on, say, Facebook or Instagram, right?

But more importantly than just showing photos, we asked the question about whether or not they would sign a printed contract without actually reading the terms of the contract, versus whether or not they always click yes or no to the terms and conditions of using various websites, whether it's social media or shopping or whatever the case may be. So, we wanted to find out.

So, we noted that there were some significant differences indicated with respect to their sharing behaviors, both online and offline.

So, to go back a little more about our method. So, again, we used qualitative and quantitative methods. So, just some breakdowns for our interviews. We had 30 participants. We did these long-form qualitative interviews, and we are going to do more long-form qualitative interviews as well.

So, we had 20 women, 10 men. We note the average age was around 26, and then we have some racial demographic data broken down as well. Then for our quantitative side, we did a survey, and we're going to talk a bit more about the results of the survey today, and there were 871 participants. Almost equal breakdown between men and women, but note the age. So, we had an age of 35.9, so almost a 10-year age difference on the survey, the qualitative side, and again, the breakdown of racial demographics.

Also important are some of the variables that we used or we attempted to measure in our survey. These variables we got from prior literature. They also emerged, again, when we were doing our qualitative interviews, and one of those important ones was social trust. Social trust was measured on a six-item scale,

and social trust is really asking the participants, you know, how they felt about whether or not they trusted that the institutions, the entities, you know, brands or advertisers, the government, news media, also, how they felt that they would -- whether or not these entities would fulfill their responsibilities related to the consumer's private information.

And, so, those are example questions on here as well. And then the second important variable we attempted to measure on a four-item scale was control, or how participants perceived they had control over their information.

So, an example question was, I can use online privacy tools to remain anonymous online. Perhaps more importantly are our main dependent variables, so we had the always click yes. So, again, we're assessing behavior, whether or not the participant always chose to click yes related to privacy policies or terms and conditions online.

And the second one was privacy concern was -- we measured on a three-item scale about whether data companies -- whether they thought that data companies would collect information about them that would make them feel uncomfortable.

And Heather is going to come and talk about

93

some of the relationships we found.

MS. SHOENBERGER: Right. So, we diverge a little bit here, where we're very positive about our findings. And also I wanted to note -- well, I'll note that in a second.

So, our always clicking yes variable was our indication of behavior, as our DB. This was a hierarchical regression, and I made it very simplified for this, because we are under a time limit. The first block was demographics. The only demographic in this particular equation that was significant was age, and it's no surprise that it's younger people that predicted always clicking yes. We've seen this in numerous reports where younger people tend to be a little bit more careless online, maybe a little bit more apathetic, et cetera.

Then we moved to a second block, and these are two variables that did come up in our surveys and also have been used in numerous studies before ours, and social trust in this particular case was not a predictor, but control efficacy was. So, even though they may not actually be able to control their data, the belief that they can predicted always clicking yes, and we believe this is the result of the confidence that people have if they believe they have control, and as a

94

result they go ahead and say, sure enough, I'm just going to go ahead and click yes because I'm confident and I trust that this is going to work out for me.

Those who had had -- oh, so the next block were all items that were derived from our interviews. Of course, some of them you've seen in previous studies as well, but all of them were derived from our interviews.

So, negative experience, those who had had fewer negative experiences, self explanatory, but more likely to click yes without reading any terms of agreement, no further investigation. Peer recommendations, we were really hopeful that a peer recommendation would kind of be an if/then rule, if a peer recommends Snapchat to me, I would go ahead and download it. That was not the case in our regression analysis. It wasn't significant.

Convenience was a pretty big variable made up of items like that policies are too long, they take -- it's faster to just skip them. They're full of legalese. Some of the information that we heard last night at the conference about how these policies are just laden with too much material for consumers to ingest, especially in an over-saturated environment with jobs and time constraints, et cetera.

95

And then the two variables that are really important to us for this study were both essentially cues. One was site appearance. If the site appeared to be safe and not weird, itdidn't raise any skepticism. Again, we've seen this in previous studies, but our participants noted this in interviews as well, predicted clicking yes if this looked safe and also was familiar.

And then just simple presence of a privacy policy or an icon like TRUSTe also predicted clicking yes. So, this was our behavior.

And at the conclusion of this, we thought, we're on the right track here, these cues are what is driving the motivators of actual behavior online, and we were really excited.

Then, we got even more excited for our privacy concern variable, a variable that has been heavily researched in this area. Many researchers have noted the -- and this panel, the panel before us noted that there is a disconnect between privacy concern and actual behavior. We may have a potential to bridge that with this research.

So, in the -- hierarchical regression is in the exact same format, higher ages and higher education, again, no surprise, predicts privacy concern. Lower social trust, trust of the institutions, predicted

96

privacy concern, lower control efficacy, both in line with previous research.

People who had suffered more negative experiences were more likely to say that they had higher privacy concern. Again, peer recommendation, we had high hopes for that, but it didn't work out. Convenience fell out of this model as a result of the two cues at the bottom, and there are definitely, within the same direction as before, if the site had poor esthetics, and it was ugly or weird, it made people feel more skeptical, predicted privacy concern. And then a lack of privacy policy or a link or an icon predicted privacy concern. And note that both of those two cues predicted both the concern and the behavior.

So, in this study, our aim was to better define the behavior of the average consumer online, and it appears that while they're not -- they're specifically not reading policies, especially when these safety cues exist. So, it leads us to have the same conversation that the rest of our panel has had, where if they're not reading the policies, can there be meaningful notice and choice? Of course, that's a question for potentially another day.

And if we make really clever use of the cues, and there may be more than the ones that we explored, both entities who collect data, so businesses, advertisers, the government, news media, who use data, can reduce privacy concern, which is something that they would like to do, encourage the free flow of data, another something that they would like to do and something that last night was mentioned that the Federal Trade Commission potentially may be interested in doing, also, and increase trust.

And on the flip side of that, consumers could rely on cues that are more uniform and meaningful, even if they don't read the privacy policies that underlie those particular cues.

So, with that, now I really have to move through this very quickly. So, there's really a three-prong approach, and we have already begun the research to sort of decide whether or not this is the right approach, but we would like to suggest guidelines for different types of data collection and use, something that has been echoed on the panel already, based on the average consumer's expectation of privacy.

So, there is some additional research to do, and then delineate those types of data collection and assign a cue, or a heuristic to each type of data that would be endorsed by the FTC. Here is the catch for people who are in the advertising industry who are in

the business of collecting and using consumer data. They would have to adhere to those guidelines in order to use the cue on their sites, which would signify safety, increase trust, hopefully, et cetera.

So, we would also do research on what icons would be most effective to consumers, and also link those icons to readable policies. Another thing that we noted was the convenience variable was made up of items like it's too long, it's full of legalese, we don't understand, and if we could make those policies readable and approachable to the consumer, something that we can do in the lab, we can test this, we could potentially also for that small sect of people who are going to read those policies, they at least will have the opportunity to make meaningful choices, and it will be short, quick, and more concise.

So, in conclusion, we are continuing to pinpoint consumer expectations of privacy in a way to develop these guidelines and the resulting cues that would align with the guidelines. As Jasmine mentioned, we are continuing to collect data both in the interview portion of the study and also in the survey, just to make sure that we have as close to a census as possible, because we are dealing with the average consumer in the United States, and we want to make sure that we get

everybody. Examine some additional contextual variables as they arise, because while the cues that we mentioned are really good predictors, they may not be the only ones.

Design policies for readability and understanding for consumers so that they have the opportunity to make meaningful choices if they do, in fact, read those. And, finally, something that we think is really important, and I will diverge for just a second, literally just a second, in Australia, there was a really great PSA to help people avoid being hit by trains, it's called Dumb Ways to Die, and it's gone viral and has actually resulted in lower train deaths, and it's really a silly video.

You can look it up on YouTube, there's these little people or little animal type things dancing around and talking about dumb ways to die and don't get hit by a train. And essentially, we're looking to do a PSA like that, that would be based on research in America, I mean, that worked in Australia, it may not work here, to allow both the consumers to understand what these icons mean, how they can use them as a way of increasing trust, and how -- and also to entice entities to go ahead and opt into this system and adopt the guidelines the FTC has put forward in a way to align

with consumer expectations.

And with that, we conclude.

(Applause.)

MS. ANDERSON: Thank you very much, Heather and Jasmine.

Our final presentation is by co-presenters Andelka Phillips of the University of Oxford and Jan Charbonneau of the University of Tasmania. I think you two win for the longest commute today. And Andelka and Jan will be presenting their work on privacy in the direct-to-consumer genetic testing space.

MS. CHARBONNEAU: Well, first off I would like to thank the FTC for the opportunity to discuss our research. We're going to talk about privacy of a specific type of data, that being genetic data, the data that results from genetic testing. So, a very specific type of data.

What we have to realize is, genetic data is the most personal data there is out there. Not only is it a unique identifier of us individually, but because of the familial nature of DNA, it can also identify our families. So, when we're talking about privacy in this context, we're talking about it in a much broader context, not just personal, but looking at the family.

We also know that this data is inherently

101

identifiable, okay? There's growing recognition that it is simply not possible to de-identify this data in a way that makes it impossible to re-identify it. It may take a good skill set, but as we get increasing numbers of genetic databases out there, as there are more public databases, we know that we can re-identify that data.

The other thing is, this data is irrevocable. If there's been a privacy breach, you can't change it. It's not like your iTunes password, you can't come up with another one. Okay? So, this is a different type of data.

Does it matter if this happens in a direct-to-consumer genetic testing situation? Well, the first thing we have to realize is the difference between traditional genetic testing and what happens when we have genetic testing in a direct-to-consumer setting.

Traditionally, genetic testing has happened within a country's health care system, and that's important because when an individual gets a genetic test in their health care system, they're deemed a patient. And by being called a patient, that enlivens a whole host of professional and regulatory oversight, existing legal duties of care, and simple things like doctor/patient confidentiality. So, all the government systems for data protection of health care kick in,

102

because that's a patient.

When we look at direct-to-consumer genetic testing, we have to realize that at its core, this is a commercial transaction that occurs in each country's marketplace, and increasingly in market space, because the majority of the activity is actually online.

When an individual engages with DTC, they engage as a consumer. What that means is that enlivens each country's consumer protection legislation. It also enlivens some particular legal protections in contract, negligence, et cetera. Okay, but a very, very different situation.

What does the general public think of when they think of privacy? At the Center for Law and Genetics at the University of Tasmania, we've been looking at genetic privacy issues for the last 20 years, and in the last few years, we've moved into DTC. Some of our early research in direct-to-consumer genetic testing suggested from the Australian general public's perspective that privacy concerns were going to be the key constraint on commercial uptake.

Interestingly, this past year, we found the same results when it comes to intention to biobank, in other words giving a genetic sample into a genetic database for nonprofit, institutional, and

103

health-related research as opposed to commercial.

We've also modeled the DTC space, and that was an interesting exercise, and forced the thinking to go broader than just the consumer/company interaction. What we realized very quickly was not only does DNA go a lot of places, that sample travels from labs to companies and who knows where, through the postal system, usually, but also those results can go places. Okay? The actual genetic data about those individuals gets spread around.

And that informed the research that I'm going to talk about today, which is an online panel of 3,000 respondents of a thousand American, a thousand Australian, and a thousand UK respondents. We've just added in a thousand Japanese respondents, which will give us some interesting contrast.

The way the sample broke down, about 10 percent of the people are actual consumers, and that equates to those early adopter categories. That leaves about 90 percent of my respondents who are the potential consumers. So, we're able to look at actual versus potential consumers.

So, what does privacy mean from the general public's perspective? Well, simply stated, if something is private, it's not shared; if it's shared, it's not private. Okay? The simple way. And that's how the

104

general public looks at these things.

Privacy issues arise from sharing. So, privacy is all about control over sharing. Providing your permission to share means that you have control over your privacy. So, that's the way the general public looks at it. If my permission is asked, then I know what's being asked for, I have the opportunity to ask questions, but I also have the opportunity to say no, and that my no will be respected. So, I have control over my privacy if my permission is sought.

So, what do consumers think about whether or not their permission will be sought? In other words, as the previous presenters alluded to, this area of perceived control. Well, interestingly, the American respondents, 47 percent, thought they had perceived control. And what's interesting is, on any dimension that I analyzed on, Americans are statistically different to the other consumer groups.

For the UK, it's 43 percent; for Australians, it's 40 percent; and for Japanese, it's 36 percent. So, that's quite a difference in terms of whether or not people think their permission is going to be asked. Are they in perceived control? If they are in perceived control, what does that mean? Well, they're more likely to purchase the DTC tests. They're more likely to

105

participate in DTC research. And that's important because that's permission-based, right? They asked their permission, but do they actually realize that what they're doing is giving nonspecific enduring consent?

They're also more likely to share broadly. They'll share with family, not friends, so there's some control. They'll share with their doctors, and that's important because DTC companies very clearly state their results are for recreation, education, or information only; they are not a diagnosis. But as Graeme Suthers said, it would be a very brave GP who would action a DTC test.

If they go to their doctors, they're back into the traditional system. They're also more likely to engage with online sharing communities. But does perceived control equate to actual control? These are commercial transactions governed by contracts and privacy policies. We did some research in Australia, looking at the privacy policies of the DTC companies operating there. Do they comply with our legislation? The short answer, no, they do not.

I'm now going to hand it over to Andelka to talk more about contract terms.

MS. PHILLIPS: Well, I've actually been looking at the contracts and privacy policies of

106

direct-to-consumer tests for companies that offer tests for health purposes. Now, as has been noted in the previous session, and also in the previous group's work, these contracts and privacy policies appear everywhere online, basically any website you use, any software update you make will be subject to terms and conditions, and they'll be presented either as terms and conditions, terms of use, terms of service, privacy statements, privacy policies, and sometimes in this context they're combined in one document.

At present, these are used to gather not just the purchase of DNA tests, but also using the website, and sometimes participation in any research the company is doing.

Now, as several people have previously noted, people don't tend to read these contracts and privacy policies, partly because there are just so many, and it would take too long. This industry is no exception to that, and I would also say that similarly to most e-commerce, these contracts are also not industry-specific, so they don't necessarily address all the issues raised by the industry and what they're doing with data.

And because of the ubiquity of these contracts, consumers often also display inattention or

107

blindness online, so they may just not notice things, we may not read them, we just click on "I agree." And this is really problematic in this context, and I think there really needs to be reform, because unlike some of the other what was said that consumers don't read these, I’ve had to read 71 contracts, and I really think that there are problems here.

So, the major privacy risks in this context arise chiefly from sharing or sale of sequence DNA, but also from sharing or sale of other types of personal data, often health data or other data that we might normally consider to be sensitive.

This is because companies are often engaging in ongoing health research, so they're collecting large amounts of personal data from consumers. There is also the risk of possible discrimination based on a person's genetic makeup.

And then there are some other risks that arise. Some of these are more future risks, so there's a possibility with the increasing use of biometrics that in the future these genetic databases could be used for identity theft, targeted marketing, the most obvious example at the moment is targeted marketing of drugs to particular population groups or even family groups.

Also there's a potential for discrimination in

108

employment or insurance, if this data is shared inappropriately. And, more remotely, there's the risk of creating synthetic DNA.

Now, as I previously noted, these contracts are not industry-specific, so often you'll encounter the same terms in these contracts that you would when purchasing a product or downloading a song online. And they also use really similar wording.

Now, in the United Kingdom and the European Union, there is strong consumer protection legislation that deems some terms in consumer contracts to be unfair and unenforceable, and at present, some of these terms would likely be deemed unfair and unenforceable. And this is interesting because I know I'm at the Federal Trade Commission's conference, but I've been looking at mainly American companies, this is an American industry overwhelmingly, but these tests are sold internationally, and people's samples are being sent across borders, and, so, there is a need for international collaboration to protect consumers in this context.

So, one of the most concerning things here is that consent will often be deemed through use or viewing of a website, and often consent to altered terms will also be deemed through continuing to use the website.

109

Now, as most of you are aware, it is often easy to use a website without ever looking at the terms and conditions. So, this is quite concerning, because the other thing that's very common, and the majority of companies will include this, and 39 percent of companies include a clause that allows them to change their terms at any time. And only a very small percentage, around six percent, will actually -- or it might actually be four percent, I'm sorry -- yeah, six percent, will notify a person directly by email of changes.

So, most of the time, companies can change their terms at any time or from time to time, without direct notice to the consumer. And this is important here because it could have an impact on what companies do with your data. They could change their policies on sharing, sale, or storage of data, and this can significantly impact consumers.

As Jan mentioned previously, too, because this is marketed as a consumer service, companies are often including clauses that say that their services are only for research, informational, or sometimes even recreational purposes. Now, in the context of health testing, I would question whether anyone orders a breast cancer risk test for recreational purposes.

And skipping on, quite a few of them also

110

share data with law enforcement, which consumers may not be aware of, and there's often very broad sharing with potential third parties that might include affiliates, and -- yes, I'm running out of time -- but I really do think there's a need to improve these contracts. And following on from the previous two discussants' work, I really think these contracts need to be written in a more easily understood way that would enable consumers to make informed decisions.

So, thank you very much.

(Applause.)

MS. ANDERSON: Thank you, Andelka and Jan.

So, now it's time for our discussion session. We'll be spending about 20 minutes, which I'll be leading with my co-discussants, Alan McQuinn of the Information Technology and Innovation Foundation, and Darren Stevenson of the University of Michigan and Stanford Law School.

So, I'll just start us off. We're each going to provide some brief comments about what we have heard, and then we will ask the presenters about their work and its implications.

So, first, to me it seems like you are all striving to answer some of the same basic questions. So, what do consumers think about privacy and why? And

111

those include things like what are their expectations and hopes about the kind of data that's going to be collected and how it's going to be used, how much control do they have, what affects their understanding, what affects their willingness to trade privacy consciously or subconsciously, or unconsciously, for some benefit, and is that contextual, does it vary by the trust of the firm or online effects.

And I noticed three common themes in your answers or your findings. The first is that notice seems to be failing. So, Andelka and Jasmine's paper talked about the ubiquity of form contracts and how companies have begun to incorporate crook clauses that don't seem to be related to the purpose of the contract from the consumer's perspective, but do give the company whose policy it is some sort of an advantage.

Serge found that about 75 percent of permissions were being requested invisibly. Ashwini found 40 percent of collection practices that she was looking at in her study were not addressed or were unclear in the privacy policies. Ashwini, Heather, and Jasmine found that consumers were relying upon things other than privacy policies to decide whether they are going to use an app, and even to form their expectations of what's happening.

112

The second theme is that companies' policies and practices aren't matching up with consumers' expectations. Ashwini found rampant mismatches between expectations and reality, Andelka and Jan found that half of DTC companies' policies allowed them to share consumers' personal information with third parties, contrary to what consumers would have expected, Serge found consumers would rather not allow so much access to their data.

And the third theme was that several of you were recommending that companies highlight unexpected data collection and use, especially when it involves sensitive information. Serge was recommending runtime prompts and indicators when apps were accessing protected resources, Ashwini recommended highlighting unexpected uses, Andelka and Jan were recommending highlighting key clauses and providing shorter, clearer notices.

Now, one of the biggest benefits that I see of PrivacyCon is that it brings all of you together, the best and the brightest, all working to understand the same issues and providing us with the benefits of your learning.

So, we are hoping that this conference is going to facilitate you learning from and building upon each other's work, and I hope that we can continue to benefit from the insights that you have given us about how best to protect consumers' privacy, and industry can hopefully do the same.

As Chairwoman Ramirez said, it's now more than ever that we need to stay up to date with the latest findings on privacy and data security research in order to fulfill our mandate to protect consumers, and your efforts deepen our understanding of our own research in that respect.

So, thank you all again for coming, for sharing your work and your thoughts, and with that I am going to turn it over to my co-discussants for their thoughts and allow them to ask first questions.

MR. McQUINN: Thanks, Kristen. Thank you to the FTC for letting me come here today and respond.

I thought that the all of the presentations were very thought-provoking and they could definitely help businesses better understand their consumers. Now, but we're here today at the FTC, and what I'm looking for is evidence of the need for public policy intervention. And, frankly, I'm not sure that there is much.

As we walk into this, there's definitely a lot of discussions over different public expectations versus

privacy, or people not understanding the legalese in direct-to-consumer genetic contracts, but is that a public policy problem? I'm not so sure.

Let me draw an analogy. Say I'm not necessarily sure what goes into my Chipotle burrito. Sure, I'm able to pick different fillings and -- I may be able to pick different fillings, but I'm not so sure how they're sourced. So, when you ask me questions about what's in my Chipotle burrito, my expectation may differ from the reality of what's in there.

Now, that's not necessarily a public policy problem, right? But what is a public policy problem is when consumers start to get sick or have food poisoning as a result of the contaminated food from the Chipotle burrito, but when I'm listening to these presentations and reading these reports, I'm not necessarily -- I'm seeing that we're talking about what's in the privacy burrito, rather than actually talking about the privacy food poisoning. That's just some food for thought, I guess. And I look forward to a good discussion.

Thank you.

MR. STEVENSON: I have no way to connect to the burrito, but we wish Chipotle well with their current issues.

So, at the risk of stating the obvious, I

think what we have here is we have evidence, empirical studies that show that consumers have expectations. All of you in this room, you guys are not ordinary consumers, because you're here at PrivacyCon, but ordinary consumers we're seeing that there are consistent measurable expectations.

I really enjoyed the studies and I encourage you all to read them if you have not read the papers. And I think most of these papers have supported this notion of contextual integrity that's popularized by Nissenbaum and others, but the idea that pre-held expectations are measurable and can be demonstrated.

Two complications come to mind. So, the first is the difference between expectations and preferences. It was clear in Ashwini and colleagues' papers, they were really careful to define what is an expectation. What are we actually studying here, and then to contrast that with consumers' preferences. Expectations being different than preferences, which we saw in Turow's work, and with his colleague, that consumers might just be resigned, so that where expectations and preferences converged, I think this is a very fruitful area of study.

So, what are we measuring when we are measuring consumers' expectations, is it what they are

just resigned to give up, or is it what they would prefer? And in the papers, a few of them kind of went back and forth on that.

A second complication that comes to mind are expectations or preferences or we'll just say consumer tastes. This is a moving target, so these are continually changing, so even though they're consistent and we can measure them empirically and the FTC can decide does it warrant intervention based off of trend, these are evolving and they change over time.

So, how can policy, which tends to move slowly, track and be responsive to something that is changing, that is dynamic? So, if we were to have PrivacyCon in three years, next year, five years, and we repeat all these studies of consumers, would we see the same expectations. So, how can policymakers incorporate this sort of moving target of consumers' expectations?

So, I look forward to our discussion here and we can open it up to questions. Or if you have any responses to our comments.

MS. McNEALY: I like the burrito analogy, but at the same time, if Chipotle has lean steak or whatever they have, right, they -- I mean, if they make representations to the consumer that it's from a certain source, then you have expectations that, hey, my beef is

117

from a certain source, and even if we don't know exactly where it's from, we have an expectation that we should get at least a product of some I guess quality, or at least we expect the regulators would enforce them, you know, would enforce the restaurant giving us a product that either won't make us sick or won't have been, you know, had something done to it by a worker there, right?

So, I think there is a certain level of protection we expect from regulators with respect to things particularly like privacy. Just I think most of us are used to jaywalkers, right? So, we're supposed to cross at the light, right? But jaywalking is more convenient. It just is. But there is an inherent risk in jaywalking, right?

So, regulators, particularly on, say, college campuses, which I think most of us are used to, have said, you know what, we see people are just going to cut across here anyway, so because there is a power dynamic that skews in favor of the moving vehicle, let's put a crosswalk here and we expect the car, the bus, the whatever, to stop and let those people who are -- who would be jaywalking, in the first place, to cross. It doesn't take away the power of the bus or in this case the corporation, but it does say, you know, let's quote Spiderman, with great power comes great responsibility,

118

right?

So, the expectation is that when the bus or the whatever sees that person in the walk, they're going to stop. Does it happen all the time? No, but I think from a public policy perspective, it's putting in -- it's proactive measures to protect people from themselves and other people at times.

So, I think, you know, from the perspective of a regulatory agency that is a consumer protection agency would want to do something proactively when there are signs of issues or trouble. I think it's perhaps incumbent upon a consumer protection agency to do that.

MR. EGELMAN: So, I guess the issue of expectations versus preferences. So, we've done some studies and we have, you know, actual data to show that, you know, to some extent, this is an issue of learned helplessness. So, people are just sort of, you know, resigned to the fact that all of our date is out there, regardless of whether that is actually the case.

So, for instance, we did a study looking at single sign-on in websites. So, when you click the, you know, use your Facebook login to log into this website, those sites can then request some data from your Facebook profile. And, so, we wanted to see whether, you know, making that more apparent to users, so trying

119

to highlight what types of data might be collected by those websites from your Facebook profile, we expected that that would have an effect on whether people, you know, used this.

And we found that that was not the case. And when interviewing subjects, they said, oh, well they just assumed that Facebook is giving away all this data anyway, so I might as well get a benefit from it.

And so that's sort of the learned helplessness issue. And I'm not sure there's anything -- I think addressing that part of it is sort of putting the cart before the horse, because I think one of the issues we need to focus on are the expectations, you know, before they're formed.

Some of that might be doing a better job of public education with regard to online privacy, other pieces might come in the form of enforcement making that somewhat more subjective. So, yes, the law moves very slowly, technology moves quickly, but I don't think -- you know, I don't think the issue is making the policies around specific technologies, the issue here is narrowing or closing the information asymmetry.

So, you know, while we don't expect people to read every privacy policy that they encounter, we have some expectations about what a business might be doing,

120

as was pointed out.

So, you know, I don't expect, you know, regardless of what they say about, you know, what farm the beef came from, I don't expect it to have E. Coli in it, and that's not something that they need to explicitly, you know, provide notice for, it just should be expected that there's no E. Coli in this beef. I'll leave it at that.

MS. PHILLIPS: I would like to say, because we kind of ran out of time a little bit, but there is really a need for more transparency in the industry we're looking at, because often if you look at website claims, there will be quite a gap between what the contract actually says and what the website is encouraging consumers to believe when they are encouraging people to purchase tests.

And the other thing is that because the industry is so new and the technology is changing so fast and it is largely unregulated, a lot of tests that are coming to market haven't been validated. So, there is a question sometimes about what the consumer is actually buying, because the value to the company is the sequenced DNA, which they are using in ongoing research often. So, they are selling a product that gives them very personal data that they use for a long time and may not be destroying ever, potentially, and the consumer, an ordinary consumer doesn't necessarily have the expertise to understand all of the risks.

And the other thing is that genetic test results are complex in nature, a lot of general practitioners have trouble interpreting genetic test results, and there's been some studies that have shown that a lot of GPs wouldn't be comfortable with interpreting a DTC test result if a consumer brings it in, but at the moment, most of the time, it's being avowed as a consumer service.

And in terms of particular worrying terms in contracts, in some countries like the UK, the Office of Fair Trading, which is now being disbanded, but is the competition and markets authority, has a history of working with industry to try to discontinue certain unfair terms as well, and that's what I would say. There are some terms that really shouldn't be in the contract, because it's making it a very unfair and unbalanced bargain, and a lot of the use of these contracts is also eroding traditional contract law principles, really.

And I think people will often tend to engage with these much -- and I think your work shows that much more differently than they would with a paper contract.

So, for browsewrap, which is where the terms

are on a hyperlink, it's akin to walking into a shop and being bound by a sign on the wall that you didn't see and walking out again. And that's really problematic. Thank you.

MR. STEVENSON: I think I will add one. So, on someone's slide there was a mention of incorrect mental models, and a lot of us think through consumer knowledge and I think the educated consumer. So, no one would argue for an uninformed consumer as the goal, but I think I want to push back a little bit on that idea that our goal or the goal of some of this work is to correct mental models. I'm curious what you guys think.

So, someone smarter than me said something like all models are wrong, some are useful. And I think the consumers sometimes have very strong or inaccurate models that are helpful heuristics, and I'm curious if you guys in this work since you're all studying consumers' perceptions, if you sort of see those inaccuracies actually beneficial or -- not that we want this inaccurate model. Does that make sense?

MR. EGELMAN: Yeah, I think that was my slide. I think that one of the bigger problems with notice and choice is that I guess there's unreasonableness on both sides. So, there's, you know, unreasonable expectations on what the consumer should know to make an adequate

choice based on the notice given to them, so it's unreasonable to expect every consumer to read every privacy policy that they encounter.

At the same time, yes, people have really bad mental models about what's happening with their data when they go online. And I think, you know, maybe there needs to be some better outreach on that issue, but at the same time, I think, then, you know, that goes sort of to enforcement, which is instead of thinking, well, did the company give notice and was it incorrect, you know, and outright misleading, but is it also adding into that equation, is it reasonable to expect that someone could actually understand this? And I don't think that's currently being taken into account.

MS. SHOENBERGER: I will also answer that very briefly. As far as using heuristics, part of the cognitive advisor sort of mental model, first of all I disagree that heuristics, using them are faulty. They're almost always correct. I mean, we rely on them all day long in various capacities.

I think what we were arguing for were heuristics that were actually giving consumers -- they were backed by informed and concise and true information that the FTC approves. And, so, by using -- by promoting consumers and allowing them to see what these

heuristics mean, promoting the keys to those consumers, it gives them a meaningful choice.

A heuristic is no longer something that risk is as much of an issue for, more it is something that can genuinely grow on them as an indicator of safety.

MS. ANDERSON: Heather, can you guys talk a little bit more about how you see that kind of heuristic coming into place, how you'd develop it based on an average consumer's expectation, given that we heard a lot of the findings are consumer-dependent, and it kind of depends on your background and the experiences you've had, your age? How would you go about trying to develop something that would be generally applicable?

MS. SHOENBERGER: Right. We are in the preliminary stages of doing that, and this would be something that we would be testing in a lab, probably, a physiological lab, looking at people's automatic responses in addition to self-report, but that said, looking at heuristics and making cues that were in line with guidelines that we had come up with is based on consumer expectations of a different type of data collection.

So, we came at it, and this is an arguable point, from the type of data collection and how it's being used, and then entities could opt in, depending on how they were collecting and using that particular type of data. So, there would almost be a continuum of the address that you could or icons or cues that you could use.

And then in order to use that on your site, or within your materials, you would have to adopt the FTC's guidelines that went with that particular icon, and we would empirically test every single element of that.

So, the icon itself might be something that you would have to test to see if it was something that caught someone's eye, or someone noted that, I think it was Turow noted that people didn't notice some of the privacy policies. That's something we could correct with better web design and better icon design.

MS. ANDERSON: Okay. We have about 20 seconds left, so I would like to give you guys an opportunity to ask the last question.

MR. McQUINN: So, to follow up on what Darren said, with how privacy concerns kind of have morphed and changed over time, the ITIF actually released a report called The Privacy Panic Cycle that kind of tracks this, but I wanted to see, several different industries up here have kind of -- that you have studied have changed over time. Some of them are new like genetic testing, and Android is on its sixth release. I'm just wondering

if you could talk about how you've seen expectations change over time.

MS. PHILLIPS: Me?

MS. CHARBONNEAU: I think one of the things we have to acknowledge is that we're moving into the commercialization of health, and we're moving into the monetization of health data. And, so, what we've found is, what's existing at the moment is not industry-specific, and probably that would be our main recommendation, that as we move into this, whether it's direct-to-consumer genetic testing, whether it's the data that's coming from your Fitbit, whether it's the information you're putting onto sharing sites thinking that you're just meeting some folks out there who have the same complaints you do, and let me tell you what happened with the latest drug, this is now being monetized, and this is now in the corporate sphere, and our protections of the relationships and the data were created for the traditional health care system. And we haven't yet made the move over into looking at anything industry-specific as we move into this new form of commercializing health care and also monetizing health data.

MR. EGELMAN: So, one thing that we've looked at is trying to relatively weigh different user concerns

based on the technologies. And, so, I guess going to this issue of what policy is needed, and policy moves slowly and technology moves fast, while people do have very nuanced privacy preferences and expectations, at the same time, there are some things that people will think of as universally bad or universally unconcerning.

And, so, we did this study three or four years ago, we came up with a whole slew of risks related to smartphone usage, such as an app that uses data for X or shares data with certain parties, and we had people rank those. This past year, we did a follow-up study to that, where we came up with similar risks relating to wearable devices and IoT, and what we found is, you know, if you categorize those risks, the results pretty much held.

So, people are almost universally concerned with things that have financial impact, and almost universally unconcerned with things that are already public, such as demographic data that would be publicly observable. So, you know, an approximation of your age, for instance.

And, so, in that regard, I don't think we need -- we should expect regulation to be really specific to the technologies, but we can come up with regulation around the various risks that most people are concerned

with, and that should last longer than specific technologies.

MS. ANDERSON: Thank you all. Unfortunately, we are out of time, even though I feel like we've just started the conversation, but I do hope that the conversation continues after this conference. I look forward to reading more of your research as time goes on.

For all of you in the audience, as you heard this morning, our cafeteria will unfortunately not be available for lunch, however there are boxed lunches that are available for purchase just outside of the auditorium. They are only taking credit cards. You may eat your lunch in the overflow conference rooms that are across the hallway. Food is not permitted in this auditorium, neither are beverages, except for water.

And remember, if you leave the building, please take time to come back through security on your way in. If you don't have electronics with you when you come back through security, that screening will be faster. You can leave your electronics in this room, I have been told there is going to be a guard and that the room will be locked. So, you can do that to try to expedite your screening on the way back in. And thank you all for coming, we will see you back here at 1:00

129

p.m.

(Applause.)

(Whereupon, at 12:20 p.m., a lunch recess was taken.)

130

AFTERNOON SESSION

(1:12 p.m.)

MS. ANDERSON: Next we're pleased to have Commissioner Julie Brill provide a few remarks. Commissioner Brill has long been an advocate for consumers in securing their privacy and data, and we are thrilled to have her here today.

Commissioner Brill?

COMMISSIONER BRILL: So, thank you, Kristen, and thank you everybody who's here, as well as all of you out in TV land. Lunch may be over, but the feast of scholarship will continue. It's really my pleasure to open the afternoon with a few remarks about the research that's on display here at PrivacyCon, but before I do that, I really have to take a moment to do exactly what Chairwoman Ramirez did, and that is to thank the FTC staff who worked incredibly hard and incredibly well to pull this together. Kristen, Dan, Justin, Maneesha, I know I'm leaving out probably 25 people, but they all did a really, really wonderful job. So, can we just have a round of applause for these fabulous people.

(Applause.)

COMMISSIONER BRILL: Great job. Now, aside from the quality of projects and presentations, one thing has struck me about today's agenda. Instead of

131

being organized by discipline, you know, computer science here, economists over there, the day is organized around the key substantive issues in consumer privacy. This thoughtful organization is leading us towards something that we need for sound privacy policy development. A cross-disciplinary, richly detailed picture of consumers and how they make decisions about technology use.

Lurking behind the main regulatory approaches to privacy, whether it's notice and choice, informational self-determination, or a use space model, are questions about individual consumers, their goals in exercising their privacy rights, and their ability to do so in the environment around them.

At a high level, I think two principles should guide policy and practice. First, individuals have to be in the loop regarding decisions about what data is collected about them and how it is used; outside the privacy sphere, companies have excelled at helping consumers manage and use highly complex systems.

Now, we heard a little bit about Chipotle and the burritos. I actually think a much better analogy in this space would be cars. Cars are now computers on wheels, but we can all drive them, because companies have kept the complexity behind user interfaces that are

132

simple to use. I think companies can do the same for privacy, but building the right tools depends on understanding which decisions are important to individuals.

Second, I'm wary of solutions that depend too heavily on any one technical measure. Now, just as an example, it's a positive development that companies are offering more services that allow individuals to encrypt their communications, and these services are getting more user friendly, but their ease of use is limited to communications that stay within one particular service.

If you want to communicate between services, you may be forced to use tools that only a few select experts can really implement properly at this time. But these principles leave many questions open, and details unspecified. What data do consumers expect companies to collect from them? How do they expect companies to use this data? What do consumers understand about what actually happens to their data? Which aspects of data processing should be under consumers' control, and how effective are the tools that companies offer to consumers to exercise this control?

Answering these questions requires a three-dimensional approach. So, I was excited to hear this morning from researchers who are using structured surveys, qualitative interviews, and looking at human computer interactions to map out what consumers understand about the data practices of the services and devices they use.

Of course, it is just as important to understand more about what happens behind the scenes, outside the view of consumers. Data and device security are incredibly important to consumers, yet assessing security remains well beyond the capabilities of most consumers, including most of us, but not all of us in this room.

So, I'm thrilled to see researchers doing a deep dive on security vulnerabilities on specific Internet of Things devices, while others are analyzing data from thousands of vulnerability reports to better understand the kinds of incentives that will spur a virtuous cycle of discovery, reporting and patching.

Also beyond consumers' purview lies the big data analytics that have developed more quickly than have frameworks for specific concrete guidance on legal and ethical issues. Our Big Data report issued last week is intended as our first step towards providing such guidance. The report recommends that companies review their data sets and algorithms to determine whether they may be having unintended effects, such as

treating certain populations disparately, and in ways that may potentially violate the law.

Our report also recommends that companies bring a broad set of fairness and ethical considerations into their use of big data analytics. The presentations in the next segment of PrivacyCon address exactly those issues.

Finally, I want to give a shout out to the institutions that have helped produce the specific pieces of research that we're hearing about today. They are just as important as the research itself. Much of the research presented today comes from universities that have made substantial, long-term commitments to examining the relationships among law, technology and public policy.

In addition to generating new research that also contains policy insights, these universities help to train students to become leaders in their fields. Technology-focused centers and clinics have sprouted up at law schools all over the country in the last decade. They expose law students to technology, and probably just as importantly, to the way technologists think.

Departments, schools, and even entire campuses that make interdisciplinary work a core mission are doing much the same for students of computer science,

engineering, economics, public policy, and social sciences. Building these programs has not been easy. It's often easier to stick closer to traditional disciplinary lines.

So, let me offer a word of encouragement. PrivacyCon is just one example of the impact that scientists, lawyers and others can have when they're trained to do ground-breaking research as well as to identify and analyze public policy questions and issues. This combination of research capability and capacity for action also describes, just coincidentally, the design of the FTC itself.

So, naturally, we are a ready audience for research that sheds light on the challenges we confront in enforcement and policy development. And I hope that the institutions that many of our presenters call home will be lasting platforms for robust exchange of ideas with the public and private sectors for many years to come.

So, with that, let's hear what you have. Thank you very much.

(Applause.)

COMMISSIONER BRILL: And, Dan, are you going to -- Dan will introduce the next panelists. Thank you.

SESSION 3

BIG DATA AND ALGORITHMS

MR. SALSBURG: Thank you, Commissioner Brill.

Could the next panel come on up. So, welcome back to PrivacyCon. Our first session today really looked at what kind of data is being collected about consumers. Our second panel looked at what do consumers expect is happening with that data, and now, this session, we're going to look at what actually is happening with the data.

So, I am really pleased to have with me researchers who are going to present three outstanding research presentations, and we're then going to discuss them.

So, why don't we get things started with a presentation from Michael Tschantz and Anupam Datta. Michael is from Berkeley and Anupam from Carnegie Mellon, and they're going to lead things off with a presentation titled Automated Experiments on Ad Privacy Settings.

MR. TSCHANTZ: Thank you. I am Michael Tschantz, and this is going to be a joint presentation with Anupam Datta, and we're going to be talking about AdFisher, a system for looking at online trackers and determining what information they are using about people to select the ads they show to people.

There are two things I want you to take away from this topic. First, it is possible to do this with scientific rigor, despite not having access to the internals of the system; and second, we can find such interesting flows of information, but we can't figure out why they happened.

So, let's get started by motivating the problem. Here's a web page, it's The Times of India. I find it an interesting example because it has a lot of ads from Google on it. Here's two. Now, Google has little pieces of code across the Internet. In fact, this web page has two little pieces of code, and these pieces of code imparts back to Google about what other web pages you visited. Google can then select the ads it shows on The Times of India based upon this information.

And this is generally true of online behavior trackers, there's many trackers with many little pieces of code all over the place. There's a seemingly endless number of companies doing this kind of thing.

But it can be concerning. I mean, suppose, for example, you want to show a friend an ad -- a newspaper article and you see nothing but ads for antidepressants, which Google may have shown you in

certain circumstances. Now, Google understands that people have concerns like this, so they and other companies have provided things like the ad privacy settings.

Here is a screen shot of my ad privacy settings, it shows various information inferred about me. Google got my age correct, but got my gender wrong. Google also allows you to go in and edit this information, so if I cared, I could go in there and provide my correct gender. Google doesn't give us a whole lot of information about exactly how this thing is working, however.

So, what we have is a situation where we have our web browsing behavior going into an ad ecosystem at one end. You have various things like ad settings sitting in the middle providing sort of a window into how that ad ecosystem works, providing inferences they create, and allowing you to put edits in, and then we see advertisements coming out the other end.

But we would like to understand the flows of information in this system better than they currently make clear from their privacy policies and descriptions of how these systems work. And this is a difficult task because the system is opaque. We don't know what's going on in that ad ecosystem. Google and other online

behavioral trackers won't share its source code with us, we can't do the traditional forms of program analysis.

So, we designed AdFisher, a system that allows us to run experiments on these kinds of opaque ad ecosystems. Let me run through quickly how AdFisher works.

AdFisher creates a bunch of fresh Firefox browser instances which simulate users. So, these could be simulating people who browse various websites. It randomly assigns them to either a control or an experimental group. These two groups of simulated users would display different behaviors on the Internet. They then interact with the Internet in various ways, and we collect measurements about how advertisers change their behavior towards these simulated users.

These measurements go into a test of statistical significance, which reports whether there's a statistically significant systematic difference between the experimental and the control group. If so, we know that whatever information describes the difference between these two groups, and in how they behave towards the ad ecosystem, is information being used by the ad ecosystems to select ads.

So, this is our main contribution, is that we brought the rigor of experimental science to these sort

of online blackbox experiments in such a way that allows us to describe cause and effects which are equivalent to flows of information with the theorem we proofed. It does it with statistic co-significance, without making questionable assumptions about how Google operates.

This is important because Google is an extremely complex system, pretty much any assumption you make about how it's operated might not hold, or perhaps it holds it for one moment in time but not later when you're running your experiment. And we provide a high degree of automation.

So, now I'm going to give you an example of one of the findings we discovered with our system. In this experiment, what we did was we fired up our simulated users and we had half of them set the gender bit to be male and the other half to female on the Google ad settings page. We then had them all browse websites related to finding jobs. We then collected the ads shown to them at The Times of India, and we found a statistically significant difference in the ads shown to the male and female groups.

Now, this in and of itself isn't terribly surprising. We know that advertisers show different ads towards men and women, but what's concerning is the nature of this difference, something that AdFisher can also share with us.

What we found is that there were a series of ads from a career coaching service that was shown almost only to the male simulated users. In fact, the ratio was so large that it's in violation of the 80 percent rule often used in employment law to detect disparate impact.

That being said, we are not claiming that this is an instance of illegal disparate impact, because this is an ad for a career coaching service, it's not actually for a job. Nevertheless, we find this ad being shown predominantly to men to be concerning.

Now, this is just one of the findings. We have another interesting one involving substance abuse. We found that if you visited a website for a rehab center, all of the sudden Google would start showing you ads for that rehab center across the web, or at least at The Times of India. And this is concerning, since it's sort of like medical information being used for determining the ads you see on a newspaper's website.

So, I've used my time to explain some of the things we know. Anupam is going to now explain some of the open questions left open.

MR. DATTA: So, I'm very excited about where this research area is going in terms of developing

rigorous science and useful tools that are beginning to find effects in the ad ecosystem, and more generally in online personalization systems. At the same time, I am deeply concerned, also, about the findings themselves that we and others in this research area are beginning to develop, and we'll hear more from the two other speakers shortly about other findings.

These studies are beginning to get a lot of attention in the popular press, indicating that these concerns are shared much more broadly in the community. But there's much more to do in this space. There are questions like how widespread are instances of the discriminatory targeting or targeting that violates privacy expectations of perhaps contextual integrity or other notions. And then there is also the question of who is responsible?

So, I want to take a few minutes to highlight that these questions are incredibly nuanced to answer in the presence of the complexities of data analytics and other pieces of an ad ecosystem. So, I'm going to focus on this question of responsibility partly because following up on the conversations from the morning, I think that detection is an important step, but we can't just stop there, we have to go towards accountability, meaning assignment of responsibility and then

institution of corrective measures.

And this is going to involve collaboration between computer scientists and legal scholars, and probably policy changes. I want to focus only on the computer science piece of it for now, but we are working on the interaction between computer science and law, in part in collaboration with Deirdre Mulligan.

So, let me highlight some of the nuances of assigning responsibility with this concrete instance of discriminatory targeting that we found. So, just to remind you, this was an instance where typing job-related ads were being served in significantly higher numbers to simulated male users rather than female users.

So, what are some possibilities here of which entity could be responsible? So, one possibility is that Google, Google's programmers intentionally programmed their targeting system to be discriminatory in this way. We considered that to be highly unlikely, but nevertheless, it's not something we can rule out because we don't have enough visibility or access into the system that they use internally.

Another possibility is that the advertiser, the specific advertiser, in this case the Barrett Group that was advertising for this career coaching service

might have indicated when they submitted their bid for the ad that Google should show this ad more to male users than to female users, and Google may have honored that request.

A third possibility is that perhaps the Barrett Group indicated that the ad should be shown to high earners. In fact, in response from questions from the journalists at Pittsburgh Post-Gazette, the Barrett Group actually said that they were targeting users who are over the age of 45 and who earned more than $100,000 because they thought that would be an appropriate group to target for people who would want to go one level up and go after the 200K plus jobs.

Now, it could be that these high earners are much more strongly correlated with the stronger correlation of the male gender than with the female gender, and Google may have inferred that and then decided that they should send more impressions of this ad to male users than to female users.

Yet another possibility is that other advertisers might be targeting the female demographic more, and there is some evidence that female demographic is targeted more by advertisers, because they make more purchasing decisions, and those other ads may have come with higher bid amounts, which took up the slots for the female users, and the males just got these -- the ad from this particular service because they were the leftover, untargeted -- there were just more slots available for the male users.

Yet another possibility, and this would be the case of machine learning introducing discrimination, is that Google's internal systems may have observed that more male users are clicking on this particular ad than female users, and since machine learning systems learn from these kinds of observations, and they're trying to optimize for the clickthrough rate, they may have served -- started serving more impressions of this ad to the male users.

So, all of these are hypothetical scenarios because we don't have enough visibility into the system to determine which of or if any of these situations -- possible explanations is the real explanation. But I wanted to highlight this to explain the nuance of this problem, that this is a very complicated problem, and if you want to go towards making systems more accountable in this space, then the researchers will need additional access to the internals of the system.

So, being able to work not just from the outside, like we have in this work, and Roxana will talk about shortly in her work as well with the Sunlight

system, they have a similar model, but people on the inside, who have more access might, if they are interested in proactively testing their systems, that additional step will be very crucial towards proactive detection of violations, as well as of identifying responsibility.

So, that is something that I would urge this community to go towards, and it's an open call to technology companies who work with researchers like us to work on problems of this form that are socially important.

So, let me stop here with the summary that what this body of work, AdFisher, and a previous result that introduces the methodology, brings rigorous experimental design ideas to this research area, which lets us discover causal effects, for example that it's really the difference in gender that caused the difference in high-paying job-related ads being targeted, with statistical significance, so with confidence that it's not just a fluke observation, but it is really how the system is behaving.

And the third kind of contribution here is to bring automation that allows us to discover these kinds of effects at scale, and this combination was the first in our work, and then the community has grown and

developed it in many different dimensions.

So, we found evidence of gender-based discrimination that was one specific highlight, and the other highlight is how browsing helped related websites have a significant effect on targeting, in particular how substance abuse -- browsing substance abuse websites results in rehab ads being targeted.

And the two big open questions that I want us to open up for discussion, and these are active areas of research in this area, is how widespread is this discrimination, and how do we go from here in assigning responsibility. And as a corollary, I would like to emphasize that additional access to the internals of the systems, people with additional access to the internals of the system, working with such people is going to be highly crucial towards achieving these goals.

Thank you very much.

(Applause.)

MR. SALSBURG: Thank you, Anupam and Michael.

Now we're going to hear a presentation by Roxana Geambasu of Columbia University titled Sunlight: Fine-Grained Targeting Detection At Scale With Statistical Confidence.

MS. GEAMBASU: Hello, everyone. I am very happy to be here.

I will now tell you about some tools that we are building at Columbia to increase the web's transparency at large scale. To motivate our work, I'll start with an example that shows just how opaque today's web is. And you probably already know that Gmail uses emails in order to target ads, but do you know how the keywords or inferences drawn from these emails are being used to target you, specifically? I'll test to see how aware you are of how you're being targeted by showing you some examples that we got from one experiment.

We created this Gmail account, and populated it with a bunch of very simple single-topic emails, and I'm showing here on the left-hand side, five of those emails out of about 300 that we created. And on the -- you know, after that we retrieved ads that Gmail showed in this account. And I'm showing here on the right-hand side, ads, two ads out of about 20,000 that we got. So, this was a pretty large-scale experiment.

And what I want to do is to challenge you guys to tell me what each ad is targeting. So, for example, what does ad one target? Which of the emails? What do you think? Just quickly. Whatever comes to mind.

AUDIENCE MEMBER: Vacation.

MS. GEAMBASU: Vacation. Well, it actually turns out that ad one targets the pregnancy-related

149

email. You know, it's pretty hard to tell, right? Nothing in the ad really tells you anything about how it's actually targeted.

What about ad two? It's about a hotel. What does this one target?

AUDIENCE MEMBER: Homosexual.

MS. GEAMBASU: You got it right. That's exactly right, the homosexuality-related email. Again, it's still pretty hard to tell. And it's not just about targeting of ads on Gmail that's hard to discern. Everything is obscure on the web.

For example, you know, data brokers apparently are using -- you know, can tell when you're depressed and apparently sell this information. Or some credit companies, for example, are trying apparently now to use Facebook information in order to decide whether or not to give out a loan.

You know, you may have heard of these things from the media, just like I did, but do you know that whether these things are actually happening, to what degree, and how those things affect you? I'll bet, you know, not -- you know, people don't know too much about these things.

Welcome to the data-driven web. Many of the web services and third parties collect huge amounts of

150

information about us, every location, every site, every site that we visit, every click that we make and so on. And they leverage all of this information for all sorts of purposes. Some in line with our interests. For example, we all love our Netflix recommendations or Pandora recommendations, but other uses may not be so beneficial for us. And the big problem is that we have absolutely no visibility into what happens with our data in this huge, complex web data ecosystem.

Who has access to what data? For what purposes are they using it? Are the uses good or bad for us? You know, how do the uses affect us, really?

And it's not just the end users that don't know how to answer these questions, but society as a whole has a hard time answering these questions. And I believe, you know, the FTC, you know, as well, from my communications with them a little bit. And that's very dangerous, because, you know, obscurity and lack of oversight can lead to abuses, either intentional or not.

So, in my group at Columbia, we are developing these new kinds of tools which we call transparency infrastructures that shed light into this dark data-driven web. Our goal is to build really large-scale infrastructures that can go out there on the web and track the flow of information and reveal it, you

151

know, on one hand so that on one hand we can increase users' awareness of what happens with their data online, and on the other hand, increase in power privacy watchdogs, such as the Federal Trade Commission, to monitor what all these services are doing with users' data and keep them accountable for their actions.

And over the past several years, we've been building a number of these transparency infrastructures and we are continuing to do so now. And in this talk, I will tell you about just one of these, in the remaining time, just one of these infrastructures, the latest essentially public domain transparency infrastructure that we have built.

Before I do that, I want to acknowledge my students and collaborators, without whom obviously I wouldn't be standing here, you know, telling you about these systems.

So, what is Sunlight? Well, it's a generic and broadly applicable system that detects personal data use for the specific purpose of targeting and personalization. It detects which specific datum about a user, such as email searches, or visited websites are being used to target which service outputs, such as ads, recommendations or prices. The ads that I showed you at the beginning of the talk, their targeting, was

152

discovered by Sunlight.

Sunlight has three unique properties compared to everything else that exists. It is precise, scalable, and very broadly applicable. We've already tried it with great success to reveal targeting of Gmail ads out in arbitrary websites, recommendations on Amazon and YouTube, and prices on various travel websites. Not all of these experiments are actually in open domain yet.

And in all of these cases, Sunlight works with high precision, about 95 percent, as well as reasonable recall. How does it work? Well, the details are pretty complex, but at a high level, the idea is intuitive. Sunlight first started by correlating users' inputs, such as emails, with service outputs, like ads, by performing experiments on accounts with differentiated user inputs.

We can actually make the link from correlation to causation if we control how those inputs are placed in the accounts. Let me show you an example quickly, just to illustrate this process.

So, remember the ads that I showed you at the beginning of the talk? I'll show you how Sunlight might have detected their targeting, but let me first simplify the example a bit, so, you know, let's keep just three emails and one ad. And let's ditch the contents of the emails and ads.

So, what we have is a main account that consists of emails E1, E2 and E3. In these accounts is ad1. And what we want to do is to explain the targeting of ad1 on, you know, these -- one or a combination of these three emails.

What we'll do is three things: First, we will create a set of extra accounts, we call these shadow accounts, say three accounts, and populate them with different subsets of the emails. We do this randomly so that the placement of the emails into the accounts is random, is done randomly independent of any other variable.

Second, we collect ads from the shadow accounts and, you know, say, for example, in this example, that shadow accounts 2 and 3 observe ad1, but shadow account 1 doesn't. Third, we analyze these observations and yield the targeting prediction, and in this case, the most natural prediction that we would reach is that ad1 targets email 3 because the ad appears in all accounts with email 3, but never in accounts without email 3.

So, that's kind of how Sunlight works. And now there is an important distinction that I would like

to make, which is that the first two stages of this process populating shadow accounts with subsets of the emails and collecting ads from them, are service-specific, and pretty much in the mind in Sunlight, the frame of mind is pretty simple, simplistic, which adds to some browser automations.

The last stage, however, the analysis of these observations to yield the targeting prediction is intellectually challenging, and that's what Sunlight actually provides. Specifically, the example I showed you here is trivial. In reality, the scale is much larger, there are a lot more emails, you know, to consider, a lot more ads to explain, there is a lot more noise and so on.

So, all of these things make targeting prediction challenging, and Sunlight addresses these challenges by designing a rigorous methodology that leverages well-known methods from statistics to provide precise targeting predictions at scale.

And it does so, very importantly, and quite uniquely, in a service agnostic way so that we can reuse the analysis across many different services, like I said before.

So, let me now show you some of these challenges, just to exemplify the kinds of mechanisms

that we use to address them. Let's look at that simple example that we had with the three emails. Look at what we did. We used three shadow accounts in order to explain targeting on three emails. That's a pretty -- that's a lot of accounts, shadow accounts that we needed to create.

What if we were trying to explain targeting on a more realistic user account with thousands of emails, and potentially other online activity, too, that compounds together with the emails to produce the ads. We would have needed to create, you know, would we have needed to create all combinations over a number of accounts that are equal to all combinations of these inputs. That is a scaling challenge, a huge scaling challenge that I think is tremendously important.

And, you know, it turns out, in fact, that we don't need as many extra accounts, we can get away with a lot fewer, only logarithmic number depending on the number of inputs that we are trying to explain targeting on, and my theoretician collaborator, Augustin Chaintreau, proved this aspect theoretically and we evaluated it experimentally.

And the intuition is that if we can assume that an ad targets only a small subset of the many inputs that we have in a main account, then we can

leverage sparsity properties, the same concept that are underlying compressed sensing, which say that you don't need a whole lot of observations in order to reconstruct accurately, you know, a sparse signal.

For those of you who are familiar with machine learning, I guess, we use, you know, that's what sparse regressions can also give, and that's what we use in Sunlight. However, these particular methods don't, you know, guarantee -- only guarantee logarithmic correctness, do not guarantee the correctness of any individual prediction. And what we want is a correctness assessment of individual targeting associations so that we can trust the results that we get from Sunlight.

And for that, what we do is we use hypothesis testing, just like in AdFisher, a well-known method that provides a quantification of the statistical significance of each prediction.

So, you know, Sunlight puts all of these things and other mechanisms together in a particular architecture that provides the unique aspect, you know, properties that I mentioned before, genericity and predictability, scalability and precision. I won't go into the details of this.

And instead, what I'll do in the, you know,

157

remaining two minutes, is I'll tell you, you know, how Sunlight can be used. Specifically, Sunlight is a transparency infrastructure which provides some valuable primitives for targeting prediction, and on top of it, we and others built transparency tools for studying specific services. And we've built a bunch of these tools, and it's actually extremely convenient to build on top of Sunlight.

I will tell you about just one of these tools that we have built, which we call the Gmail Ad Observatory. It's an online service that enables studies of targeting of Gmail ads on users' inboxes. Here's how it works: A researcher or journalist supplies a set of emails on which they want to detect targeting. The Gmail Ad Observatory uses the sort of Gmail accounts in order to send emails to a separate set of Gmail accounts that become then the shadow accounts from which we extract the observations, or collect the ads and infer the targeting.

The Gmail Ad Observatory then collects the ads periodically and supplies them to Sunlight to get the further targeting. And what we did, so this is kind of the tool that we built, and what we did was we used this tool to run a 33-day study of ad targeting in Gmail. A pretty large-scale study. We got overall about 20

158

million impressions of ads and, you know, about 20,000 unique ads.

And what we found, we found a bunch of things, I'll show you just one result, which is a contradiction of one particular policy, or statement that Gmail makes in one of their FAQs. Specifically, what they say is that they do not target ads based on sensitive information, such as religion, sexual orientation, health, or sensitive financial categories. Well, guess what? We actually found examples, and a lot of examples, of ads that target each and every of these, you know, specific topics.

And I've already shown you, for example, the ad that targets the homosexual, you know, homosexuals. You know, let me show you another example from the health category specifically. You know, there are some senior-related, a lot, actually, of senior-assisted living ads that target Alzheimer's. Other ads, many ads actually that target Alzheimer's in general.

The results are an interesting ad that you can see there that targets the depression-related keywords, the "is he a cheater" ad for a cheating spouse or site, apparently. And there are a number of ads, as well, you know, in our example, that target the keyword cancer. I'm showing here just one of them. We found a number of

159

other ones.

MR. SALSBURG: You want to just take a sentence to wrap up.

MS. GEAMBASU: Yeah, that's right. So, to wrap it up, I've told you about our agenda of building generic and broadly applicable transparency tools which enable oversight at scale. These tools can be used to study targeting phenomena of various kinds, like ads targeting, for example, but not only, also price targeting, and I have actually a demo of that, if you guys would like to see it later.

Thank you.

MR. SALSBURG: Thank you, Roxana.

(Applause.)

MR. SALSBURG: Our final big data and algorithm research presentation will be from Daniel Hsu of Columbia University. It's titled Discovering Unwarranted Associations in Data-Driven Applications with the FairTest Testing Toolkit.

MR. HSU: Thanks, Dan.

Okay, so I'm going to tell you about another tool that we've been developing at Columbia, and also at EPFL and at Cornell Tech. A lot of collaborators on this project.

So, I should preface this by saying that I am

160

sort of an outsider in this community. I mostly do research in machine learning and on the algorithms that are used by Google, by Yahoo, by Microsoft, for doing the data analysis, for maybe doing the targeting, and so this is kind of a, you know, has a different perspective. I'm going to give sort of a different perspective on this problem, and but, you know, you're all well aware of the kind of issues that come up with a lot of these data-driven applications. So, maybe you probably heard of this study that was done about detecting sort of differences in prices from Staples' online store that are based on where you live, and this turned out to have some kind of correlation with the income of potential customers, and this was sort of an interesting finding, but with sort of more interesting from our perspective is that this was an unintended consequence of the sort of pricing mechanism that Staples was using.

And, so, here's another example of this kind of data-driven application that may have some kind of unintended consequences. This was in the case of Google's image tagging application, where if you were to upload photos onto Google's social network services, Google would try to automatically tag your images with various things, like say there's a car here, here are

161

your friends, and there was very unfortunately an incident where people found that some African American users, their pictures were being tagged as gorillas, and this was definitely not what Google was intending, right? This is not something that, you know, they wanted to happen.

So, you know, these are sort of problems that arise when you are creating these kind of data-driven applications. And what we want to argue in this work is that these are, you know, these are bugs, and sort of developers should be testing them, testing for these kinds of bugs and trying to debug them, to correct these issues, sort of in the same way that they would try to correct or do debugging to find potential functionality bugs, performance bugs and so on.

So, this is where our work comes in. We know that, you know, this is not an easy -- this is not sort of an easy problem to solve, these bugs are pretty nefarious, they're pretty hard to detect. So, what people might suggest is that, okay, you should take some preventative measures, but these, we know, also have a lot of limitations.

So, one thing you might suggest to do is, okay, maybe we should just completely ignore certain attributes about the data when we are designing these

162

data-driven applications so that we do not sort of create these kind of unwarranted associations in the service outputs. But we know this doesn't work, because there are always sort of other attributes that may be associated or correlated with the sort of sensitive attributes like income level or race. This, indeed, is what happened with the Staples pricing application where location just happened to be sort of correlated with income level. So, that might not work.

Another thing that you might try to do is to apply some kinds of sanity checks like to see if there's some kind of statistical parity in your outputs to make sure that if you look at, you know, at race, you're not sort of -- you're sort of at parity across the different race attributes, but we know, again, this is not -- this can be insufficient as well, just because there could be, you know, sort of smaller subpopulations, you know, with a particular attribute that end up having a strong association with a service output.

So, these are really hard problems for developers to solve. And, so, what we think we're trying to argue here is that developers really do need new tools to help them find these kinds of bugs. So, detecting these kinds of unwarranted associations is already a hard task for them to do.

163

So, this is where our research comes in. We've been developing this tool kit that we call FairTest, and as we call it, a testing suite for data-driven applications for developers to integrate into their tool chain to try to, you know, check your application, to do debugging, to run every time you compile to make sure that the application is working as they would want it to behave.

So, the way we kind of characterize it, or caricature a data-driven application in a data-driven application is somehow takes user data as inputs and there is some kind of output that the application provides, maybe the service prices, image tags, and recommendations, and so on, looking for some kind of function of these outputs.

And, so, maybe things like the user inputs, might be like the locations of the users and their profiles, whether they click on various things on the website, and like we said, the applications outputs are like the prices, the image tags.

So, FairTest comes in by something that you could strap onto your development tool chain and look at these kind of user inputs and the application outputs and try to check for various kinds of unwarranted association between the output and sort of protected

164

attributes that you wouldn't want to have some kind of strong association there.

And, so, FairTest is a tool for automatically doing this and it does this with some kind of data, and the hope is that it will at the end produce some kind of bug report that the developer will be able to look at.

So, what the developer would have to do is to sort of specify which of the sort of user input are the ones that are -- that we want to check for a strong association with. These are what we call the protected variables, protected attributes. These might be things like the gender or the race of the user, and then there are other many very likely, but many other attributes that are used by the application, and these are things that we're going to use to sort of try to define or to search -- to define various kinds of contexts in which there might be some kind of unwarranted association. And then the last one I will talk about in a little bit.

So, the goal of FairTest, again, is to define these kinds of context-specific associations between some kind of protected attributes and the application output, and then the bug reports is something that we'll apply some statistics or machine learning in order to produce something that the developer can understand in terms of what does -- which kind of context, what kinds

165

of associations were found by FairTest and to sort of rank them by the association trends or the statistical significance, so that is something that the developer can actually look at and understand.

So, let me say a little bit about how FairTest works. It's sort of, at its core, it's a machine learning algorithm or machine learning application, so FairTest itself is some kind of data-driven application. And the way that it works is that it starts by collecting or you start by providing it some kind of source of data. And here is where it's really important for the developer to really be -- to have some kind of source of data that is representative of a population of their user base, and this is where it's sort of difficult for maybe other parties to have access to this, but a developer presumably, you know, they're working -- they're like at Google or they're at Microsoft, so they have access to this kind of data already.

So, when they have this kind of data, they can really sort of check their application on the real user population to really discover the effects that have some meaning in terms of the actual users.

So, FairTest relies on this kind of data. And, so, what we'll do is something very similar to how

166

AdFisher and Sunlight operate, we will split this data into two parts, one we call the training data and the other part we call the test data. And we use the training data, so the part of the data set to sort of find these kinds of associations through some kind of clever machine learning algorithm. And then what if we find these kinds of sort of associations between protected attributes and physical application outputs, we'll use sort of remaining data, so sort of segregate it to actually validate these things and to measure their effect sizes and to check really are these things harming sort of a large segment of the population and is it very significant, and so on.

And this is where there's a lot of sort of technical machinery coming from machine learning. And then at the end, there are some, actually a lot of the work that's here is to make these kinds of findings sort of consumable by the application developer, so it's something that's interpretable and that they can actually use to help them maybe debug their application.

Let me give you an example. We actually applied this tool to a couple of sort of applications, some real applications that are sort of data-driven applications. So, one of them is the first one I wanted to tell you about is this sort of health care

167

application. This was actually sort of something that was produced by one of these machine learning contests or data science contests where some company, in this case it was Heritage Health Company, they ran this kind of competition where they tried to get -- you know, they provided some kind of data about patients going to hospitals and some sort of description of the patient records, you know, how many times they've been to the hospital before, you know, what were their symptoms, things like this. And the task was to use this kind of information to predict whether or not the or how many times the patient would come -- would visit the hospital in the next -- the following year. Sort of this kind of re-admission rate prediction.

So, what we did is we looked at the winning entry to this competition. It was a pretty good entry, sort of an application that was able to correctly predict with some pretty high accuracy, I think around 85 percent accuracy, whether or not the patient would be re-admitted into the hospital the following year.

So, this was the data-driven application. It takes these kind of inputs, age, gender, number of times they've been to the hospital and so on, and then it tries to predict whether they will be re-admitted to the hospital.

168

So, what did we find by applying FairTest here? What we found was that there really are some specific contexts where there's an association between the age of the patient and how badly the predictions -- how bad the predictions were. Sort of the rate of error rate or the size of the error in the prediction.

And, so, this was a -- this is sort of a contextual association that we discovered. It was not for the entire population, but sort of for some well-defined segment of the population. I think it was something like male patients who have been to the hospital at least or who have been to the ER at least like twice in the past year and so on.

But when we -- but within this subpopulation, there was a really strong effect, and a really strong association between age and the error in the prediction.

So, this is an interesting finding. We think that this is actually, you know, sort of important in a social sense, because this is something that could potentially really lead to actual harms, for instance, if this application was actually going to be used for insurance purposes, to do something to adjust your insurance premiums and so on.

So, these are associations that can really have some impact on the patients that they are -- or on

169

users that they are -- on users of the system.

I want to tell you about sort of another application, this is not a real application, but it's sort of a historical application, but it's something that would illustrate sort of a different capability of FairTest. So, this is a very well-known data set, sort of the application you can think of as the graduate school admissions application. What it does is it takes people who apply to Berkeley graduate school and decides whether to admit them or not.

So, this is a well-known data set from the '70s. If you don't know what happened with this data set, what happened was that they discovered out that there was this kind of gender bias on the admission rate at Berkeley, so men were being admitted at higher rates than women. And, so, indeed, FairTest can be used to discover this kind of association.

But what it can also do is it can try to explain where this association comes from. And, indeed, this is what this paper by Bickel, et al. In 1975 discovered was that, well, once you condition on which department the applicant wanted to get into, then the effect, either goes away or the impact maybe reverses, that women in specific departments would be admitted at higher rates than men.

170

So, this time what we wanted to do here is to illustrate here how FairTest can be used to sort of help a developer actually debug their system and try to explain what was going on, what was going wrong in their system. And maybe there's this other capability in FairTest for doing this, we call it the sort of providing some kind of explanatory variables, and this was really aimed to this a real system or a real tool for developers to use to debug their applications.

So, let me just make a few closing remarks. So, we also applied FairTest in a couple of other applications. You can read about it in our preprint, which is available on the web. So, I already mentioned this other feature of explanatory variables. There's this other sort of big issue out there in data analysis, which is that of adaptive data analysis, where you want to be able to reuse a data set many times. This is something that we're starting to look at and integrate into FairTest, and this is sort of open source software that can be used by developers right now.

So, just to sum up, really what we're trying to advocate here is that we really need to empower developers with sort of better statistical trainings, better statistical tools in order to make these kind of data-driven applications more fair, more socially

171

conscious and so on, and we think that they're just a good way to start here.

Thank you.

(Applause.)

MR. SALSBURG: So, joining me on the stage now are discussants James Cooper of George Mason University Law School and Deirdre Mulligan of UC Berkeley.

So, we've just heard three presentations about tools that are designed to shed some light on how data is collected from consumers, how this results in them receiving targeted ads, web content or result in discrimination.

So, let me turn first to James and Deirdre. What are the common themes you see running through these three presentations?

MS. MULLIGAN: So, I teach at the School of Information at Berkeley, and I spend -- one of the departments in my -- one of the programs in which I teach is a master's in data science, and we teach about privacy, we teach about security, right, these are people who are going to be doing data analytics, and one of the areas where we've been lacking, both methodologies and tools, is to deal with issues of fairness, right? How do we think about the biases in our data, how do we think about the biases in our

172

algorithms, and most importantly, I think what -- in particular, and I'm kind of most deeply engaged with Anupam and Michael's work, because we have some collaborative work that we're doing, how do we think about bias in systems where there are multiple inputs? And, so, it's very difficult to track an output back to a single actor's decisions.

And, so, as somebody who is working in that sort of program, one of the things that I think is most important about these tools is on the one hand, we have our last presentation, FairTest, which is actually trying to empower people who want to avoid, right -- all algorithms have biases, all data -- if you design an algorithm without a bias, it has no purpose in the world, right? Let's be clear, right? It has a bias, it's just that we want to avoid certain bad outcomes.

And the question about how we empower people who are designing systems to proactively avoid those outcomes is something that we need research on technical systems that people have called for, oh, we need access to the algorithm, we need access to the data as though if they can look at it, they're going to understand it.

And that just isn't the case in many instances, right? And, so, we actually need technical systems, we need the use of statistical machine learning techniques to police machine learning systems.

And this is particularly important because I think what all of them are highlighting and really focusing on is not -- I mean, we're concerned about intentional discrimination, but what I think many of us are worried about exploding is disparate impact, right? It's that nobody is intending for a bad thing to happen, but because what machine learning enables, what makes it different from what's gone before, is that the meaning of information emerges, right?

And, so, it turns out that these three pieces of data add up to some particular protected trait. And as machine learning techniques continue to uncover the way in which we have correlations that equate to these different things, we're in this -- we have this ongoing need to try to figure out proactively how to avoid those sort of problematic correlations.

So, I think they're all working on this shared problem from two different sides, right, that there's a long history of testing and we think about discrimination, housing discrimination, sending people out in the world. And, so, I think the AdFisher and the Sunlight are working on that side, right, can we test from the outside, and then I think that Daniel's work is really nice because it's saying, for the people who are

trying to do good, trying to avoid bad outcomes, can we empower them with tools that are based on the same sorts of statistical techniques that we need to police machine learning. So, I think they're really powerful in that way.

MR. SALSBURG: And, James, what do you see as the common themes?

MR. COOPER: I would agree with what Deirdre said. I mean, there are obviously, I think the common themes are pretty self evident. I mean, the co-authors are kind of back and forth on two papers, and all the papers kind of describe algorithms that do various similar work, and I think valuable work, as Deirdre pointed out.

So, yeah, I mean, I don't really have much to add beyond that.

MR. SALSBURG: So, Deirdre pointed out that in the real world, there are lots of inputs. I mean, consumer profile consists of it could be a million data points, or more. How can your tools account for that? When you're creating user profiles, is there any way to know what would really be happening to a consumer?

MS. GEAMBASU: So, this is a real problem, a very, very big problem. I would quote it as the biggest problem in web transparency work, to date, in my

opinion, which is to actually emulate real users with controlled experiments. All -- both of the -- both of AdFisher and Sunlight rely on controlled experiments with fake accounts that, you know, are assigned fake input sets or inputs.

And that results in some targeting, we are seeing, all of us, some targeting, but it's not necessarily true that it's realistic kind of targeting of the kind that real users would actually see.

We may be losing a lot of the targeting that real users see. We may actually have targeting that real users never see. And so on.

And I think that's a big, big problem. I think we need research in designing tools that leverage direct user -- data from real users in order to achieve some of the goals that we have in our system, transparency goals that we have in our systems.

That said, you know, I think, for example, I, because I've been working so much and focused and been invested so much in scalability, building scalable systems that can take many inputs, but millions, not the size that real users produce, certainly, you know, we've been focusing on that. And Sunlight does scale pretty well with respect to, you know, many, many -- trying many, many inputs and discovering effects on many of

these inputs. But there are big limitations still even there.

I also wanted to point out, because maybe the audience didn't realize. So, FairTest and Sunlight were actually, we're both collaborators on both, we just split the talks so that we wouldn't have to create, you know, to talk both about for each one of them.

MR. DATTA: Maybe one quick thing that I would add here is there are two ways to go about getting access to real data. So, one is to actually work with the technology companies who have that data. And, so, we have an ongoing collaboration now with Microsoft Research where we are actually beginning to get started with working with the internal data that they have about their users.

The other way to do it, or at least one other way to do it, is to try to get data from real users through crowdsourcing. So, there is the recent interesting paper from AT&T Research and collaborators elsewhere, which tried to do that, so the way they do their experiments is to just crowdsource it and collect data from users about their browsing profiles, and then compare it against the same user without the history. Some amount of the history. And then see if there's a differential treatment.

So, that's beginning to get towards experimental findings that have some amount of real user data.

MR. SALSBURG: James, do you have a question you want to throw out?

MR. COOPER: Well, sure. It's sort of a question and a comment. I'm an academic, so, of course, I'm going to spin my comment to say what I want to say and then ask you. So, one of the issues, and I guess this applies probably more to Michael and Anupam's paper, but I think to all the papers, is, you know, if we think about the transmission of your findings into policy, I think one of the touchstones of policy, at least in my view, should be harm.

So, I guess I think about your -- the finding of the job search ad, different for men, different for women, you know, and if you look at the statistics, let's assume that the data is there and there's a statistical difference and we can even say it's causal. Digging down deeper, you know, what's the real-world impact of that in the sense of, so, the click-through rates are maybe, what, one out of a thousand, if you're lucky, right? That's the average, right, one out of a thousand?

So, let's say one out of a thousand people who

visit this website, they would click on that, and these are people whose profiles have visited other job searching websites. You know, so my point, or my comment there would be to what extent -- they're not going to be limited, this isn't really necessarily, hey, I've gone to a thousand job websites, but now I've gone to The Times of India and I'm just going to take a job, I'm going to follow my career based on this ad that's served to me, I think that's probably not likely.

And then, I visited both those websites, I'm not sure, I'm sure you have, I don't know how many have, but the ones at head-hunter website, I'm not sure, it's got the nice banner, 200K plus, but it's a head-hunter. I don't think -- I'm not saying it's -- I'm sure it's legit, but I'm not suggesting the FTC look into it or anything, but compared to the other one, where the women were served more often, I think that was Jobs Near You.

And you go on that and the first page, click-down menu, they're not blue collar jobs, they're accountant, lawyer, bio. So, if you look at what would be the real-world impact, if you could imagine the two random people, the man and the woman. The woman who says, well, I didn't see the head-hunter ad and so I'm just going to go with jobs for me.

So, I look and I think about the real-world

impact. I mean, you did detect -- so, you did find, you know, a statistically significant difference between men and women, but at the end of the day, before we get into issues of harm, which I think should be a touchstone of any policy, especially here at the FTC, you know, where is -- do you need to find more? I mean, is there actually some sort of evidence of harm here?

MR. TSCHANTZ: Well, there's the saying amongst advertisers, which is I waste half of my budget, I just wish I knew which half. So, I really don't think anyone can look at any one ad and necessarily know what its entire impact is, but we do know that advertisers -- you don't see Coke ads on the TV because they expect you to stop watching the TV and run out and buy a Coke, right?

And these ads can be functioning in a similar way. It's about creating an impact upon people that lasts when they see something over and over again, or don't see something over and over again.

So, we're concerned about the women not being exposed to the encouragement to seek high-paying ads just as much as we're concerned about whether any one person clicks on that ad or not.

Now, I do think you raise an interesting point about the fact that this firm putting up this ad, you

know, I looked up its -- some customer reviews on it and it didn't really have the highest customer reviews. So, if we look at just the lack of perhaps women developing a business relationship with them, then it might be actually in their favor that they're not seeing this ad.

So, I don't know. You are correct, we can't pinpoint and measure the exact amount of harm, but we do know that men and women are being --

MR. COOPER: Or any harm. I would just kind of go that far.

MS. MULLIGAN: So, I think there are a few things to highlight. One, there was another example brought out about proximity to work, I don't remember whose paper it was in. Was it in Daniel's?

MS. GEAMBASU: Proximity to the location of a store.

MS. MULLIGAN: No, no, the proximity to work. It may have been in the FTC's Big Data report that just came out. And it's an example that's been used before, if you were looking for a potential employee pool, right, that you wanted to advertise to and you said, oh, well people who live closer tend to be better employees, and you might look and find out, well, that has a lot to do with income or whatever, or it could be a proxy for something else.

181

And we do, when we're thinking about employment, equal access to not just employment opportunities, but also we think about the advertising of those employment opportunities as something where we're concerned about racial disparities and gender disparities and how we're making information about opportunities available, as a legal matter we're concerned about that. So, let me finish, hold on.

And, so, kind of setting aside this particular example, right, which we agree is problematic for many reasons, and I think, you know, one of the most interesting things that this particular example of the head-hunter ad brought out, which Anupam noted, is that the most likely, we think, or at least a highly likely reason that men were seeing this more than women is that people were willing to pay more to sell women -- to show women advertisements for, you know, hair care products and other things, right?

And the point being that if you were a company and you were trying to use this to make information available about employment opportunities, you don't have complete control over who sees them full stop. Right? And when we're thinking about anything that repliers -- where you as an advertiser want to be attentive to who's getting access to your ads, because you're interested in

182

making sure that they are equally available to a population, define them whatever way you want, and you realize that there are other people whose bidding and decisions are interfering with your ability to know whether or not they're going equally to men and women, or they're going equally to people of different races, or whatever. You begin to say, wow, how do we think about causality, right, and how do we think about the relationship between bad outcomes and infrastructure, because it becomes an infrastructure issue.

Even if you are in the Staples example, Staples had access to their data, they were making decisions, they had access to lots of stuff, and they weren't seeking to have a particular bad outcome from your description, Daniel, yet they didn't do enough work or they didn't think through what was going to happen, right? So, again, it's about how do we create an infrastructure and tools.

MR. COOPER: Two things. To the extent -- my only point was using findings like this to inject into policy and the potential enforcement actions. You know, because that seems to be sort of the undercurrent in the papers, in at least two of them where, well, here's a Google privacy policy and, wait, my ad suggested there's tracking, which, you know, could lay the predicate.

183

So, my point is there seems to be a lack of harm. Now, in the Staples example, to say it's an unintended outcome, I think is completely unintended. I mean, that's just channel conflict mitigation. That's just the idea that I've got a brick-and-mortar store and I don't -- I mean, so that has nothing to do --

MS. MULLIGAN: Their intent wasn't to disempower people.

MR. COOPER: No, absolutely.

MS. MULLIGAN: And that's my point.

MR. COOPER: But I guess when you said they didn't intend the bad outcome, to them it's the correct outcome because it's the correct outcome based on that's the local pricing, I'm not going to undercut.

So, it has everything to do with competition and it has -- I mean, that has nothing to do with, wow -- you know, because there's actually -- you know, you think about, there's really no model that would set price discrimination and say, let's charge the poor people more than the rich people, and that's when -- you know, when I go to the movies and I hold up my George Mason ID, I try to cover the faculty part of it, right? And that's why, because they charge the students less, oh, you're faculty, sorry, you pay full price.

MR. DATTA: I have a brief, brief comment on

184

the question. So, for the job-related advertising example, I think this is where I was positioning this in that open problem of examining how widespread this phenomena is. This one particular ad is not enough for us to change how public policy works, but if -- and, you know, part of what Roxana is doing is building these infrastructures that allow examination of the entire Internet, possibly, a much more broader variety of sites, at scale, over many, many months, and you might -- if then she finds that there are many, many instances of these kinds of ads, maybe not this particular questionable ad, but from legitimate services that are showing up repeatedly in a differential treatment form, differential than the disparate impact, then the establishment of harm comment that you are saying is absolutely valid, that additional layer of analysis will not come from the kind of tools that we are building, that has to come from, you know, people like you and the regulatory agencies will look deeper, dig deeper into -- dig deeper into is this really a legitimate disparate impact, additional harm consideration.

So, absolutely on board with you on that, in addition to the other comments.

MS. GEAMBASU: So, I just wanted to add something very, very brief, I completely agree with

Anupam. What I wanted to note is that this research is at the beginning. This kind of research into building infrastructures that can, you know, tell what's happening is at the beginning. And as a result, we know very little.

We have a bunch of examples, right? That's pretty much what we have. I have great hope for this field, especially because more and more people are coming into it, that we'll develop the kind of infrastructures that we will need in order to actually make impact on -- you know, in the legal domain. But right now, you know, I think we know too little in order to do that.

MR. DATTA: Short of having proof of existence is useful as a starting point. We don't have evidence that it's widespread. That's ongoing work.

MR. SALSBURG: I guess the good company that wants to ensure it's not discriminating can use Daniel's tool, and the others can get caught by the two other tools.

MS. GEAMBASU: That's exactly the way we were thinking and why we've been developing both from the exterior, right, for the other thing, and tools for the developers to actually help them, you know, figure out what to do when the pressure is on from the exterior.

MR. SALSBURG: So, we have 50 seconds, that gives each of you about 10 seconds to do a final thought.

MR. TSCHANTZ: Just to complete my thought, you know, we've decided in employment men and women should be treated the same. So, to me the fact that they're not being treated the same is in and of itself a harm. Maybe it's not to you, but that's my opinion.

MR. DATTA: So, I would say that we need a complete accountability tool chain that goes from detection to responsibility assignment to correction mechanisms, and there is an emerging body of work on each of these pieces of the puzzle. Our focus here has primarily been on detection. There was a small amount of explanations in the last talk, but there is a huge set of open questions related to responsibility assignment and corrective measures.

MS. GEAMBASU: It's okay.

MR. SALSBURG: Well, with that, we will wrap up this session. So, thank you all so much. The cafeteria will be open during this break, if you want to go out and stand in a long, long line. So, we will be back in 10 minutes.

(Applause.)

(Whereupon, there was a recess in the

proceedings.)

SESSION 4

ECONOMICS OF PRIVACY AND SECURITY

MR. MORIARTY: Welcome back, everyone. My name is Kevin Moriarty, I'm with the Federal Trade Commission and this is session 4 on the economics of privacy and security.

First we have Jens Grossklags from Penn State University presenting An Empirical Study of Web Vulnerability Discovery Ecosystems.

MR. GROSSKLAGS: So, welcome to the first of two talks in this session that are actually about security. This is joint work with Mingyi Zhao and Peng Liu at the College of Information Science and Technology at Penn State University.

So, my talk is about the topic of bug bounties and vulnerability discovery that is mostly conducted by external researchers that are often called white hats. In 1995, the first bug bounty program was founded by Netscape that invited external security researchers to scrutinize its services. Since then, we had a number of other company-sponsored programs emerging that were run in an independent fashion. However, more recently we actually observed the emergence of so-called bug bounty platforms, and two of them are HackerOne and Wooyun, which are the focus of our study.

189

Wooyun was founded in 2010 and is mostly focused on the Chinese market. HackerOne operates in Europe and in the United States mostly and was founded in 2013.

So, the motivation for our study is to better understand how these bug vulnerability ecosystems actually operate and whether they make a significant contribution to web security. We also want to provide useful data for the policy heads, for example on the limits of vulnerability research and practice.

Our approach is to do an in-depth empirical study of these two ecosystems, and in our paper we take a very broad approach in the sense that we try to understand how organizations, white-hats, black-hats, the public interact on these third-party vulnerability platforms, but in the presentation, I will mostly focus on the perspective of companies and organizations.

So, the two programs that we look at have a couple of common aspects, mostly that's very popular, a lot of white-hats are interacting on them, and also a lot of vulnerability reports are made, but otherwise there are a couple of important differences. The first one is that HackerOne is “organization initiated” in the sense that these companies ask HackerOne to run a particular program for them versus on Weeyun, hackers

190

can actually submit -- any kind of hacker can submit any type of vulnerability about any website to the platform.

So, this is a very fundamental difference. There are some differences with respect to the bounties. Otherwise, also Wooyun is different in the sense that it doesn't delay full disclosure policies, so irrespective of the wishes of the company, after 45 days, the whole technical details of the discovered vulnerability will be communicated to the public.

So, there are some differences in the type of data we have about the platforms, so we cannot always directly contrast and compare the two, but what we can do is in five broad categories provide somewhat of a comparison on how these platforms actually operate.

The first one is participation. What we observe here is that on HackerOne, the number of public programs that are run is limited to about 100, and all of those are IT companies. In contrast, on Wooyun, we see a much broader portfolio of companies that are more or less coerced to participate on the platform. And interestingly, we see here a lot of organizations that typically are not known to run bounty programs by themselves, like government institutions, education institutions and also financial institutions.

So, the first takeaway that we have is that

191

the white-hat-initiated model allows for much broader participation and which may be good in the sense of web security. The more limited participation model or platform such as HackerOne, of course, raises then the question of how these platforms can actually encourage more companies to participate.

A second issue that we want to then explore is the quality of the submissions of what we observe here, in particular on the platform of Wooyun is that we have a very broad range of types of vulnerabilities that are submitted, and in 44 percent of these cases, these are actually classified as high-severity vulnerabilities.

On HackerOne, this is a little bit harder to determine from publicly available data; however, if you actually heard rumors of bounty amounts that are paid through the white-hat hackers, and also look into the policy statements, by combining these two data points, we can actually also then infer how many vulnerabilities are of high or medium severity, which is plotted on the slide.

So, here, we can also conclude that across these two programs, white-hats actually make significant contributions to the security of these websites by contributing high-severity vulnerabilities.

So, but more importantly speaking, the

192

white-hat model, white-hat-initiated model that is in Wooyun seems to harvest more of these vulnerabilities in an efficient fashion. Now, the question arises how well actually these different platforms and in particular the companies associated with them can actually respond to the submitted vulnerabilities. And here we see some interesting differences.

When we look at Wooyun, we can actually see that in particular, those very popular companies as measured by a measure of the Alexa rank, we see that most of them can actually adequately respond to the submitted vulnerabilities and handle them. In contrast, less popular and smaller websites very often are actually not capable to do so. So, in fact, about 25 percent of the submitted vulnerabilities remain entirely unhandled by the organizations to which they are targeted.

On HackerOne, in contrast, since these are company-initiated programs, we see a very quick response time. Within four and a half hours, we see the first response to submitted vulnerabilities, and most of them are actually handled then within 30 days.

So, an interesting takeaway that isn't on this white-hat-initiated model on these platforms, we see that a lot of companies that are coerced to participate are actually also not prepared, which is something that we have to take in consideration, versus -- and that, of course, then raises the question of balance. Well, should we actually coerce these companies to participate, is it a reasonable activity that we should be engaged in.

The next question that I'm approaching is, of course, of world interest, what impact do actually these kind of bounties have? Here is the first overview that we are seeing. So, what we are seeing here is a subsection or subassembly of companies participating on HackerOne. We see on the left side that some companies are actually not paying any bounties at all, versus others pay actually pretty substantial bounties for a submitted vulnerability.

On average, this doesn't really help us yet to determine what actually the really significant impact is, and for that purpose, we actually conducted a regression analysis in which the dependent variable is a number of vulnerabilities submitted, and the independent variables are the average bounty paid by a particular program, the popularity of the program, and a measure of the overall activity of the white-hats on the platform in the particular period.

So, what we are seeing here is first I want to

highlight the top part of the table, is that about a $100 increase in the expected bounty pay towards white-hat researchers, we see about three more vulnerabilities reported to the programs. What we also see is that programs that are more popular are also receiving more vulnerability reports, and that, of course, has two factors.

One, more popular websites, of course, receive more attention, but often they are also more complex, they offer more services to the users, so they likely have a larger attack surface in the sense for white-hat researchers to find potential vulnerabilities.

So, the takeaway here is that white-hats do not necessarily always focus on monetary compensation. In fact, what we observed is that 20 percent of all contributions on HackerOne actually go to those services and programs that actually do not pay any bounties at all. So, pay nothing actually serves as a potentially viable approach.

In contrast, what we also observe is, well, a higher bounty amount at the end of the day is still associated also with a larger number of vulnerabilities that are submitted by the white-hat researchers.

So, which brings me then to the last question, the one about security improvements. So, what do we

actually get out of it? And in order to assess it, while we do not have an inside look into the organizations, we are using the trend of vulnerabilities submitted over time. So, the argument here is that if you have a declining trend of vulnerabilities, everything else, keeping moderately equal, then we would argue that perhaps at this particular website, the security is overall improving.

So, when we take a first look at the data and we see that it's actually rather spiky, so it's not immediately apparent by looking at these graphical depictions what kind of trends are emerging; however, one thing that we see in the top three graphs for HackerOne is that this seemingly initial spike, where once a program is opened, a lot of vulnerability researchers are submitting for vulnerabilities that they have stockpiled or that they have been essentially energized by the opening of the program to do immediately a lot of research that led to additional submissions.

Wooyun is a bit more noisy. So, in order to get a better understanding of how overall these trends shape out, we conducted a statistical test that's called the LaPlace Trend Test, and we focused here on programs that have a certain amount of minimum activity that were

running for at least four months, had at least 50 vulnerability reports submitted to them. And what we see here is actually that two contrasting trends.

So, for HackerOne, we actually observed that over time for the majority of the programs, we see a decreasing trend of vulnerability reports. In contrast for Wooyun, which is a white-hat initiated, this coerced kind of company participation model, we see exactly the opposite. Mostly an increase in the vulnerabilities reported.

So, it would be reasonable that, and we could argue, well, despite monetary or perhaps because monetary incentives are in place, we actually see nonetheless there's fewer vulnerabilities on HackerOne. So, despite incentives, fewer vulnerabilities, we argue that this is indicative of actually improved web security practices at these participating companies. And keep in mind, again, that these participating companies are mostly IT companies in the case of the public HackerOne programs.

We also see this initial spike, which from a web security point of view might be really welcomed news, if, indeed, it's indicative that a lot of the stockpiled vulnerabilities are actually removed from the knowledge of white-hat and potentially black-hat

197

hackers.

We see an opposing trend for Wooyun programs, and our interpretation of that is that, well, this likely has to do something with the lack of preparedness of these organizations when it comes to receiving these vulnerability reports.

So, for example, they may not have a well-developed, secure software developing life cycle, good integration between the security team and these external security developers, and many other factors might actually play a role here.

Which already brings me to the last point. So, we believe that it's instructive to conduct a really in-depth analysis of these programs to better understand what contribution they can actually make to its overall web security and practice. And it's definitely helpful that these two programs provided this as public data which we can study in detail. There are many more results which we actually have in the paper, in particular pertaining to how white-hats actually behave.

For example, we can showcase in our paper how white-hats learn from one another by investigating the reports of their fellow hackers. We can also study what kind of discovery participants have in place. For example, are they focusing on specific programs, or are

198

they applying the same kind of technique across very different websites.

So, there are a lot of interesting additional results, if you haven't already accumulated our papers, I encourage you to take a look at them.

In total, I believe that the jury is still out about which of these two participation models, the white-hat-initiated model or the company-initiated model are really giving us the best advantages.

On the first glance, it seems that the white-hat-initiated model really has strong benefits in terms of participation, so we see many more white-hats, many more organizations that are involved in these kind of ecosystems, but on the other hand, a lot of these participating organizations are not very well prepared when it comes to receiving these kind of vulnerability reports, and actually then improving also the security on their websites.

So, there are various kind of pros and cons that we can observe. One issue is clear is we can jump start or further engage in the discussion what kind of contributions overall these bounty programs make to the security of these websites. Our initial assessment is positive, but I think we can go into further detail during the discussion, and that brings me to the end of

199

my talk.

Thank you very much.

(Applause.)

MR. MORIARTY: Thank you, Jens.

Next up we have Veronica Marotta and Alessandro Acquisti from the Carnegie Mellon University.

MR. ACQUISTI: Thank you and good afternoon. This is joint work with Veronica, Kaifu Zhang and myself. If you know some of our previous work, you may know that often we use behavioral economics to try and understand how people make decisions about personal information. The study represented today is actually about traditional microeconomics, and it is about understanding the allocative and welfare impact of targeted advertising.

However, there is still a behavioral angle, at least in the motivations behind our work. In behavioral decision research, it is very well known that how you frame a certain problem influences the way people will think about this problem, and we make decisions about it.

And currently, we believe not only the age of the data, but also under a very powerful frame, the frame that personal data is the new oil, and we are all going to benefit, perhaps in equal parts, from the

200

collection and sharing and analysis of our personal information.

More specifically, there are a number of frames which are quite common in the public debate over privacy. For instance, personal information is the lifeblood of the Internet. So, the increasing -- the increasingly sophisticated collection of data is necessary for us to have free services online. Or loss of privacy is the price to pay to excel the benefits of the data, or sharing personal information is an economic win-win, which benefits equally data holders and data subjects.

Well, in our broader research agenda, we are interested in investigating all of these frames to see how actual empirical evidence there is supporting them or not supporting them.

The paper we are presenting today tackles the last frame, and more specifically, relates to the impact that targeted advertising has on the circles of different stakeholders. Consumers, advertising firms, and intermediaries, the ad networks. And Veronica will guide you through the model.

MS. MAROTTA: So, thank you, Alessandro.

So, the specific and integral question we are interested in addressing is to what extent the more and more precise information about consumers is due to increasing in their welfare, what Alessandro just referred to as the economic win-win, versus a change of allocation of benefits among the different stakeholders, including companies, consumers and online intermediaries and platforms.

Now, in order to address this question, we rely on economic modeling to be the multi-stage three-players model of online targeted advertising that compare different scenarios that differ in the type and the amount of consumers' information that is available to the different players during the targeting process.

Now, specifically differently from previous work, we account for the important role played by the intermediary in the advertising ecosystem, and we focus on a specific mechanism of realtime bidding. Realtime bidding is a technology introduced to facilitate the location of program modeling advertisements online. Let me explain quickly how it works.

We have different players involved. On one side, we have publishers, namely websites that wish to sell advertisement space that is available on their sites. On the other side, we have advertisers, companies that wish to advertise their products online. But those two players do not need to communicate

directly, they can rely on the intermediary, the ad exchange, that facilitates the location of advertisements and the targeting process.

So, the mechanism works as follows: When a user arrives to a publisher's site, a signal is sent to the ad exchange that is subsequently broadcasted along with user data, maybe IP address, user cookies, geolocation, to interested advertisers, and hold an auction for the location of the advertisement.

So, on the basis of the information that the advertiser receives, they form a bid. So, how much are you willing to pay to show an advertisement to that user. And submit that bid to the ad exchange. Commonly, the ad exchange uses second-price auctions. This means that the highest bidder wins the auction, but he pays the second-highest bid. So, once the bid -- the winner is determined, it is allowed to show the advertisement to the user.

Now, on the basis of this mechanism, we built a model that focuses on the interaction among three main players: The advertisers, the intermediary and the consumers. We assume that advertisers are profit-maximizing agents, they want to advertise their product to consumers that will like and therefore buy their product. Nevertheless, they cannot target

consumers directly. They need to rely on an intermediary that facilitates the location of the advertisements.

We assume that the intermediary itself is a profit-maximizing agent that receives a payment every time he holds the auction for the advertisement's location. Finally, consumers have product preferences, but they need to know which seller is selling which products. So, in this sense, advertising plays informative role to consumers.

Further, we assume that consumer can be categorized by two categories of information: Horizontal information, capturing consumers' preferences and tastes; and vertical information, capturing differences in purchase power.

Now, these three players interact in our model in this way: At a given point in time, a consumer is online and he may be categorized by these two pieces of information, horizontal and vertical. The ad exchange receives the signal about a consumer, observes the information, and holds an auction for the location of an advertisement to that consumer.

On the basis of the information that it receives, advertisers form their bid. The auction is run, the winner is determined, and it is allowed to show

the advertisement to the consumer. The consumer sees the advertisement and makes his purchase decision.

Now, it should be noted here that the outcome of this process crucially depends on the information that is available during the targeting process. So, therefore, realize how the outcome for consumers, advertisers and intermediary changes when different types and amounts of consumers' information are available.

We considered specifically four cases. A case where only the horizontal information is available, a case where only the vertical information is available, a case where both pieces of information are available, and a benchmark case where no information about consumers is available, so an extreme full privacy case.

For each of these cases, we derive what's the firm's best strategy, and therefore what's the firm's profit; what's the intermediary payment received from the location of the advertisements; and what's the consumer's choice and surplus.

Now, in the interest of time, I will not go through the mathematics of the model, but I would like to show you interesting results that we obtained by simulating the model. So, what we do, we run competition simulations to analyze how the outcome in terms of consumer's surplus, intermediary's profit and advertiser's profit changes when in the four different informational scenarios.

Let me start from the consumers. Now, the graph that you see here, the X-axis captures how heterogeneous consumers are in their preferences, while the Y-axis captures how heterogeneous consumers are in their purchase power.

Now, important to note, low values means high heterogeneity; high values means high homogeneity. Now, the different colors corresponds to one of the different informational scenarios that we considered. Specifically, each region captures under which scenarios the consumers are better off.

So, we have two predominant colors here. The green region captures all the combinations of the model parameters for which consumers are better off when only the horizontal information is available during the targeting process. So, what's the intuition there? In their region, consumers are more heterogeneous in their product preferences, therefore revealing the horizontal information actually ensures that consumers see the advertisements for the products they like the most. So, there is a better matching between consumers and companies.

The yellow region, instead, captures all the combinations of model parameters under which the consumers are better off when none information about them is revealed. So, in that region, consumers tend to be more homogenous, so brands don't matter as much. So, the targeting is not as available to consumers.

Now, we can construct a similar graph for the intermediary's profit. Again, we have two main regions. The yellow region, again, is the combination of model parameters for which the intermediary's profit now is highest when none information is revealed about the consumer. So, we said in that region consumers tend to be more homogenous.

So, what happens is that if advertiser had that information, they will tend to bid lower to show the advertisement, lower in the intermediary's profit, but if the information is not revealed, then the advertisers have to bid an expectation, so they may overbid, increasing the intermediary's profit.

The red region, instead, is the combination of model parameters for which the intermediary's profit is highest when the vertical information about the consumers is available. In that region consumers are more heterogenous, and so revealing actually the vertical information during the targeting process

intensifies the competition among the bidders. They may tend to bid more aggressively.

So, if we put together these two pictures, we see that we have situations in which the interests of these two players are actually aligned through the yellow region, but there are also situations in which they have contrasting interests. So, we may think a situation of an intermediary that may have power over the information about a consumer, and may decide to act strategically, either by revealing the wrong type of information, say green versus red region, or revealing too much information when instead consumers would have been better off with less information being revealed.

Now, finally, we can use these simulations to understand and analyze how the allocation of the benefits among the different players changes under the four scenarios, so we can construct a pie chart like the one that we are seeing now for an information case, where we see the percentage of the value generated with auditing process that is captured by each player.

So, we can have a pie chart for each scenario, and what these pie charts show is actually part and very similar to what we just discussed. Consumers in blue tend to be better off either in the no informational case or in the horizontal information case, while the

intermediary, in red, seems to capture a decent amount of the benefits in all the cases, with the vertical information one being by far the best case.

For firms instead individually, it's always better off to have at least some of the information about the consumers with the complete information case being in this case the best scenario.

So, if you want to summarize those findings, we find that consumers are generally better off either when a specific type of information about them are available, or, in general, when less information are available. And that there exist situations where the interest of the players, say the intermediary and consumers, may be misaligned, and therefore a strategic intermediary may choose to selectively share consumer data in order to maximize its profits.

So, I will leave Alessandro to some final remarks.

MR. ACQUISTI: Thank you.

So, there are a number of extensions we are planning or working on, probably the most important is the empirical validation. In fact, if representatives of ad networks are in the room or following from via webcast, if you want to disprove or prove our results, we would love to work with you.

Now, going back to the broader picture from where we started. On the left, you have the three frames I started from, and I claim that they have something in common, which is very little empirical validation. So, I am not claiming that they're necessarily wrong, I am claiming that we really don't know how true they are.

So, on the right, instead, I have three broad research questions that I believe are critical to really understand to what extent they advertise the new oil and to what extent the benefits of this new oil are allocated, fairly or not, to the different stakeholders.

How is the surplus generated by data allocated? If we use privacy technologies to find a nice combination of protection of data and sharing of data, are there some costs, and if so, who is suffering those costs? Individual consumers, because they may get less targeted advertising; society as a whole because maybe the next medical researcher investigating cancer cannot find a cure because he or she doesn't have enough data; or it's just the issue of decreasing the rent exacted by oligopolies that are in the industry. Very different scenarios, and therefore also very different policy conclusions.

And, finally, under what conditions consumers

still benefits from trades in their data and in what conditions they do not, because I believe that the answer is not binary, it's not always good or always bad, it is very match context-dependent.

Now, this is a work in progress, in fact, this work in our agenda. However, if you are interested in the ground material in this area, and by this area I mean the combined use of privacy, you can find on SSRN a semifinal version of a paper that Curtis Taylor, Liad Wagmanz and myself had accepted and it is forthcoming in the Journal of Economic Literature. It's The Economics of Privacy and we will leave with you this.

Thank you very much for your attention.

(Applause.)

MR. MORIARTY: Thank you, Veronica and Alessandro. Next is Catherine Tucker from MIT to present Privacy Protection, Personalized Medicine and Genetic Testing.

MS. TUCKER: Okay, thank you very much for having me. So, I'm Catherine Tucker and I am an economist who studies the economic effects of different types of privacy regulation using real-life data. And what I'm going to be presenting today is joint work from Amalia Miller where we investigate how different forms of privacy protections affect consumer takeup of genetic

testing.

And because I know that a lot of you are here to think about advertising and more mainstream issues, I want to make a pitch for why this is interesting, before you all go to your electronics.

The first reason, so why we think it's interesting is that, first of all, this is a technology with a huge upside, as I'll get to later. Secondly, it's also a technology where I think even the most cynical person about privacy would say there are potential privacy consequences of this data being created. Sometimes when you're thinking about targeted advertising, it's hard to actually articulate the privacy form, which is why we actually think about health and financial examples, but when you think about genetic data, it's not hard to come up with examples of harm.

So, for example, I took a 23andMe test. I will share with you, I found out, rather depressingly, that I've got a few times more than average chance of getting macular degeneration later in life, that means I won't be able to see too well.

Now, the reason I feel confident announcing it in this audience is because ultimately I have tenure at MIT, I probably have the least potential consequences of

anyone in the world of releasing that kind of data because I have a job and I have health insurance, but there are potential -- you do not have to go far to think of potential negative consequences of that data, and as the previous presentation on genetic privacy articulated, I think very well, there are also issues to do with identifiability, the fact that this data is persistent, and the fact that potentially this data has spillovers to family members. So, it's really quite a lot of privacy consequences.

The other reason I think this paper that we're setting is useful is simply because there has been a lot of experimentation about different kinds of regulation which allows us to have more of a horse race than we usually do when trying to evaluate how well privacy protections work.

Now, I said there was an upside to this data, I just talked about the downside to it being created, but there's a huge upside. And the upside is the promise of personalized medicine. And the typical statement made in favor of the personalized medicine is that for the average drug, based on your genetic makeup, it won't work 25 percent of the time.

So, we can imagine if we actually have genetic data, we will be able to identify effective drugs, and

213

save many drugs and save money at the same time.

Now, as well for these claims, I often find it useful to sort of bring it to life with a very personal example, which is the example of Angelina Jolie and how genetic tests and the actions she took from it based on it.

So, Angelina Jolie did genetic testing, she found out that she unfortunately had a mutation in her genes, which meant that she was likely to get both breast and ovarian cancer; as a result had a double mastectomy and a hysterectomy.

Now, this is obviously a strident and decisive medical action, but in principle it's going to reduce her chances of getting cancer by 70 percent. So, this is the kind of data which actually leads to extreme forms of action in a medical sense, but there's a huge upside in health outcomes in terms of it being created.

Now, what we're going to do in the study is look at state laws' experimentation with different types of privacy regulation from 2000 to 2010. And what's nice about this variation is you always worry in any empirical study where the variation is coming from, why are the states actually experimenting in this way as an underlying reason.

From what we can see, it was pretty random,

214

driven by individual state senators who got a bee in their bonnet. And what we found also is nice, is that they're experimenting with many different types of privacy regulation, and we're going to bucket them in the study into three buckets, which are informed consent, regulating data use, and establishing property rights.

And I want to -- in the past, what I've done is I've said, well, you know, the great thing about this is it actually emulates different countries' approaches to doing privacy regulation. If you sort of think EU and OECD approaches, more associative informed consent, maybe the U.S., you can say we've thought about restricting data use, and then there's sort of this economist's dream of establishing property rights.

Now, I say that in the past, the reason I no longer push it is I mentioned this once when I was giving this talk in Paris, and this person from the Ministry of Culture in France stood up and said, how dare you say that, in France we regulate privacy in every single way you could possibly imagine, so it's not just one, but in general, what's nice about it is at least we've got a horse race for different ways we might think about regulating privacy.

Now, we're going to have data on people's

215

decisions to get to use genetic tests. We're lucky we have a national sample that was done every five years in the peer review study and they're going to be asking 30,000 people about whether or not they had a genetic test in each sample.

Now, it's a great data set in one way, and they focus on the decision to get a genetic test for working out whether or not you have genetic susceptibility towards breast and ovarian cancer. And the reason I say this is a very interesting genetic test is there is actually something you can do with this information to save your life if you take the test. So, potentially, this is a hugely valuable health -- piece of health data to create.

Now, the negative is that this is a technology in its early stages, and so as a result, we're only seeing a little bit of takeup in our sample, about less than one percent.

Now, what we're going to do in the paper is use standard econometric techniques to relate the decision of these people in our sample to go and get the genetic test, to what the state privacy regime was like in that particular year.

Now, I realize this is not an economist audience, so what I want you to think of this is as the

216

statistical relationship that we do where we're controlling for just about everything that you might think of going on in the background. We're controlling for the year, we're controlling for the state, we're controlling for everything about the patient.

Now, if you like equations and subscripts, the paper has got plenty of those, so I direct you there.

For now, though, for this audience, so what I decided to do is to present the main results in a bar chart, and the big punch line is, is that when we bucket up our state regulations in this way, what we find is that when you have informed consent, and that's informed consent where we're telling people how the data is going to be used, we get a reduction of a third or in terms of how many people are taking a genetic test. Now, this is a large proportion, but remember, these are quite small numbers, so the baseline is small.

Now, when we have a usage restriction, that is we say, oh, the state government says this data can't be used to discriminate, say by employer, say by health insurance companies, that really has no statistical effect that we can measure. The thing which has this big boost, or positive effect on the decision to get a genetic test, is whether or not you actually give individuals control over how that data will be used in

217

the future.

Now, when you get results like this as an economist, you're always going to worry, well, where are they coming from and what's the explanation? So, one explanation which worried me was maybe it's not about the patients, maybe it's about hospitals and whether or not they're offering the tests.

So, we went to the -- clicked more data to test this, and we found that's not really the explanation. It is the case that if you have these consent laws, hospitals react negatively, that's not a surprise. I've found that in the past. Basically it's because you have to construct an entire parallel system.

However, what was important about this study was we didn't find -- what we found was a negative reaction by hospitals in terms of whether they offer genetic tests to giving patients property rights. Again, maybe not surprising, why would you set up a genetic test in a facility of your hospital, probably to do some research, and this is going to restrict your ability to do research, but it suggests that the main effect of having these individual controls positively affecting outcomes is not driven by supply site, but instead driven by patients.

Now, more proof of this is what, again, a

218

typical thing we would do in economics is we're always going to worry about, well, you're saying this about patients, but could there be another explanation of something else going on in the state? We tested for this by looking at alternative explanations.

One such test was we looked to see, well, if we look at the decision to have an HIV test, which you might say you think of as similarly sensitive to having a genetic test, could we see any influence of the genetic laws on that decision? We found absolutely nothing, which suggests it's not driven by underlying tastes or privacy in that state.

Similarly, we couldn't find genetic law effects on flu shots, which suggests it's not driven by tastes for preventative care.

So, what is really going on? I've pulled out hospitals, I've pulled out disparate things from spurious correlation to the state, and I think what we're going to argue is that ultimately it makes sense when you understand how this privacy information is delivered.

Genetic testing is unusual in that you have genetic counseling where you sit down with a genetic counselor and you will discuss these privacy policies for perhaps 20 minutes, as well as the positive and

219

negative consequences of taking a test. So, this is very different from the typical online environment. We know that consumers actually found out about some of these laws.

The laws they don't find out about, though, are the anti-discrimination laws. These are usually part of the conversation, and I think that explains really the lack of the fact that consumers just aren't reassured because they don't find out about these laws actually existing.

On the other hand, when you go through the typical forms or process where someone is given informed consent and told how the data could be used, but not correspondingly given control, we are going to argue that highlights a sense of powerlessness which perhaps can explain some of the negative effect, whereas when you restore control to the patient over how their data might be used in the future, then you have reception control which tends to be a positive effect, which might encourage them going ahead with the test.

Now, we have some more material in the paper where we try and prove that this really is about privacy concerns in that we show that these effects are going to be higher in situations where there's more likely to be bad news if you have a genetic test, that is there's

220

reason to think you're going to have bad news from the test; however, we also show there is absolutely no effect if you've already got bad news. That is if you've already had cancer, the bad news is out there with your medical record, none of these privacy laws are actually going to drive any of the effects.

I'm also going to show the effects are largest for people who in their surveys took various privacy protecting actions such as refusing to state income. So, again, let's sort of draw it back to privacy rather than someone else explaining my results.

So, let me just sum up what we found. So, I want to emphasize, I think it's important for every empirical study there's going to be limitations, and certainly on this study, we do our best to try to make it causal; however, you can always come up with a whole bunch of explanations. We don't actually sit there in that patient and genetic counselor room when they go through the privacy policies, so we're speculating on the mechanism based by reviewing the privacy policies we've seen in different states, and the other biggest advantage is there was a study in the early stage of diffusion, and so this is going to be representative of the individuals who embraced new technologies earlier.

Having said that, I do think there is

something to be learned, which is that where the states give more control over how private information is shared, we do see an increase in genetic testing, and we see this increase particularly for people who are worried that there may be bad news from the genetic test.

Now, we found that in general informed consent, that is giving people information about how their data will be used, without giving them corresponding control just deters patients, and both patients and hospitals from having genetic tests and offering genetic tests.

Lastly, we found that data usage policies have absolutely really little effect, and so it's either good or bad news, depending on how you look at it. I was quite positively encouraged because usually when I run a statistical relationship between a privacy regulation and economic outcomes, I find a negative effect, so I was pleased to find nothing bad.

On the other hand, these laws are designed to help people, and perhaps my research suggests that they're actually just not being publicized enough to reassure patients.

So, with that, I will say thank you very much, and I thank you again to the organizers for giving me

the chance to speak.

(Applause.)

MR. MORIARTY: Thank you, Catherine.

Next up is Sasha Romanosky of the RAND Corporation presenting Examining the Costs and Causes of Cyber Incidents.

MR. ROMANOSKY: This has been a long day, hasn't it? Thank you all for sticking around, and thank you to the FTC for hosting this. It's great to be here.

I'm Sasha Romanosky. I will present some empirical work related to cyber events, and I'll define those in a second, but I want to explain a bit of the motivation, or at least two motivations behind this work. One is you've probably heard of this executive order by the president a couple of years ago to try to improve critical infrastructure, and as part of that developed this beautiful framework for cyber security. So, if anyone has any questions about how to protect their systems, you can go to the standard and take a look and it will tell you everything you want to know.

The trouble with that is that it's a voluntary standard. It's not meant -- certainly not meant to be regulated in any kind of way, I think despite some of the criticisms people have had. And, so, the question then becomes how do you get firms to adopt? How do you

get firms to adopt these standards? We think they're underinvesting in security, so how do we get them to increase their security?

That's a great question. And, so, the story behind this empirical work is trying to understand the incentives of firms. Are there incentives? Do those incentives exist for them to adopt more security or an appropriate amount of security or a fair amount of security or an efficient amount of security? We're going to look at that.

The other motivation, for anyone who has had a conversation with me over the past few years, knows that I am keen on cyber insurance, and the kind of empirical work that cyber insurance can perform and how they can, at the end of the day, help assess the risk of firms. Really, that's what they're interested in, is understanding the variation of risk across their firms to price that. And the kind of de facto policy that they are creating now with these policies.

So, with those motivations, what I look at, the data set that I have comes from a company called Advisen, which is based in New York and provides loss and incident data to insurance companies. They have been creating the data set on cyber events for a number of years now, but traditionally they look at loss of

property, other kinds of general liability that firms will face, these are corporate data events related to loss and litigation.

Most of the data sets that you see up there relating to cyber events include 5,600 observations, we have a data set of 12,000. So, as far as I know, this is the largest data set of cyber events, data breaches and privacy violations, which is very nice because it allows us to do some analysis to try and understand better different kinds of patterns and the risks that we will talk about.

I am separating the different kinds of events. When I say a cyber event, they are generally broken into these categories as I'm defining them. There are certainly other ways of categorizing them, and that's perfectly reasonable. For the purpose of my talk here, I'm separating them into data breaches, we normally think about as an unauthorized disclosure of personal information; security incidents, attacks against a company for the purpose of causing harm to that company; for example, a denial of service, or a theft of intellectual property or an outage of a system; and privacy violations, so this is what I'm calling an unauthorized use or collection of personal information; and then other sorts of phishing and skimming attacks. I think for this audience, we'll be mostly interested in the data breaches and the privacy violations.

One differentiator between these, the data breaches and the security incidents, what we might think of as acts caused to the firm, so they are bearing and they are suffering these attacks, as opposed to privacy violations where the firm is engaging in some kind of activity.

It's always useful to understand the data-generating process to understand where the data are coming from and what's included and what's not included. And, so, to be clear, these data come from public sources. There's no proprietary information. And Advisen has a wonderful team of analysts that go out and scour new sites, national and local news sites using Freedom of Information Act requests, they find the information using Lexis and West Law and other data sources. So, they have amassed this wonderful collection.

So, a cyber event will occur to a firm, a condition on that it will be detected by the firm, either by the firm, by a third party, by a consumer, by law enforcement, somehow it's being observed by the firm. We, of course, have no information about those events which are not detected, that's just not in our

data set.

Given detection, it is disclosed to the public. So certainly, of course, there is not always a requirement for a firm to disclose an event. There are exceptions, even with the breach notification laws. So, we do not observe those that are not being disclosed.

Conditional on disclosure, we would hope it would be reported within this data set, and of those events that are recorded in the data set, some will lead to a legal action, either private, public action, civil or criminal.

To give you a sense of the overall totals, we see that data breaches have, in fact, been increasing over the past few years. So, these claims by others that there are more breaches now than there were before do seem to be true; however, we find that they are increasing at a decreasing rate.

As opposed to security incidents, privacy evaluations, and these phishing and skimming attacks, which represent a much smaller proportion of the overall incidents. So, we see the first takeaway from this is that data breaches really represent the majority of these events. Interestingly, security incidents seem to be increasing at an increasing rate over the past few years.

Now, as far as I know, there have been no changes in regulation requiring disclosure, an increase in disclosure of security incidents, and so conditioned on the same level of reporting and detection. What this that might suggest is that firms are being attacked more now than they were before.

In regard to the insurance industry and trying to understand the risk of their insureds, one way to understand that is to look at analysis by industries. We might want to understand what kinds of industries suffer the greatest number of attacks or pose the greatest risk. And, of course, there are many ways to think of this. We could look at total number of events by industry, but that gives us an incomplete picture. And, so, we might look at the incident rate, proportion, the percentage of firms within a given industry that suffered the greatest number of attacks.

And then we could also look at lawsuits as just an aggregate, and litigation rate. We could also look at cost of events. I won't go through all of these in the interest of time, but I'll show you these. So, as a function of total incidents, the finance and insurance industry suffer the greatest number of incidents, followed by health care and government, education and then manufacturing.

But as a function of incident rate, government agencies, so these are states and local DMVs, law enforcement, courts, suffer the greatest incident rate followed by education.

Let me just skip through these. And then we look at the legal actions. So, of the 1,700 or so legal actions that we have recorded in this database, 300 or so are criminal actions, and some filed in federal court, some filed in state court, but really the bulk of these legal actions are private actions brought -- private civil actions brought in federal court. And these will be allegations of negligence, all sorts of common law and statutory allegations.

So, negligence liability and strict liability and breach of contract, unjust enrichment, a whole smattering. From previous research, we found almost 80 -- over 80 unique causes of action brought by plaintiffs in these suits.

When we look at the litigation, the total number of litigation and litigation rate, we see the privacy lawsuits have been increasing dramatically over the years, whereas the data breaches have been held steady. Now, these represent -- specifically these privacy violations -- in regard to the lawsuits, the privacy lawsuits, the allegations represent claims of typically unsolicited email or spam or faxing, unsolicited telemarketing, or recording, either video or audio recording.

And overall, the litigation rate for data breaches and security incidents has been decreasing over the years, which confirms some of our previous work, and so right now we're looking at a rate of about three or four percent.

What we also show here is that you'll notice that the litigation rate for privacy violations is really quite high, 95 percent, and I think this is really just more of an artifact of the data. I think while for the data breaches we can understand the sample of breaches and identify which of those have been litigated because of the breach notification laws, but for privacy violations, we don't really have that same denominator. We don't really understand the total number of violations, and therefore the percentage of which would lead to litigation. I think in our data set, all we're really finding is that we're only observing a privacy violation when a lawsuit is occurring.

Now, the next question, we're going to look at some cost data, and so I will couch this by saying that these are estimates of costs. They don't include --

certainly don't include lots of other information. They are all firm-based, so typically first-party losses, second- and third-party losses. So, all the costs that a firm would incur because of a data breach that you could imagine.

So, the cost of notification, the cost of forensics, the cost of repairing any IT systems. In some cases, they represent a dollar figure loss, like a financial loss. The third-party losses represent the loss -- the cost of litigating. Litigating the lawsuit, any kind of consumer redress, or financial sanctions imposed by regulating agencies.

So, given all these costs, the big question is how much does the data breach cost. And, so, Poneman has done a great -- have produced some great surveys over the years trying to estimate these costs, and what they come up with are typically figures of $5 million, $7 million as the cost of the data breach.

I might argue, though, that these -- this is an improper measure because they're looking at the mean, the statistical average, and so because of the variation of the distribution of these costs, a median is a better metric. So, not every data breach is a target of $270 million and rising. Not every breach is Sony, not every breach is JP Morgan or Home Depot. There are many firms

that don't lose that much money.

And, so, what we find here is that most companies lose less than $200,000. And, so, if you were to ask me the question of how much does a data breach cost, I would say less than $200,000. And, so, this is getting back with the incentives that firms may or may not have investing in security and privacy protection controls.

The median cost is a little bit higher for privacy violations, and that's still something we're exploring to try and understand exactly why, but I think the takeaway here is that this $5 million, $7 million cost is overblown.

We also wanted to look at repeat players. So, this notion comes up quite a bit in different conversations of what is the impact to firms that suffer multiple kinds of events? Are they bearing higher litigation rates, are they bearing a higher cost, how often do they occur?

What we find is that in our data set, almost 40 percent of firms are these so-called repeat players suffering multiple events. And that's quite a bit higher than I would have thought beforehand. I think that's quite extraordinary, in fact.

And, indeed, in the information and financial

insurance sectors, almost 50 percent of them are repeat players. I think that is quite interesting, also. The figures here that I'm showing, $9 and a half million versus $4 million are the mean, and what it's showing you is that the cost for the repeat players is almost twice, a little over twice than the nonrepeat players, those that suffer just a single event. Now, the medians are exactly -- showed exactly the same thing, that the cost is higher for these repeat players.

What I then also wanted to do is try to understand, okay, well maybe $200,000 is actually a lot for these firms, so what does this represent as a function of their revenue? So, what I did is went through all of the data to try and understand what do most companies lose as a function of their revenue and then try and couch that relative to other kinds of losses in different sorts of industries.

So, we wanted to look at retail, there's hospital, bad debt, global payment fraud. So, what you could imagine is that Visa and MasterCard have a certain tolerance for fraud, for bad debt, and that through either an organic process or some calculations, they have settled on some percentage. And these numbers come from industry reports, showing 5.9 percent, 5.2 percent, 3.1 percent for fraud. Cyber events, less than half a percent. So, it's saying that cyber events cost less than half a percent of a firm's revenue, a great deal less than these other industries.

In addition to that, in other works by some colleagues, Lily Ablon at RAND, we conducted a survey using the American Life Panel, a great survey instrument that RAND has available to it to try and understand consumer sentiment towards breach notification. How do they feel and respond to firms getting these notices of a data breach?

And what we find is that for the most part, they're really quite content. They're really quite happy with responses that they're getting, with the timeliness, with the information presented in the notifications, and really have generally no concerns. There was a small percentage, though, of them that change firms, but by and large they are really quite happy.

So, consumer sentiment, if it is, in fact, high, coupled with a small cost to a firm because of these events, really may suggest that firms have very little incentive to change their practices. Thank you very much.

(Applause.)

MR. MORIARTY: Thank you, Sasha. And thank

you to everyone for those presentations. They were wonderful, and very varied.

So, I want to recap them briefly, but first I want to introduce Doug Smith who is from the Federal Trade Commission, and Siona Listokin from the George Mason University School of Policy, Government and International Affairs.

So, we had four very different presentations. Jens presented an evaluation of two bug bounty programs and offered conclusions about how they can be effective to identify, resolve and reduce vulnerabilities. Veronica and Alessandro proposed an economic model for advertisers, platforms and consumers and concluded that the allocation of the benefits of sharing consumer information tends to benefit the platform and the advertiser. And if I'm wrong about any of these recaps, you can tell me in just a second.

Catherine presented an evaluation of the rate of genetic testing in states with privacy laws that fall into three different general categories, and concluded that states where redisclosure is restricted have the highest testing rates, and that states with informed consent decreases the rate of genetic testing.

And, finally, Sasha looked at one set of data and offered conclusions about the median cost of cyber

events, putting it around $200,000 and less than what other studies have found about the cost of cyber events.

So, to start, I just want to turn it over to Siona to offer some thoughts and start the questions.

MS. LISTOKIN: So, Kevin had asked me to talk about themes in this panel, and I would note that the title of the panel is The Economics of Privacy and Security, and I think that's about as close as we will get to a theme. Lots of variation here. Papers covered some of the most important or touchstone topics in privacy, so health data, online advertising, and, of course, security.

I would also point out that the panel had a lot more focus on how firms respond to incentives and not just consumers. And, finally, a lot of talk, or the papers here really are a cross-section of stages of research design. So, if you think about economics of privacy, you've got a model that extends existing theory, descriptive papers using new data sets, and explanatory or causal papers.

So, my meta theme here is that the field of economics privacy is alive and well and quite robust, but that's going to be my question. So, extending Commissioner Brill's comments after lunch, and the conclusion at the end of Veronica and Alessandro's

paper, in this field, what's your wish list?

And this is for everyone. Where do you see the gaps in this literature, specifically as it would relate to policymakers and industry practice? So, not just advancing academic research. I'll start with Veronica and Alessandro, but I'm interested in everyone's thoughts.

MR. ACQUISTI: One comment, and I'll piggy-back on our last slides about the piece in jail, which contained a letter from SSRN. In doing that review the leaders from economics and privacy, we identified three ways of research. The field is not novel at all, actually it started in the late 1970s, early 1980s, with Chicago School scholars like Paulson and Stigler. So, there's a beautiful pedigree and also quite a bit of work starting back 40 years or so.

However, only at the time there no models, like the late '70s, early '80s, no models or microeconomics in the field of privacy, it was more about using economic concepts, such as asymmetric information, models of that, and apply them to privacy. What we have now is lots of careful modern work, and what we start to see in maybe the last five years, 10 years, in terms of the work of folks like Catherine Tucker and others, it's beautiful empirical work.

So, in terms of my wish list is to see even more empirical work, and in order to have more empirical work, sometimes we need data from the industry. So, if the industry is serious and believes really that data is the new oil, and that more transparency is good for everyone, then we should address the problem of information asymmetry that Serge was referring to as one of the questionable problems we had in the previous panel, which is we really want more data from the industry regarding exactly what they do with the information they collect.

So that even if people, the end users, may disregard their privacy policies, they may not care about what companies are doing, researchers can actually study the data and then come out and aggregate it and understand what is really happening and then come up with maybe policy recommendations.

So, my wish list is more empirical work and more transparency from the industry side.

MS. LISTOKIN: Go ahead.

MR. ROMANOSKY: I mean, I would echo that, right? I think there has been a lot of time spent doing what a colleague would refer to as admiring the problem, and I think that's useful, and I think that's good, and I think that only gets us so far.

I like empirical work because it speaks to evidence for something. It gets us past normatives and values and what should there be, and it really speaks, it really helps answer the effect of, you know, what will be the effect of A on B. And certainly, you know, the causal inference is the gold standard. So, in order to do that, you know, the point is exactly true, you know, we need the data, right?

And sometimes that takes us being very creative on finding it in clever ways like the previous panel, the researchers did themselves, coming up with these experiments, which I think is beautiful. And sometimes it takes paying for it which, you know, is okay, too. But certainly I think we need the empirical work.

So, I would echo everything Alessandro had to say, especially in the wonderful accented way that he said it.

MS. TUCKER: Well, I'll just add to the accents. So I, you know, unsurprisingly, I agree for the need for empirical work. What always strikes me is if I'm a policymaker trying to decide if I want a minimum wage, or what the level of the minimum wage should be, I could draw on hundreds of economic studies that have measured in hundreds of different ways how minimum wages

affect wage levels.

However, if I'm a policymaker making really important decisions about whether to regulate privacy or data, I'm instead relying on just a handful of studies in potentially none generalizable spheres. So, really it's almost personnel and numerosity.

MR. GROSSKLAGS: I want to add something. On my wish list is perhaps a better understanding of the long-term consequences of both the loss of privacy and potential security compromises, and some work that Alessandro and I have done goes in that direction to understand how people perceive privacy decision-making over time, but what we could not assess in a robust manner is what are actually the potential losses that we may face down the road. And I think this is a very critical issue when it comes to genetic privacy, but also to consumer privacy.

A similar issue also rises in the context of security, where actually the most interesting things might happen in the context of what we do not observe. Right? So, you saw it in Sasha's chart, we could only analyze the data that was detected. So, what about all the security breaches that we do not observe and that we know nothing about?

Similar with respect to my presentation, there

is the behavioral white-hats, which we can now analyze in a reasonable fashion, even though this was one of the first works doing that, but what we do not observe is the behavior of black-hats, and there we still have a lot of work to be done in terms of investigating them and getting maybe qualitative data, but also tying together data sets such as Sasha's with, for example, analysis that we have done to kind of be able to infer where vulnerabilities have been known by the black-hat community that had not been discovered by the white-hats.

MR. ACQUISTI: May I add something? Jens has really said something really important about long-term effects, and here is the dilemma that is a field day that privacy faces. In my belief, the most interesting implications of data sharing and data protection are long-term and indirect. But generally, as economists, we can publish and do rigorous work when we have shorter, indirect effects. It's very, very difficult to do studies and find causal links over long spans of time when there could be a data breach now which only has an effect seven years later, and you are not going to satisfy reviewers in a rigorous type journal with analysts who try to find those kind of effects.

So, this data is far afield. I don't think

241

there is any simple methodological solution to that.

MS. LISTOKIN: Thanks.

MR. SMITH: I guess it's my turn. I have a question for the group, and I'm actually going to ask one of Catherine first. Catherine, so, you know, what your research showed is that different laws have different effects on consumers' choices in this particular context of sort of genetic privacy.

So, I was curious, sort of how you think this research -- what implications it has for other areas of privacy and data security?

MS. TUCKER: Okay. So, what was nice about this setting is it allowed us to have more of a horse race, where we had the same thing we were trying to explain in lots of different privacy regimes. Now, the extent of -- the reason I find it useful or reassuring is it helps me believe some of the other research I have done in other areas, which have been more case by case.

So, some of the research I have done, for example, in targeted advertising, which a lot of people have talked about today, is emphasize the negative effects of informed consent, but also positive effects from improving consumer perceptions of control. But I was always nervous because those were two very separate studies, different at different times, different

242

spheres, even different countries, and so I found it reassuring to actually use this horse race to make me think, well, perhaps there is something more generalizable we can say about the effectiveness of different privacy regimes.

MR. SMITH: Thanks. And then, the question I have for the group, is actually just a little bit of a follow-up on one of the things Siona pointed out, which is that you guys are looking a lot at how firms' choices have in this arena, so what is your -- what do these papers, and research in general, suggest about what the private sector is getting right? What is it getting wrong? You know, what can this improve on our understanding of what kind of market failures might be most concerned about in this area?

Why don't we start at this end I guess.

MR. GROSSKLAGS: What is the private sector getting right? I think one -- one observation that also Alessandro and I have made over the time is that we see a lot of entities, private entities entering the market with privacy-enhancing offers, but not really picked up in the marketplace to a sufficient degree. And, well, the good news is that we do see these offers, we see a lot of technological solutions that are eventually picked up by startups, but what we see less is an

243

adoption by the big players because of a lack of incentives.

Targeted marketing or advertisement is just too enticing to give it up in exchange for a more privacy-friendly practical solution. So, that's a fundamental conundrum that we are presented with that is very hard to sidestep.

Nevertheless, I think it's very important that we see these new offers in the marketplace, and I hope more of them are actually picked up and practiced.

MS. LISTOKIN: What are they getting right?

MR. ACQUISTI: Well, getting back to Jens' point about offers in the marketplace, one reason for optimism is I think the sense of privacy enhancing technologies, PETs. So, almost every time I'm invited here to the FTC I end my talk talking about PETs because I really strongly believe that technology is not just the problem, it can be the solution. Privacy technologists do not stop altogether the flow of data, but rather modulate, right? They are sharing the protection.

So, the reason for what privacy firms can actually make, this may be wishful thinking, but may be provocative in deploying PETs, anticipating otherwise very little regulatory integration, so that they can

244

still do much of what they are doing now, but in a more privacy-preserving manner.

Now, truth to be told, some of these technologies are still in its -- in their infancy. For instance, metamorphic encryption are still very promising, but we still don't know how efficient and practical it would be, but the promising is enough for the moment, and I do believe that we could in this race for privacy, we can actually have the cake and eat it, too, because of these technologies.

MR. ROMANOSKY: In terms of what are firms getting right, God, that's such a good question, and I wish I had a better answer than the one I'm about to give. I think -- so, I think what we can rely on is that firms will -- firms will -- firms will operate based on incentives. And, of course, the goal then is to tweak the incentives such that they become aligned for all of the players, right? So, that's not new.

And what that means is that, you know, look, if privacy really is a big deal, then consumers should really act like it's a big deal, and if -- and only until they do will firms have incentive to take it seriously. So, I guess I need to -- I guess I would say that consumers should take it seriously, and act like it, and then firms will take it seriously.

Now, if there are -- if there are market failures for which consumers can't -- you know, can't impose any kind of effect on the firm, then that's where, you know, that's where regulation or policy or FTC actions can come into play.

Go ahead.

MS. TUCKER: No, I just wanted to build on that, because I think what I often see in the discussion is this underlying assumption that it's never in firms' interest to regulate on privacy. And so that -- and therefore, government has to intervene. But I think there are instances that we see in research where there are incentives for firms to actually improve privacy protections for consumers.

For example, the provision of user-centric controls, and so I sort of see that as a beam of light in a rather cynical world.

MR. ROMANOSKY: Yeah, and I think, I mean, it does touch on the world of information disclosure and choice and notice, and poor choice and notice. You know? Poor choice, over the past five, six years, has taken a beating, hasn't it? But it's relied on this notion that, you know, firms don't behave the right way, consumers don't behave the right way because they don't have the right information, and only if we could give

them the right information would they make the proper choices.

I'm just not sure that's true. At least let me say it this way, that maybe firms, at least in my case, with the data, that maybe firms do have the right information, maybe they are aware of all of the risks that using and collecting the data have, and that maybe they are making rational choices. And for them, investing a certain amount, which we may think is underinvesting, isn't the proper amount, but maybe it is actually the right amount as far as they're concerned.

MR. GROSSKLAGS: I just want to also add that this panel was also about security, and I think one thing that firms do right is participating in bug bounty programs, and really taking serious efforts in hardening their web security, but also other security aspects. And I think they're still quite a step away from anything approaching full security, but I think having a multi-dimensional security program, including bug bounty programs, is definitely a step in the right direction.

MR. MORIARTY: Jens, on a related point, I wanted to ask you, there was a notorious blog post by the chief information officer of Oracle where she sort of denigrated the value of bug bounty programs, and basically the analysis was, look, it's very expensive to

go through all these bug reports, it very rarely yields useful information. Your study does show that there are benefits to participation, but I think the question that she was raising is, are the benefits of participation greater than the benefit of just using that same money to pay another engineer to evaluate internally the controls in your software?

MR. GROSSKLAGS: So, bug bounty programs are certainly not the first security measure that any kind of company should implement; however, as you saw in one of my early slides, actually very mature companies from a security point of view were the ones running their own bug bounty programs, like Facebook and Google and so on. So, from that perspective, it was worth their while.

And certainly one of the main selling points is that it provides a different perspective, in addition to running software security tools, having internal security researchers, in the sense that white-hat researchers have perhaps somewhat more of a view like a black-hatted organization, they are more creative, there are protocols in places where other security researchers would not look, and this is certainly a big selling point to inch the security of your website even a couple of steps further.

Also, I think there's a lot of criticism about

these bad ratios between the reports and the data that is actually then useful, and I think when you actually look very closely at it, a lot has to do with the matter of duplicate reports. And, well, I mean, this is actually white-hat researchers doing their job. If the reports have not yet been disclosed, then, well, they will report oftentimes the same kind of security weaknesses to the particular entity, and, well, taking this into account, then actually the error rate is not that high.

Last point here is that here, actually the involvement of bug bounty platforms can really have a positive impact, because they can introduce measures such as reputation mechanisms, adding security walls and so on that actually then also instill some part of competition between the white-hat community participants so that they are more inclined to actually provide high-quality data to the participating companies.

MR. MORIARTY: All right, we have 20 seconds left. So, Siona, do you have any final thoughts?

MS. MAROTTA: May I add something?

MR. MORIARTY: Yeah, Veronica, please.

MS. MAROTTA: Yeah, I wanted to clarify there are a number of findings where we don't find that intermediary is always bad, but sometimes it does do the right things for the consumers, so there are, like, policy in this case is more nuanced. So there are cases in which the intermediary is -- the interest of the intermediary is aligned with the consumers, but there are cases which, instead, its incentives may be contrasting with those consumers.

MR. MORIARTY: And your paper is not currently up on our website, but I believe that we will have it up following this presentation so people can, if they want more information about your findings, they can look at it there.

Well, thank you all for participating in this session. We really appreciate it. Thank you.

(Applause.)

(Whereupon, there was a recess in the proceedings.)

SESSION 5

SECURITY AND USABILITY

MR. ALVA: Our last panel of the day will look at issues around security and usability as it relates to privacy. So, I would like to welcome our first presenter, Sarthak Grover. He is a Ph.D. student at Princeton.

MR. GROVER: Thanks, Aaron.

Hi, everyone. I'm Sarthak, and I'll be presenting our work on The Internet of Unpatched Things. So, the main aim here is to basically look at the current state of devices. We basically ended up studying network traffic from a bunch of smart devices which are really popular, and we want to talk about how these devices may potentially leak user information.

My aim is to encourage you to think of how we can improve policies to stop this potential leak of information. So, how is the smart home or the IoT environment very different from the conventional mobile or computer environment? The problem here is that we have a lot of manufacturers, and we have small startups coming up with their own devices. They may be hiding device programmers.

Apart from that, these devices have low memories. They might not have capable hardware to

enforce security protocols, which we use in computers and phones. This makes it pretty difficult, first of all, to implement security protocols on these devices, but apart from that, a big issue is that the current smart phone model works like this:

Your devices inside the home send all your information to the cloud, to a particular server. In fact, if you have two devices in the home and they want to talk to each other, currently they talk to the cloud, and the information will get back to the home, and something will happen in the home then.

So, what we have is a pretty bad combination. You have hardware which is incapable, and you have information which is always being sent to the cloud. So, this combination results in potential privacy problems.

Now IoT devices which are sending network traffic without security protocols may end up leaking some information about the user. They may end up leaking information about what device is being used inside the home, and they may also end up leaking information about whether the user is home or what he is currently up to.

So, in a sense, what I am saying is anybody sitting on your network port may be able to find out

what you are doing inside your home, and this is a big problem. So, what our aim is right now is to basically take up a few devices in our case study and study what kind of personal information or user activity information they leak to the cloud.

So, what we did was we basically bought some popular devices. We went to Amazon.com, we searched for some very popular home network devices which people are currently using in their smart homes, and we ordered them. What I am going to show you is results for network traffic analysis for five particular devices: A camera, a photo frame, a hub, an Ubi smart speaker, which is like an Amazon Echo, basically, and a Nest Thermostat.

So, what we are interested in right now is what kind of information these common devices leak to the network. And the first device I pick up is the digital photo frame by PixStar. So, what we found out was that all traffic from this photo frame is sending clear text. There is absolutely no encryption happening, all right?

The cool thing is that this device can actually talk to your Facebook or RSS feed, so it's downloading photographs in the clear; and also, whatever action you take on this device -- for example, you press

253

a button, say you press the play radio button -- that will actually go in a clear HTTP packet which somebody, again, on the network can read.

So, if there is somebody sitting outside on the network, like somebody in the ISP or a malicious passive listener, they can see what you're doing through the photo frame. Apart from that, it's also capable of downloading radio streams, again, in the clear.

So, an example of what kind of information we see, so these are snapshots from Wireshark basically, and what we saw was that your email, which you configured your account with, is actually being sent in clear text. What this means is that this photo frame is potentially leaking account data, and anybody on the network port can actually have a look at this email.

Secondly, if you press a button on this photo frame, say you a press the "List Contacts" button or the "Radio" button, anybody, again, on the network IoT can have a look at what you currently pressed. So, somebody on the ISP can go, like, this person is currently listening to the radio from his digital photo frame, though I don't know why you would listen to the radio from the photo frame anyway. So, basically what I mean to say is that you can find out about the user's activity, as well as some other information, just by

254

looking at the network information.

The second device we picked up was a Sharx security camera. It's a pretty common camera which is used for security monitoring in homes. It has, like, motion detection. What we saw was that all the traffic, again, was being sent in clear text.

Now, this security camera actually requires a login. So, if you want to view the stream, you are supposed to enter a password, but that doesn't mean the stream itself is encrypted. In fact, anybody sitting on the network can still have a look at where the stream is going and what the stream is.

Also, if you go to the web interface and you press a button, whatever you did will still go in an HTTP GET packet, again unencrypted. So, videos are being sent as JPEG frames. Also, if you have pressed the FTP button, then all your data is being uploaded to the FTP, again in the clear.

And this is an example of what things look like. So, the FTP is actually using some pretty random ports, so you can't really rely on the network to secure you, again, because these are nonstandard ports which are being used by the device. And things -- like, this is basically private data which is being uploaded.

The third device we ended up looking at was

255

the Ubi. So, this I think is like a precursor to the Amazon Echo. Basically, it's a small voicebox which you can talk to, interface with other devices. For example, we had this Ubi interface with the Nest Thermostat in our houses.

So, what we saw was that all voice to text first of all gets converted -- all voice you talk to through the Ubi will get converted to text on the device itself, and then text is sent in clear, again, to a server outside. The server here was the Ubi.com.

Apart from that, the Ubi also has certain sensors, for example, light sensors and temperature sensors, which are still sending readings in the clear. And the interesting thing over here is when we interfaced this device with the Nest, it used encryption and spoke over HTTPS, but when it was talking to its own server, it was using HTTP, and everything was in clear.

So, clearly, this device actually has the capability of enforcing security, but somehow, the policy, whatever policy they came up with, they did not enforce encryption for their own device streams. Only when they're talking to the Google API, they enforce this encryption.

So this is an example which shows how sensor readings were available. Now, these sensor readings can

256

leak information about whether the light is on in the room or not. In a sense, somebody on the network, who's on the path, can know whether there's a user inside the room or not based on the luminosity value.

Furthermore, when you were chatting with the device, all the text was converted to clear text and then sent to the network. So, here you can see an example of what the chats looked like when I monitored them on the laptop gateway.

The next device we looked at was the Nest Thermostat. Now, we're actually coming to the more secure devices and the big ones, too. The Nest Thermostat from Google actually was pretty secure. All the information was over port 443, basically using encryption and HTTPS.

Now, what we also found out was that some of the updates incoming were in the clear, and we weren't sure why, so we contacted Nest about this; found out it was a bug, and they fixed it.

So, here is an example of what we found initially. Outgoing traffic was secure, but incoming traffic, some of the updates were not secure. They were in the clear text, they had some information regarding the location, and when we told Nest about it, they thanked us, and they fixed it.

257

All right. And the last device which I'm going to talk about is the Smartthings Hub by Samsung. Again, a pretty popular hub from a pretty big company. The good thing was almost all the traffic coming out of this device or going into the device was totally secure over DNS. There was no clear text or port 80 traffic at all, and the flows were all to Amazon AWS instance.

But the interesting thing is, even though this device is in itself secure -- and, in fact, I see this as the model of future IoT devices, which are completely secure, there is still some background information, like three or five packets every 10 seconds, going to smartthings.com, which can somehow let you fingerprint the device.

The good thing is that the Smartthings is a hub. Basically, what that means is that you have other sensors attached to Smartthings over some other protocols, like ZigBee or bluetooth or Z-Wave, and you don't have a direct view of the sensors, so Smartthings itself makes all the information coming out of the house secure and then sends it out. But a person sitting at the ISP level can still find out that, you know, you have a Smartthings Hub inside the home.

So, this brings us to my conclusion and some implications on the policy. Basically, I don't want to

258

sound pessimistic or dramatic, but that's what the heading is, "Be Afraid!" We know it's very difficult to enforce security standards on smart devices. Inherently, I mean, there are multiple manufacturers. There are only a few big ones, but a long tail of small ones. Smart devices come up on Kickstarter, and people buy them. It's difficult to ensure that they all follow the same standard.

These devices are also sometimes very low capability. They don't even have a way to implement DNS on the packets they're sending out, and they also end up using nonstandard ports and protocols. But the good thing is that we are trying to make an effort. For example, I found this handout outside regarding "Building Security in the Internet of Things," and that's good, because it means we're trying to enforce security at the building block level itself. So, maybe the new devices which come out would have security inherent in them.

The second thing is, okay, so we fixed devices which are coming up now. What about devices which are already present? How do we get people to patch them up or fix them? So, for example, we want to encourage people to look for bugs, and one way would be bug bounties, but as we've seen in previous talks, you know,

259

bug bounties may work for the big guys, but the IoT domain has a lot of small manufacturers coming up, and we don't really know if bug bounties will work for that, and we don't know if the device will be popular enough to have users which actually look for vulnerabilities in them.

So, the main point is, how do we enforce such kind of things? Like, who is responsible here? Will the Government try to enforce bug bounty programs or is it the manufacturer which goes, if you find a bug, we will give you money?

And lastly, who pays for this patching in the update part? If this is using your network, is the user responsible for anything which goes wrong or the ISP or is it the manufacturers?

So, I want to end with some of the work which we are currently up to right now. So, we've talked about how we can improve future devices in terms of their security and privacy policies. We've talked about how we can improve current devices by trying to find bugs in them and vulnerabilities in them.

The approach we are taking right now is how to improve security and privacy policy on the network. Basically, we are trying to offload policy to the network layer. For example, in case of a smart home,

260

all our information is going to go through a gateway which is inside the house. This gateway might be provided to us by the ISP or it might be our own, but maybe there are parts of security that we can implement at the gateway itself. Maybe we can tell the gateway to enforce certain standards regarding the network protocols which are being used by devices or, at the very least, this gateway would inform our user that, hey, there's something wrong with your devices or this device is not using the right security standards.

So, what we are looking at currently is can we offload device security to a gateway or the network layer? And secondly, how much information about the user behavior is actually leaked to outside the home network?

All right, thank you.

(Applause.)

MR. ALVA: Thank you.

Next we will have Professor Vitaly from Cornell University and Cornell Tech.

MR. SHMATIKOV: Hi. So, I am Vitaly, and I'll be talking about mobile advertising today. Mobile advertising is pretty big these days. If you look at modern app stores, you find that a significant fraction of apps are free to the user, and the way they make

261

their money is by incorporating advertising.

So, it seems like a very reasonable question to ask, what information about the user is actually available to advertisers? That is, if an advertiser submits an ad through a mobile ad network and that ad gets shown on a user's phone, what can an advertiser find out about the user of the phone on which the ad is being shown?

So, that seems like an interesting question for which, apparently so far, there hasn't been a good answer. Very few people have investigated this, so this is what we decided to investigate in this project, to look at this. But in order to understand this, we first need to understand how mobile advertising actually works from a software perspective. So, it requires a little bit of reverse engineering of how mobile software actually shows ads to users, how it actually works.

Mobile advertising is a little bit similar to web advertising with one crucial difference. So, in web advertising, you typically have a web browser and the web browser is just showing an ad. That has been studied a lot, and even today, we have heard a lot of talks and conversation about web advertising.

In mobile advertising, things are a little different because there is something in the middle;

262

namely, there is an app library. So, the way mobile advertising works is that apps that are supported by advertising, they typically include a little piece of code called an ad library, and it's that piece of code that's actually showing ads. It's not the app itself. It's the ad library inside the code.

And it's actually very common for modern apps to incorporate multiple advertising libraries because they make more money that way. So, maybe between a third to half of all apps that are ad-supported actually include multiple advertising libraries for multiple providers that are being used to show ads to users.

So, the question I'm asking, just to repeat it, what do these ads that are being shown inside these mobile advertising libraries, what do they actually know about the user or what they can find out? In order to do this, we need to look kind of at them structurally, this whole ecosystem, and I promise I'll try to make it as painless as possible, although investigating it was fairly painful and involved a significant amount of reverse engineering, but it roughly looks something like this.

There are three kind of big parties in the picture, so there is that app, which is being shown on the phone. There is the advertising service, which is

263

supplying ads to the phone. And then there is the advertiser whose ads are being shown. And there has been a lot of work previously on trying to understand what information is available to the advertising service, but instead we are looking at what's actually available to the advertiser, and it's not the same question because there is a big difference between the advertising service and the advertising.

The advertising service is typically a reasonably respectable, reasonably reputable company that's maybe owned by, you know, Google or Twitter or some kind of recognizable entity. It is a large business. They have a reputation at stake. They make a lot of revenue.

Whereas advertisers are people who actually supply these ads that are being shown. Who knows who they are? I mean, this is dynamically determined. They are phished in real time, sold by auction, syndication, in all sorts of ways. These advertisers are not necessarily trusted. It's very hard to determine what information they're trying to extract.

And that's why mobile advertising libraries go to fairly significant lengths to protect users from malicious advertising and from snoop advertising -- from advertising that stealthily tries to extract information

264

about users. They use a variety of technical mechanisms to achieve this, and I'm not going to go into them. You can read our paper if you want to find out more about this.

The short summary is that what they try to do is they show every ad that they show to the user inside a little browser instance. So, there is a little web browser inside every advertising library, and they create a quote of this web browser every time they want to show an ad, and they show an ad inside of it.

And the good news about it is they, in effect -- they can effectively rely on security and privacy protections inside web browsers to protect phone users from malicious advertising. So, technically this is known as same origin policy, but you can think of it just as a way of sandboxing untrusted advertising to make sure it doesn't have any access with the underlying phone and cannot learn anything it's not supposed to learn from the phone.

And mostly it works, with one little exception. Mobile ads these days need access to what an Android phone is known as external storage, and the reason they need to do this is for rich media, because people who view advertising, especially people who supply this advertising, they want rich experiences, they want video, they want images, and because of that, they need to cache a lot of information on the device, so they have access to external storage.

But to be safe, they allow ads to load files from external storage but not to read them. So, it cannot read it. It can just load it and show it to the user without being able to read it. So, that looks fairly harmless, except that Android external storage is kind of this weird thing. In Android external storage, there is really not a whole lot in the way of access control protections.

That is, if there are multiple apps running on the device and they store files in external storage, they can read each other's files. And that, you know, may not be ideal from the security perspective, but this should not really imply a whole lot about security and privacy of mobile ads, because as I told you, mobile ads cannot actually read other apps files from external storage. They can try to load them and try to show them to the user, but they cannot actually get access to them directly. They cannot look at their content.

So, so far so good. So, it seems like this whole way of protecting users from potentially malicious mobile ads is fairly carefully designed and carefully thought through, except that there is this one little --

I keep talking about this one little weird thing. They cannot read them, but they can try to load them. Why is this interesting?

It turns out that by trying to load a file that doesn't belong to them, mobile ads can learn a little bit. They can learn literally one bit of information. They learn if a particular file exists on the device or not. They cannot read it. They just learn if a file with a particular name exists. That seems like, okay, all right. That's fascinating. Why am I talking about this? Because that's really a very small amount of information.

So, now let's look at how this information might be used by a mobile ad. So, let's take an application which actually has nothing to do with mobile advertising, it's just a popular application in Google Play store that happens to be a drug shopping application. So, this allows people to go and look and reach pharmacists, drugs. You know, if somebody is taking a particular medication, they can find a pharmacy nearby where the price is lowest on it.

So, in this particular case, you can see there are some mitigations. There's particular things -- actually, the fact that a person is taking one of these might be considered sensitive because these have to do

with anxiety and various psychological disorders.

So, what this app does, if a person is regularly shopping for a particular drug, they need it faster, it takes the picture of the pill, like the literal picture, like I'm showing here, and stores that picture in external storage on the device so that next time it's faster to show this picture.

Now imagine that there is an ad running in a different app on the same device, okay? It has the app that's showing the ad, but a totally random ad. It has nothing to do with that pharmacy shopping app that I showed you before; however, as I told you before, an ad being shown on it has the ability to ask a very simple question. Does a file with a particular name exist on the external storage? And in this case, it's asking for a file whose name corresponds to the image of one of the anxiety drugs.

So, what can a mobile app -- and this is a question to you guys -- what can a mobile app learn from the answer to this question? So, all it learns is one thing, if the file with a particular name exists on the device. What can the app learn by knowing the answer?

AUDIENCE MEMBER: That you order or you use or have some interest in this drug.

MR. SHMATIKOV: The only reason -- if the

answer to that question is yes, the only reason a file with this name would have existed on this device, if the user used that app and searched for that drug. There is no other reason. So, then if they see an -- if an ad sees that the file like this exists, it cannot read this file. All it needs to know that this file exists. It learns with hundred percent certainty, because this name is unique, that the person has been shopping for a particular drug.

And this turns out to be a pervasive problem, because -- and remember, this ad is being shown in a totally different app. It's not even being shown in the pharmacy shopping app. It's just being shown in some ad which is shown -- which is running on the device, maybe even later, not even at the same time as the pharmacy shopping app.

And this turns out to be a generic problem, is that if there is an app that need not even be advertising supporter, that puts files under external storage, like a lot of them do, in a way that depends on the user behavior, then an ad shown in any app on the same device can determine that this file exists, and from the fact that this file exists, it can infer what user behavior led to the presence of those files.

So, I have shown the example with drugs -- and by the way, this violates nothing in the security policy, because all the security policy says is that it cannot read these files, and it cannot. It does not read the file. It just learns that the file exists.

And this -- actually, it turns out that this affects all kinds of mobile apps. Here is another app. This happens to be a mobile web browser which caches visited pages and files with predictable names. Actually, the names of the files are derived from the URL of the pages that the user visited, and it's vulnerable to exactly the same attack.

A malicious ad running in another app can look at the presence of certain files on the device, and it can figure out which sites the user visited recently, because the only reason a file with that name would appear on the device is if it were cached by the user's mobile browser as a result of a previous visit to a particular website.

And in our paper, we have many more examples about the inference that it could be done this way. We actually did an analysis of several very popular advertising libraries, including AdMob and MoPub which, for instance, had a very significant fraction of Android apps. They all, at least at the time of our study -- I'll tell you in a second what happened later -- have

this vulnerability, meaning that a mobile ad shown in any of these libraries could infer information about the user by presence of cached files raided by other apps.

We also looked in our study at other issues like the leakage of location information. I'm not going to go much into detail about this. I will just show you this picture, and the only thing I want you to admire about this picture is how complex it is, because it shows, like, how in five stages, literally, in MoPub, information about the user's location can be extracted by an ad, but it works pretty reliably and doesn't result in a mobile ad running in mobile, can create, like, very nice trajectories of user movement like this, which immediately reveals a ton of information about the user, including actually the user's identity, if one of these happens to be like a single-family residence where the user lives. So, this is really fine-grain information that can leak out through these interactions.

Okay. So, what are the lessons of this study? As far as I know, this is the first reasonably comprehensive study of, first, how advertising libraries on Android try to protect users from malicious mobile ads and snooping mobile ads with intermediate success, as you can see. It shows -- and this is a slightly more

technical result, but nevertheless important -- that standard web isolation policies that are used in web browsers, here are exactly the same things I used in the mobile context, and they are no longer sufficient because they no longer prevent leakage of sensitive information. Something more subtle is needed here.

We actually, when we first did this study last summer, we didn't make it public right away because we actually wanted to work with developers of these advertising libraries and companies that deploy them so that they can fix at least the most severe vulnerabilities that we identified, and, in fact, some of them did, in particular AdMob, which is the biggest Android advertising service -- actually, they're owned by Google -- and they fixed that in the latest release of their AdSDK.

Some library developers told us to go away and not bother them anymore. I hope they won't do this after this talk. And if you want more detail, we have our paper online. It's written for a technical computer science audience, but I hope at least the big themes will come across from that.

Thanks.

(Applause.)

MR. ALVA: Next we will have Florian Schaub.

He is a post-doc fellow currently at Carnegie Mellon.

MR. SCHAUB: Hello, everyone.

I'm going to be talking about a project called the Usable Privacy Policy Project, and this is a large-scale project funded by the NSF under it's SaTC program, and I'm a post-doc on this project. Norman is actually the lead PI on this project, and it's a collaboration with many people at CMU, Fordham University, as well as the Center for Internet Society at Stanford. And you can read more about the project at our website, usableprivacy.org.

I'm going to give you a short evaluation and then give you an overview of what we do in this project in different parts. So, we look at privacy policies, and privacy policies originally had this promise of service providers would disclose the data practices so users can then make informed choices about which service providers or websites they trust with their data, but the reality looks a little bit different, because privacy policies play different roles for different stakeholders.

So, for the service providers, it's not really about informing the users. Most of them, when they draft a privacy policy, the goal is to demonstrate legal and regulatory compliance and in this way limit their liability. And regulators are happy about this. They use these privacy policies to assess and enforce compliance.

So, there's actually a nice and strong interaction between those two players, but that means the user kind of gets left out. And as a result, these privacy policies are long, they're complex, they're difficult to understand, they're full of jargon, they don't really offer many choices to users, and I think we all know by now that users mainly ignore them.

This puts us in this really weird situation where these policies outline what companies do with our data and what we allow them to do with our data, but this information is not used by the users or made apparent to them.

And there has to be much work on overcoming the status quo here. Proposals like layered privacy policies, showing short summaries of policies, graphical approaches, as well as machine-readable privacy policies, but many of these approaches don't go anywhere really because they lack industry support and there are not sufficient adoption incentives for companies to actually implement those solutions that have been proposed.

This is where really our project comes in,

because we are looking at semi-automatically analyzing these natural language privacy policies that most websites, most mobile apps already have, and we analyzed them to then extract key data practices out of these policies, and we do this by combining crowdsourcing, machine learning, natural language processing, and this way enable large-scale analysis of privacy policies.

And at the same time, we look at modeling users' privacy preferences and concerns so that we can actually provide them more effective notices that focus on those information aspects and data practices users really care about and give them information that is actionable.

Our project has many tightly interconnected threads, and I'm not going to try to untangle this for you right now. Feel free to look at our report to get a deeper insight there. But basically we have two goals. One goal is we want to better inform users, we want to give them notices that actually inform them and provide them with choices, and we want to inform public policy by showing issues with privacy policies, as well as showing ways of remedying those issues and also providing -- hoping better notices could be provided.

And to identify data practices of interest, we approach it really from three different perspectives.

So, part of our research team looks at legal analysis. Joel Reidenberg and his team analyze privacy harms in litigation cases to see what issues come up the most. We conduct user studies where we determine what are privacy preferences, concerns and expectations of users. Ashwini this morning talked about some expectation work in that context.

And we also look at the policies themselves. So, how are they written? How are data practices actually expressed in those policies? We have some work going on right now that looks at quantifying the huge ambiguity and the vagueness of privacy policies.

To analyze these policies, we started by building an annotation tool that basically allows us to give policies to crowdworkers or other annotators, and this kind of tool shows them the policy on the left hand and then a question on the right, and we ask them to answer the question but also mark text that basically provides the evidence for their answer. And this is really important, because this text selection, in combination with the answer, then helps us build machine-learning models and frame machine-learning classifiers.

And by showing these questions or tasks, annotation tasks to multiple annotators, we can actually

get quite robust results; however, you know, just giving this to untrained crowdworkers and saying, oh, well, 10 people say that's okay, is not really a good idea. So, we collected studies to compare the performance, annotation performance of experts who either write policies or have long experience in analyzing policies, graduate students in law and public policy and untrained crowdworkers recruited from Amazon Mechanical Turk, and we asked those people to annotate different privacy policies. The crowdworkers and skilled annotators and grad students annotated 26 policies, and then six of those policies were also annotated by experts.

I'm not going to go too much into the details for the sake of time, but one of the interesting results is that even the experts don't always agree on the interpretation of a privacy policy, and one reason for that is that the policies are vague but also that they are sometimes contradictory, and there are just too many different contexts handled in a single policy.

The good news is that for data collection practices, those are relatively easy to identify and to extract, they're usually in one part of the policy, but data-sharing practices are a bit more complicated. They are spread out throughout the policy. Sharing is mentioned in many different contexts and parts of the policies. So, it's kind of difficult to extract finer nuances reliably.

Now, when we compared the performance of the crowdworkers to the skilled annotators, we actually find quite encouraging results. So, when we hold the crowdworkers to a certain quality standard, 80 percent agreement, which means eight out of 10 crowdworkers need to come up with the same interpretation, then we actually find that in a large number of the cases, these crowdworkers agree with the interpretation that our grad students find as well. So, they come up with an accurate interpretation.

In almost all of the other cases, they don't reach agreement, which means they don't give us wrong answers. We have a very -- this dark bar shows us, it is the percentage where they come to a different conclusion than the skilled annotators. So, this is great. So, either we get an answer from our crowdworkers, and then with a highlighted group that's actually correct, or we don't get an answer, and that tells us that the policy might be too -- might be vague on the particular issue we're trying to analyze.

So, this shows that accurate crowdsourcing of privacy policies is feasible, but privacy policies are still long and complex. So, we look at leveraging

machine learning and natural language processing to further enhance those extraction tasks and make it easier for crowdworkers to complete these tasks faster without loss of accuracy.

And one approach we've tried to do or we've been developing here is predicting and highlighting relevant paragraphs. So, we take the answers we have from our skilled annotators, and we use that to train logistic regression-based relevance models for different types of data practices we want to extract, and then we highlight the top five, top 10 paragraphs that most likely contain answers or information about the data practices we want to extract. And what we find is that that really helps the annotators to come to conclusions faster without losing -- without affecting the accuracy.

And we did additional experiments where our analysis -- where we looked at do they actually just focus on those five paragraphs or do they also read other parts? And they do read other parts of the policy, but it helps them to focus their search and find parts in the policy again.

Another thing we do is we split up this relatively complex task of reading a privacy policy by splitting the policy up in smaller paragraphs and then giving a crowdworker only a single paragraph. We can

further split those tasks as well. So, rather than asking them multiple questions at once, we first ask one set of crowdworkers to kind of label in what category -- what category of data practice is described. Is this a sharing practice? Is this a collection practice? Is this maybe about user access?

And then in follow-up questions, we can ask more details that is the particular aspects for that kind of category. And that means that the task interfaces we can show to crowdworkers are a lot more compact and they can complete those tasks faster and with lower errors.

And based on that, we have developed an annotation scheme that really makes use of this approach. This is an interface not for crowdworkers. We are using this with law students, but the next step is to then break this up again with the product I just outlined.

But this is a very fine-grained annotation approach, and we're currently collecting data from law students. We already have over 100 policies annotated, and this provides a really rich picture on how information is represented, how data practices are represented in the policies.

We're going to release a data portal to allow

exploration of this data on Privacy Day this year, January 28th. So, visit our website towards the end of the month.

And the nice thing about this data is it's really helpful to train machine-learning and natural language processing models and drive research in this area.

Ultimately, what we would be hoping for is that we can actually automate the extraction, and one approach we've been working on here is paragraph sequence alignment. So, if I have a paragraph in one policy, an Amazon policy, and this one is about collection of contact information, and then if I compare that one to a paragraph -- to other paragraphs in other policies, there's a high likelihood that I can find similar paragraphs that also describe the collection of contact information, and this way we can basically reduce which paragraphs we might not have to show to crowdworkers and this way automate some of the annotations and analysis.

Now, once we have all this data, we want to provide notice to users, and here we focus on making sure the information we give users is actually relevant. So, we highlight unexpected practices, practices users care about, and information should be actionable. If users can't make a choice, then there's no point in showing them information, because they're just going to be helpless. We heard about this this morning, about users becoming resigned because they can't make any choices.

So, what we do is -- what we want to do is we want to show them the choices that are made available in the privacy policy -- there aren't that many -- but because we can scale up this analysis to many websites, we can also show more privacy friendly websites as alternatives to users and in this way offer them choices that go beyond the policy of what a single website might offer.

And we're currently in the process of developing a browser plug-in to basically make this technology available to users, and the idea is that we display a limited set of relevant practices, and we're going through an iterative design process at the moment with focus groups and online studies, but hope to be able to release this plug-in to the public this summer.

So, in conclusion, what we do in this project is we semi-automatically analyze privacy policies, and we do this with crowdsourcing, with natural language processing, and machine learning. And the goal of our project is really to enable large-scale analysis of

these privacy policies.

So, at the moment we are annotating 100 policies. In a year, we are hopefully annotating a thousand policies and we are doing it at the same cost or even cheaper. That's the idea.

And at the same time, we're really interested in understanding what users care about so we can on the one hand focus the analysis, but also help regulators focus their activities potentially to look at those issues users care about or are concerned with.

We want to show ways to effectively inform users about the data practices that are currently lost in those policies. No one is going to read the policies. So, if you want to make those policies usable, we need to extract the information that is really relevant to users and show it to them in a format that actually makes sense to them and actually allows them to act on it.

Thank you.

(Applause.)

MR. ALVA: And our last presenter of the day will be Norman Sadeh. Norman is a professor in the School of Computer Science at Carnegie Mellon.

MR. SADEH: Well, good afternoon.

I think very few people in this audience

probably appreciate how much progress is actually being made -- we have been able to make over the past few years in both modeling and predicting people's privacy practices, and so this talk is about sharing some of these results with you and showing you also how this effectively supports revision of developing personalized privacy assistants; in particular, the success or at least the early success we've had with mobile apps in particular, and how this, we believe, can be extended to IoT.

So, this is joint work with a large team that will be acknowledged on the very last slide. I don't think that I am going to have to work very hard to convince this audience that people care about privacy, and yet as we all know, also, people are often very surprised when you tell them, for instance, what sorts of apps they have downloaded on their mobile phones and what information is being collected or shared by these apps.

This is just an example of an early study we conducted in this space. The biggest offender in that case was an app that some of you at the FTC are quite familiar called Brightest Flashlight, and 95 percent of the people who had that app were extremely surprised and very upset to find out what information that app was

actually collecting.

And, so, as we all know and as Florian just emphasized again, very few people read privacy policies, and that's sort of the reason why we have this level of surprise. Also, as I think we are also realizing, many of us have tons and tons of settings and just don't have the time to configure all these settings.

For instance, if you are a smart phone user, and as most smart phone users, you probably have somewhere between 50 and 100 apps on your phone. These apps typically will require between three and four permissions. These are permissions to access some of your more sensitive information. If you do the math very quickly, you realize that this would require people to configure somewhere around 150 different settings. How many people are willing to configure 150 settings on your cell phone? Not that many.

And, so, with this in mind, and obviously with a recognition of these challenges both already on the fixed Internet and in the mobile space, the natural question is, well, if this already doesn't work on the fixed web, if this already doesn't work on the mobile web, what are the chances that it's going to work in IoT, with the Internet of Things?

And, so, our vision in this space, as I said, is this idea that perhaps personalized privacy assistants could be developed that will actually reduce the burden and allow you to manage your privacy better across these different environments.

And, so, the idea is that this personalized privacy assistants in particular will learn over time your privacy preferences and will be able to automatically configure many of those settings based on various correlations between how you feel about sharing your information with one app versus another app; based on also understanding what your expectations are, going back to the presentation that was given this morning by Ashwini Rao, who has been looking at these issues.

In particular, for instance, if you think, as Florian also mentioned about privacy policies, when you read these privacy policies, they tend to be very long, very verbose, but very often at the end of the day there's only a very tiny fraction of that text of that policy that matters to you, and perhaps even a tinier fraction of the text that pertains to things that you didn't already expect.

And so perhaps these personalized privacy assistants could help us actually highlight -- could help us by highlighting those elements of policies that really would be a surprise to us, that perhaps would

lead us to modify our behavior as we enter a smart room, for instance, in an IoT context.

Perhaps these personalized privacy assistants could also help motivate users to revisit some of their settings to verify that they still feel the same way. Privacy preferences are not fixed. They might change over time based on your experience, based on what you learn.

And, so, again, what I would like to do is I would like to share with you some of our success that is actually supporting some of the early elements of this functionality. What you are seeing here is effectively an early model that we built about how people felt sharing their information with various mobile apps for various types of purposes, whether the app requires this information for internal purposes, for sharing with advertising networks, for profiling purposes, or for sharing with social networks.

I'm not going to describe this chart in great detail, because time is limited, but effectively what we are supposed to see here is that people don't always feel the same way on average when it comes to sharing their information. There are clearly differences between sharing location information at a fine level versus sharing it at a coarse level. There are

differences when it comes to sharing, for instance, access to an SMS functionality and certainly depending on whether you are going to be doing that for advertising purposes versus using it purely for the purpose of the app that you are trying to download. People are going to feel very differently.

What this figure, however, doesn't show is how difficult it is to actually configure settings, and the reason why it's difficult to configure settings is that this chart here, as you see it, is not the whole story. The whole story actually comes out when you start looking at this other chart, which shows you the standard deviation when it comes to these preferences.

And, so, the story here and the reason why privacy is so complex is that we don't all feel the same way about these issues. If we did, it would be simple to come up with defaults and use these defaults for the entire population, and then we would be done, and perhaps even the FTC could jump in and say, well, nobody feels comfortable about this; therefore, we are going to outlaw it. Clearly that is not the way we operate.

And, so, the reason why this is complex is because we have this diversity in preferences. Some people are quite fine with their fine location being shared with advertisers and others object. The good

news, however -- and this is a result that has come out from research over the past years -- is that very often it is possible to organize the population and their preferences into a fairly small group of people, fairly small groups of people that feel very much the same way about these issues.

And, so, what I want to share with you here is, again, an early example of our work in this area, where, again, looking at these mobile app permission preferences, we're able to organize a population of users in just four groups, and just based on these four groups and what we're able to predict based on the preferences within each of these four groups, we're able to show that it might be possible to predict somewhere between 75 and 85 percent of their pricing preferences when it came to configuring their permission settings.

And, so, this is very, very simple technology. I'm going to show you that we've been able to go much farther than that. But that gives you a sense already for how easy it is actually to predict many different settings that perhaps people would want to have.

So, this next chart here shows you the next step in our research in this area, where we looked at actually a population of 240,000 users. I should actually say a population of 3 million users, but we had

289

to clean up the data quite a bit, and we eventually zoomed in on the fraction of the population that was most engaged with their permission settings.

So, these were users who were using a variation of the Android operating system. It was a nerdy version of Android where users could actually configure many different settings. And we were able to show that through profile levels or through personalized learning, we could just, by asking people a very small number of questions, effectively predict most of the settings that they would need to configure on their smart phones for the apps that they were going to download.

So, for instance, if you were to ask them just six questions, you could effectively reach a level of accuracy of about 92 percent. If you're willing to double the number of questions you are asking, you're getting close to 95 percent.

Now, we are not suggesting in any way that you should fully automate the setting of privacy permissions. We strongly believe in dialogues with users, but there are situations where it's extremely clear of how the users feel about some settings, and there are situations where you can determine that actually your model is not good enough to predict what

290

those settings should be, and that's where you should ask the user, right? And that's effectively what we are advocating.

And, so, we have gone one step further this past summer, and we actually piloted this technology with real users on their actual cell phones. So, we developed profiles. In this case, we came up with several different profiles and asked people to download this very early version of a personalized privacy assistant.

This assistant would ask them between three and five questions based on the actual apps they had on their cell phones, and based on their answers, it would recommend a number of different settings, as you can potentially see on the right-hand side of the slide in front of you.

And, so, we ran this, and to make a long story short, we ran this for effectively a period of 10 days. The last six days of the study, we actually tried and see if we could nudge users to modify the settings that they had adopted based on recommendations made by these assistants. We tried very hard with nudges like the one you see here. These nudges are very effective, by the way.

So, when it comes to getting people to rethink

291

their privacy preferences, when it comes to motivating them, we've actually got an entire study that shows that those types of nudges work very well.

And, so, here's what we found. We found that among the recommendations made by the personalized privacy assistants for the mobile apps, about three-quarters of motivations were adopted by users, and we also found that even after they adopted these recommendations and modified their settings based on a recommendation, even though we were trying very hard to get them to revisit these settings, they would not change them. That means that in this case, about 5.6 percent of those recommendations were later modified, despite nudges that we're sending them to revisit and rethink their settings.

Now, how do we know -- you might say perhaps they were just lazy, perhaps they ignored your nudges. Well, we had intentionally come up with recommendations that were ignoring a number of other settings. And, so, the nudges also covered settings that we had not covered in the recommendations. And those settings, users were actually modifying. So, we know that they were actually truly engaged. And, so, this suggests to us that these recommendations are actually pretty close to how people feel about these issues.

292

And, so, we strongly believe that this is the way to go for mobile apps. The question is, could we go one step further and could we generalize this to IoT? And, so, we have started to work in this area. The vision here is that you would extend this to deal with smart spaces.

And, so, what we are doing right now is we are building an infrastructure where owners of different resources, resources that are going to be sensing different aspects of your behavior -- cameras, location, present sensors and the like -- those resources have to be defined in the register by the owners, the people who own these various resources.

You know, if you enter a room like this, there are actually a number of different people who might potentially have deployed different resources already today that collect some of your information. For instance, it could be the case -- I hope it's not the case -- but it could be the case that the WiFi routers in this room perhaps collect your information.

These WiFi routers are not necessarily owned by the people who operate the building. Perhaps they are owned by the FTC, perhaps they are owned by a third party, I don't know, and perhaps it's better not to ask. But on the other hand, the HVAC system in this building

293

might be owned by an entirely different entity, and that HVAC might be collecting information, too.

So, I think that the owners of these resources should be able to very simply declare where these resources are deployed and what information these resources collect and all the other sorts of attributes that you would ideally want to see in a privacy policy.

So, we're developing an infrastructure where through a series of dropdown menus, people can specify different elements of their resources without requiring them to do any programing and look at what it takes to turn these -- this information into machine-readable privacy policies.

The idea is that users then, with their personalized privacy assistants, would be able to enter the space, discover relevant resources. Their assistants would determine, based on their expectations and their preferences, what, if anything, they need to be warned about or informed about. And if there happens to be settings in the ideal world, we would like these personalized privacy assistants one day to also configure these settings. We're not there yet, but that's effectively what we're aiming for. So, this is roughly how this is hopefully going to work one day.

So, let my try to quickly recap and also make

294

some connections with public policy in this space. So, we truly believe that this approach to effectively leveraging machine learning, in particular, building personalized models of people's privacy preferences and expectations, is one way of making notice and choice practical, right?

Today, the number of systems that you are encountering, especially in the IoT context, is just way too great for anyone to imagine that users are going to be able to read policies or configure settings. There is really a need to help users and to really do so by, number one, building models of what they care about; how they feel about different sets of issues; try to effectively alleviate burden in that context; and also make it much easier for the various owners of different elements of the infrastructure in the IoT context to participate within this infrastructure.

So, as was pointed out by Sarthak, I think, in the first presentation on this panel, one of the challenges of IoT is a diversity of players. If you think about the way you interface with a fixed Internet, most of your interactions are mediated by the browser, and so it's efficient in principle to just configure settings in your browser. On the mobile web, by and large, the cell phone mediates your interaction,

295

Android, and so it's sufficient to configure a number of settings at that level.

In IoT, it's a very different story, right, where you have a number of different players that might contribute different elements of the infrastructure. Many of these players might also be smaller entities don't have the sophistication that Google or Microsoft or Facebook might have. And so we really need to move towards an open environment, with open APIs, where effectively people will expose settings that will enable one, through personalized privacy assistants or equivalent technology, to effectively configure many settings on behalf of the user. And so that's really our vision in this space.

You can think of two different ways of deploying this personalized privacy assistant technology. One is to effectively rely on companies like Google or Facebook, each one of them potentially developing its own personalized privacy assistant, building models of the users. You can imagine also some potential tensions or potential conflicts of interest when it comes to theories that we have to come up with based on various guarantees.

Or, you could imagine a more ambitious effort where you might say, well, after all, there are some

296

interesting correlations between the way you feel about your settings on mobile apps, when it comes to sharing information with mobile apps, and perhaps your settings on Facebook and perhaps your settings in your browser.

And, so, rather than asking you these five or ten questions in each one of these environments in order to determine what your privacy preferences are, how about just asking you these questions perhaps just once and using your personalized privacy assistant, that cuts across all these different environments, interact with these open APIs to effectively configure many of these things on your behalf.

So, that's our vision in this space. It's not guaranteed that these APIs will be made open. In fact, today, they are not. They are very much part of the strategy that some of these larger entities have when it comes to building their own systems, but we would like to effectively build an effort towards perhaps convincing these larger players that they would all benefit from opening up these APIs, and perhaps people will ask me questions later on so I get to say more about this, but I think I have run out of time.

So, thank you very much.

(Applause.)

MR. ALVA: We'll conclude today with our final discussion of the day. Unlike previous sessions that have focused mostly on privacy, this session has focused on security and usability, research as it relates to privacy.

Sarthak discussed security issues related to IoT devices, and how they may affect privacy in the home. Vitaly presented on ad libraries and how the lack of tailored security controls in some contexts could result in disclosure of users' information through shared external storage.

For usability, Florian shared about an entire line of research going on around using machine learning, crowdsourcing, and other methods to make privacy policies more usable and for consumers, for businesses, as well as maybe for regulators. Finally, Norman presented new ways to understand and manage users' privacy expectations through personal privacy assistants.

Overall, this session has provided some new views into different strands of privacy research to consider. Now we will add to the policy conversation through our conversation here.

I want to welcome Geoffrey Manne, the Executive Director of the International Center for Law and Economics, as well as its founder, and Davi

Ottenheimer, who holds many hats in the security community, including authoring a book on big data security.

Geoffrey and Davi will provide some thoughts on this session as it relates to privacy for a few minutes each, and we will start there.

Geoff?

MR. MANNE: Thanks, Aaron. So, I thought the papers presented some really interesting things, and as did the papers throughout the day, and since this is the last session and I have you here, I am going to talk a little bit more broadly at first, anyway, than just about the papers today, but in a way that's consistent with what Aaron was saying, which is to say that the papers are interesting, there's some really important stuff here, but as is so often the case, the problem is deriving the appropriate policy implications from it.

One of the things I would say is that it's a little bit unfortunate, we don't have more economists and engineers talking to each other. As you might have gathered from the last panel, an economist will tell you that merely identifying a problem isn't a sufficient basis for regulating to solve it, nor does the existence of a possible solution mean that that solution should be mandated.

We really need to identify real harms, rather than just inferring them, as James Cooper pointed out earlier, and we need to give some thought to self-help and reputation and competition as solutions before we start to intervene.

Now, it is certainly something in the nature of a conference like this, and for that matter, the kinds of papers that people are writing, because journals don't publish papers saying there's nothing wrong. They publish papers saying, you know, there's a problem, and perhaps suggesting solutions to them.

So, we've talked all day about privacy risks, biases in data, bad outcomes, problems, but we haven't talked enough about beneficial uses that these things may enable. So, deriving policy prescriptions from these sort of lopsided discussions is really perilous.

Now, there is an additional problem that we have in this forum as well, which is that the FTC has a tendency to find the justification for enforcement decisions in the things that are mentioned at workshops just like these. So, that makes it doubly risky to be talking even about these things without pointing out that there are important benefits here and that the costs may not be as dramatic as it seems because we're presenting these papers describing them.

So, to think about the potential vulnerabilities that we talked about on this panel, the question to me becomes should they lead the FTC to any kind of enforcement if companies don't engage in the type of security that was recommended in some places or even any security at all?

And, again, this is an FTC workshop, so counselors out there are actually going to have to wonder if their companies are now on notice and if the very selection of papers for presentation here perhaps indicates anything about the FTC's enforcement agenda.

But here's the thing, having a possible vulnerability and acting unfairly under Section 5 are not the same thing. And, by the way, that's essentially, I think, the holding in the ALJ's decision against the FTC in the LabMD case.

Also, in terms of the desirability of enforcement, I think it's important to note that a couple of papers in this session and elsewhere throughout the day have suggested either that self-help is or can be working -- Norman's paper most obviously and immediately suggested a version of that -- or that despite the potentiality of all of these problems, something is actually preventing these vulnerabilities from being dramatically exploited.

Self-help has direct legal implications, say, for a deception claim, where it matters if it's available, but both self-help and the limited exploitation of risks are important in the economic calculus of the desirability of enforcement.

So, I want to end really quickly by saying I have more specific questions and comments about the papers when we discuss, but overall, I would just like to say that I think that last point is an area in which we're lacking in research, and I would like to see significantly more research on the implications of the availability of self-help, and what are the incentives for consumers themselves?

We've spent all of our time talking about the incentives of firms and the implications of legal liability on firms, but what about the consumers themselves? What about self-help? And how does and should the FTC take account of those?

MR. ALVA: Thanks.

Davi?

MR. OTTENHEIMER: All right. Well, I feel like somebody has given me a big basket of balls to juggle here at the end of the day. I will try to make sense of it all, a little bit of a show.

Teeing off on what Geoff just said, the idea

that there are these experiences we can have and we can learn from and there are these things we can discover through hard science is a fair split, and I will apply it now to the talks we heard today, the four talks.

I think that goes back to the question, should you study computer science or should you study social science? Should you have an applied approach to risk or should you have an academic approach? And a lot of times people forget that there's something in the middle. There is a fair balance between the two.

So, it was interesting to me to hear the first speaker talk about one end of the spectrum, which is essentially unit tests of these devices, these IoT devices; and then the second speaker took us through an integration test, a scenario asking what are these devices like in the wild? Let's look at how they're used by people, the economics, essentially the social science of how they're used. So, those are two ends of the spectrum, essentially.

And so then the third and fourth speakers brought in the middle ground, where you have somebody saying, well, maybe we can use this analytic exercise to help people make small rational decisions, right? So, you reduce the decision set criteria so people can choose from something realistic. So, you're not forcing

people to make big analytic analysis; it's really small. And that's kind of the two ends that I see.

And then the fourth speaker, even more interestingly, has a shared model, where not only are you making things easier to decide, accuracy and choice, but you're encouraging, nudging people. So, you're bringing an economic model towards the middle, towards simpler decisions with nudges. So, that's kind of how I see the four put together.

And I guess I have a ton of questions for all the speakers, but we don't have that much time, so I'll give it back.

MR. ALVA: Thanks.

So, I wanted to ask -- since we're running out of time, I wanted to ask a general question across all of the presenters. If there is one policy message that you think currently your research is engaging in, as you discussed in your presentations, but is lacking in technical measures that would actually help you implement the policy goal you would like to see, what are those shortcomings and how are you or would you like those shortcomings addressed?

And this question is open to any of the presenters.

MR. SADEH: Okay, it's a tough one. Clearly,

one has to be realistic about what can be done and how much room for maneuver I guess the FTC has in this space, but I suspect that the FTC can play a role in bringing together key stakeholders and encouraging dialogues.

And so, for instance, the issue that I was alluding to at the end of my talk, for instance, in terms of opening APIs, clearly this will never be something that, you know, one would ever be able to mandate, but perhaps efforts can be encouraged by bringing together key stakeholders.

At the end of the day, when privacy is prevented the right way and when people are looking at this rationally, everyone can benefit from better privacy, including vendors that, you know, are sometimes presented as if they didn't care about privacy. If you look, for instance, at what is happening today in the mobile space, it's very clear that everyone has come to realize that they don't want to be seen as the people who don't care about privacy, and that creates strong incentives for them to rethink the way in which they've been approaching some decisions in that space.

So, I think that perhaps the FTC can, on the one hand, continue to do what it's been doing very well, I believe, which is to encourage best practices, that it

305

has done, for instance, for mobile apps, as it has done more recently when it comes to IoT security, and perhaps also convening meetings and encouraging efforts where people look at opportunities for perhaps developing common standards, not trying to impose any standards, and, you know, standards are very challenging and very tricky efforts, but at least trying to bring together key stakeholders and trying to get them to think about where they've got effectively common interests and where they might benefit from perhaps developing some open APIs.

MR. SHMATIKOV: I think transparency is very important. Better understanding and better disclosure of how information is collected and shared between various players in the picture is crucially important, because what we have in mobile space today is these old permission models. They capture something about security of these devices. They capture virtually nothing about privacy.

There is a lot of information collection and sharing and information used between all kinds of artists -- platform operators, ad libraries, ad builders, advertisers -- that simply exists outside the existing permission models that a lot of privacy work focuses on. So, to the extent FTC can help shed light

306

on this and ask for more disclosure of information collection practices and information flows in this massive mobile ecosystem, that would be an extremely useful service, because that is not happening today.

MR. GROVER: So, I would totally agree with that. Transparency is the big issue, and maybe, like, the FTC can, in terms of IoT devices or mobile apps, say unless you follow a certain set of policies, we won't allow you to sell these devices to others.

But the problem comes back to a point Norman mentioned, that in terms of IoT devices, there are not really open APIs, and, I mean, who basically sits there and looks at all of this? Who does the analysis when you don't really have access to the code? And when the software and the hardware are basically integrated, you don't have choices in case you feel like something is wrong in the software. You aren't really able to replace it with something else.

So, transparency is the main issue, and it should be encouraged, but, quite frankly speaking, I don't know how to go about it. That's the problem.

MR. MANNE: So, but, you know, one of the things -- I mean, there's always tradeoffs. It may not surprise you all that Leonard and I wrote a paper called "The Cost of Disclosure," so I agree transparency tends

307

to be a good way of achieving these things, but it's not costless. As Norman had on his last slide, he pointed out that if we have open APIs, we're going to be empowering the groups that are collecting these massive amount of information through open APIs with an enormous amount of information that creates perhaps even greater vulnerabilities than the ones we're protecting. So -- and there may be other examples like that, too.

So, my question really is, before we settle on transparency, even, as the right sort of, you know, optimal kind of solution here, we should be aware that there are costs to that as well, and, again, potentially we're creating more risks than we're solving.

MR. OTTENHEIMER: That's right. I put it as, transparency to whom? So, transparency -- you know, you're building trust relationships, so it's transparency to somebody that you essentially trust to give you the right answer, and given that they have the information. So, I've done audits over 20 years, and I can tell you, just being able to see into something doesn't mean you're in the position to make the decision on it, which is sort of what the presentations were about to some degree.

We give the people the information. The people are positioned in a way that they can’t digest it

308

because they don't have the analytic capability at the time they're given the information. That's why I'm saying balance. If you take the sort of unit tests, you can say that's inadequate because you have a compliance checklist. If you take the environmental or the integration test, you can say, well, that's not fair because that's not a typical use case.

So, somewhere in the middle is proper use of the device prepared for a use case, and that's I think a good fit.

MR. SCHAUB: So, I think concerning transparency, an interesting point to think about is essentially the privacy policies we have right now, they're not written for users, and they're not meant to provide transparency for users, and we need to realize this and I think this needs to be more clearer in regulation as well, that if we want to inform users and achieve transparency for users, then we need to come up with user-facing notices that are actually made for users.

That could include requiring user evaluation of those notices. Are they actually effective at communicating what they are supposed to communicate? And we've been doing a lot of those studies at CMU, and we find most notices are not effective, and it's really hard to design an effective notice.

MR. OTTENHEIMER: Here's the interesting counterpoint. The more information that becomes available, the more behavior changes. So, if you actually -- I could show you exploits, for example, to your model that show as you get this in position where your machine-learning algorithms are working and you're actually getting the answers you want, the people writing the policies will change them just so you can't see them anymore.

So, the transparency has been to be in concert with the right trust model where people want it to be shown in the way that it's comfortable for them; otherwise, they adapt and your transparency backfires.

MR. ALVA: Norman, did you want to address the transparency --

MR. SADEH: Well, I would like to respond to the last comment. So, I think it's clear that privacy is an arms race, right? So, I think that Davi and I worked with Florian on the project that he described, but the day that site operators, for instance, start modifying their policy based on our technology, because of the success of our technology, will be a very good day. We're not quite there yet.

If that day happens, we will actually have the

ability to probably identify that, and that might potentially be something that the FTC would be interested in. Whether the FTC would actually be able to do very much about it or not, I'm not sufficiently versed into the legal ramifications of that, but I suspect that it would have something to say if you can establish effectively a pattern where, once you effectively are able to capture some practices that are not necessarily putting these companies in good light, they start modifying the way in which they're presenting the text, I suspect that's something -- you know, something that could potentially be done.

MR. SCHAUB: And it's also quite imaginable that it could go the other way, so that companies actually improve their language to be better presented by these independent mechanisms, and we have conversations with many different companies that say they would actually welcome having such kind of technology out there, because they do invest a lot of money and time in having privacy policies that are descriptive.

But it's basically in vain at the moment because this information is not used, and it's not clear to users that this is the case. So, I think this could go both ways, but it's going to be interesting to see

how it plays out.

MR. MANNE: I mean, my sense would be that the primary reason for the unintelligibility of existing disclosures of privacy policies is the legal risk and for that matter even regulatory enforcement. So, there's -- you know, if we're going to identify -- if the problem is we don't have disclosures that actually inform the users, then we haven't really -- to me, we've largely identified a really important disconnect between how we're regulating and, you know, the power of users, which goes to the point I was making before, which is I really liked what you were describing.

The sort of app that you guys created seems to me like it has, you know, amazing potential, but once we have something like that, think of what that does to the need for additional forms of regulation. I mean, you might still need some deception regulation, but you've done a really good job now of actually giving users what they want, and because users are so heterogeneous, because types of data are so heterogeneous -- I think, by the way, on your paper, there's a big difference between an email address being accessible and the content of a communication even with a computer device.

A real problem with overgeneralization -- and this may be actually partly reflected in the bad privacy

policies -- a problem with sort of an overgeneral response, like a network-level response to the problem you were identifying, is that -- well, I don't know enough about the engineering, but at first cut I would say it doesn't differentiate; it just imposes a single policy on everyone, regardless. That's really unlikely to be the right outcome, but that is a problem with, you know, sort of the more -- the blunt -- relatively blunt policy tools that we have.

So, you know, again, I think there's real value in empowering users as long as that leads to a reduction in the incentive of these more blunt tools to come in.

MR. ALVA: So, we have about 45 seconds left. I wanted to ask the presenters, if you had the ideal privacy agenda in your research, what would it be -- in one or two sentences -- going forward?

MR. SADEH: I think I have outlined our agenda and there were three presentations today, so I strongly believe in this vision of personalized privacy assistants. It's clearly not something where we are entirely there yet, but we have some very promising results.

If I can take another 30 seconds --

MR. ALVA: No, sorry. But thank you, though.

MR. SCHAUB: So, I think what's also important, when we are starting to look at this, providing information and integrating these dialogues into the users, into action flow. So, rather than having a privacy notice or privacy policy somewhere else, when the user interacts with it, make it part of the interaction.

The mobile platform developers are doing a good job doing this already or starting to do this already. You have those just-in-time dialogues that pop up, and they don't disrupt the interaction. They actually help it, and they actually encourage the app developers to build dialogues around it that tell you why this notification is going to pop up and why they want your location. So, that's great. I think that's a good direction to go in, to think -- we're doing quite interesting research on these things.

MR. SADEH: And we're not biased.

MR. ALVA: I'll stop you there. I encourage the audience to ask Vitaly and Sarthak after this, but I wanted to conclude.

So, the FTC's new chief technologist started on Monday, and so I wanted to welcome Lorrie Faith Cranor from the FTC, and we also thank Carnegie Mellon for allowing her time on leave for her to be here with

us.

MS. CRANOR: Thank you. I will keep my remarks brief, since we are over time.

First of all, I wanted to thank all of the FTC staff who did such a wonderful job organizing this event. Can we give them a big round of applause?

(Applause.)

MS. CRANOR: So, this is my fourth day, so I had nothing to do with it, but these guys did a really great job. I also want to thank all of you for coming and for participating.

A few notes on some things that I heard throughout the day. It was a lot to absorb, and I was busy scrolling notes and trying to synthesize what I heard.

So, I think some of the key areas that I heard, there's a lot of really interesting empirical research that is being done and some of the areas that it's being done in that we heard about. We heard about survey and interview research about what consumers understand and especially what they expect and what they desire.

We also saw that some of this research is then being used to find ways to actually assist consumers, figuring out ways to reduce the number of notices that

they need to see, and configure their settings automatically.

A question that came up in almost every panel, I think, was a question about how we can make transparency and notice and choice more effective. We heard over and over again how ineffective it seemed to be, and we heard some ways forward, some paths to maybe making it more effective.

We also heard about measurement research that looked at a variety of things. We heard about measurements on the extent that people are being tracked and what technologies are tracking them. We also heard about statistical and machine-learning research to understand how algorithms impact users.

And our speakers observed that in order to have algorithmic transparency, it's not enough to just know what the algorithms are, because that doesn't really tell us very much. What we need is systems that help interpret the results of the algorithms and show us the impact of those algorithms.

We saw some research that the models have investigated the impacts of different approaches to privacy protection and could help shed light on the effectiveness of different approaches. We saw research to understand the impact of incentives and approaches to

cyber security.

We also saw that many of the researchers who spoke here had developed some tools that had been very useful in their own research, and many of them had actually offered to make their tools available to other researchers who could also use them. And I think the community is developing a tremendous tool set that should enable a lot more research to happen going forward.

We also heard from research an eagerness to partner with companies to do empirical research. Some people noted that in order to do the research they wanted to do, they needed information that only the companies have, and so there was an invitation to partner with them.

So, those were kind of the highlights of what I heard today. I'll be very interested in hearing from all of you about what you found useful. We're also interested in getting feedback on this event. Should we do it again? If so, should we do it exactly the same way? What should we do differently? We would be very interested in hearing that from you.

One of the things that I would like to do while I'm at the FTC is to try to better bridge the gap between academic research and policymakers. And I think

317

the privacy area is an area where there's a real need to inform policymaking with research. And, so, as such, I look forward to continuing the discussions that we started here throughout the year.

Thank you.

(Applause.)

(Whereupon, at 5:43 p.m., the workshop was concluded.)

318

C E R T I F I C A T I O N   O F   R E P O R T E R

CASE TITLE: PRIVACYCON WORKSHOP
DATE: JANUARY 14, 2016

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the notes taken by me at the hearing on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: 1/20/16

SALLY JO QUADE

C E R T I F I C A T I O N   O F   P R O O F R E A D E R

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

SARA VANCE

S