FTC PrivacyCon 2017
January 12, 2017
Segment 1
Transcript

KRISTIN COHEN: Good morning, everyone I just want to welcome you all to our second Privacy Con my name is Kristin Cohen and I am an attorney in the Division of Privacy and Identity Protection at the Federal Trade Commission. My co-organizers for today's event are Pader Magee, also in the Division of Privacy and Identity Protection, and Justin Brookman, who is in the Office of Technology Research and Investigation. Before we get started, I just need to review just a few administrative details that you've heard at every one of our workshops. Please silence any mobile phones. If you need to use them during today's event, just please be respectful of the speakers and other audience members.

Please be aware if you leave this building at any point during the day, you will have to go back through security. So please leave enough time. The restrooms are right outside the auditorium. The plaza east cafeteria will be open today until 3:00. If you eat there, you do not need to go back through security. But please keep in mind that only water is permitted in the auditorium. Most of you received an FTC event lanyard, and we do reuse those, so please return them at the end of the day. If an emergency occurs that requires you to leave this conference room but remain in the building, please just listen to the announcements over the PA system.

And if you need to evacuate the building, please leave through the 7th Street exit, and after leaving the building turn left on 7th street, and across E street to the FTC emergency assembly area and remain there until instructed to return to the building. And of course if you notice any suspicious activity, please alert the building security. Please be advised that this event may be photographed and it is being webcast and recorded. By participating in this event, you are agreeing that your image and anything you say or submit may be posted indefinitely at ftc.gov and on one of our Commission's publicly available social media sites.

We are very happy to welcome those of you watching via the webcast. We will make the webcast and all of the workshop materials available online to create a lasting record for everyone interested in these issues. You'll notice that this year, we do have microphones set up, and we do hope to take audience questions. So we encourage you to think about your questions during the panels today and to line up if you have questions. We will also be taking questions via Twitter. So just tweet your questions using FTC with the hashtag Privacy Con.

And new this year, we are having a lunchtime poster session to showcase additional research. So we really encourage folks to get their lunch in the cafeteria. Head over to the rooms across the hall, and learn about great privacy research happening in our community. And lastly, I just want to give some thank you's. I want to thank the researchers and panelists for taking part today. We are very grateful for your time. Also this program would not be possible today without a lot of work from our FTC colleagues. First I'd like to thank those that assisted us in reviewing all of the terrific submissions we got. That includes Lori Cranor, our chief technologist, Aaron Alva, Joe Calandrino, Lerone Banks, Tina Yeung, Tim Daniel, Marc Luppino, Michael LeGower, and

Whitney Merrill. And we also want to thank those moderating panels today, including Mark Eichorn and Jessica Rich.

We want to thank those who are heading up our poster session, Whitney Merrill and Tina Young. And of course those who put this whole conference together, Fawn Buchard, Crystal Peters, and Bruce Jennings, alongside our paralegal support from Carry Davis, Jonathan [INAUDIBLE] Joseph Kennedy, David [INAUDIBLE] Jennifer [INAUDIBLE] Kethan Dahlberg, and Omar [INAUDIBLE] And support from our division of privacy-- I mean our division of consumer and business ed, Jessica Skretch, and from our office of public affairs, Nicole Jones. So thank you all for all your hard work today. And now it is my great honor to welcome the chairwoman of the Federal Trade Commission, Edith Ramirez, to give opening remarks.

EDITH RAMIREZ: Thank you very much Kristin, and good morning, everyone. And welcome to our second Privacy Con. When it comes to privacy, technology has always presented a challenge. How can we make the use of the tremendous benefits of technological innovation while ensuring that our privacy is protected? This has been true from the snap camera of Warren and Brandeis' time, to the drones of today. The last several decades have brought change at a breakneck pace. The rise of the personal computer in the 1980s, the internet in the 1990s, the smartphone in the 2000s. And this decade, the internet of things. This dizzying array of technological advances is only going to continue to grow.

Last week, I had the opportunity to be in Las Vegas at the consumer electronic show, and had the chance to walk the showroom floor. There were smart cars that use technologies to sense driver emotion and deploy sensory outputs like sound, scent, temperature and light, in an effort to promote mental awareness and potentially reduce incidents of road rage. There were organic light emitting diode TVs as thin as cell phones that are capable of controlling the light of each individual pixel. There were passenger carrying helicopter drones that can be used to transport organs for transplants. Drones that can fold up and fit inside your pocket, and others that are outfitted with connected virtual reality goggles that promise a whole new experience.

From the robotic vacuum cleaner that also serves as a mobile home security camera and an air humidifier, to the smart trash can that can scan barcodes of disposed items in order to build a shopping list of items that need to be replaced. Almost all of these technologies however will rely to varying degrees on the collection of consumer information. And data collection is growing exponentially. Experts estimate that by the year 2020, there will be a 4300% annual growth in the amount of data that's collected. Just around the corner are huge advances in artificial intelligence fueled in part by IoT data.

A recent White House report notes, for example, the data generated by artificial intelligence technology can enable enormous advances in health outcomes. It can improve traffic management technologies, resulting in efficiency, lower emissions, and energy savings. But if all of this innovation is going to achieve its potential, consumers need to be assured that the risks do not outweigh the benefits. And today I'd like to describe some of the risks of this new technological landscape, and how privacy is helping the Federal Trade Commission to address those challenges.

Some of the risks of these new technologies are similar to ones that we've seen before. For example, traffic management technologies might only prove useful if they use data that includes a person's geolocation information. Now, we've long recognized that geolocation information is sensitive, and should not be collected or used without a consumer's opt in consent. Risks of unauthorized exposure of geolocation information include , stalking revelation of political, health, and religious affiliations, and even burglary. As this example shows, the possibility of unexpected uses for information must be weighed against the benefits.

But in addition to some of these familiar challenges there are also new ones. One is the ever-growing number of actors that have a role in collecting, compiling, interpreting, and using data in a world that relies and operates on big data, IoT, and AI. There are consumer facing companies, a device manufacturer, a smart hub platform, or a publisher website or app. There are behind the scenes technology companies, software vendors that connect IoT products to the internet. And of course, the numerous analytics and advertising companies. This vast array of entities makes it difficult to provide consumers with informed choices. And this challenge is exacerbated when non-consumer facing entities increasingly handle consumer data. This also raises concerns about whether all of these actors are appropriately protecting the security of consumer's personal information.

Second, with the new technologies, privacy and security failures aren't simply about threats to personal information. They can also include threats to health and safety. Particularly in relation to certain health devices and connected cars. For instance, the failure of security of IoT devices, in particular the ease with which IoT devices can be recruited into vast botnets to be used in DDoS attacks, could pose substantial risks. To meaningfully thwart potential botnet armies, a significant majority of manufacturers would have to act collectively to improve security.

Third, by relying on algorithms based on big data techniques and machine learning, companies may disadvantage certain populations. As we note in our big data report issued last year, even large data sets may be missing information about certain populations. Such as those who have unequal access to technology, or are less involved in the formal economy. And big data analytics can reproduce existing patterns of discrimination, or reflect the widespread biases that exist in our society. For example, an algorithm that isolates attributes of good employees or good students, may simply be replicating biases that existed in previous hiring or admission decisions.

The only way to keep this balancing act in equilibrium is to earn and maintain consumer trust. And this is where the FTC comes in. So what do these emerging technological developments, and the challenges that they present, mean for the FTC? It means that we have to continue to be nimble and smart to keep Pace And we have to leverage our resources. At the FTC, research and data play a key role in helping to guide our work. This is precisely why Privacy Con is so important. The research this event generates directly informs three critical areas of our privacy and security agenda.

First, we use research-- both research presented at Privacy Con as well as other research-- to identify potential areas for investigation and enforcement. For instance, tech researchers brought to our attention the practices of InMobi and Turn, two companies that were the subject of recent FTC enforcement actions. In our action last year against InMobi, a mobile advertising network,

we alleged that the company tracked the locations of hundreds of millions of consumers without their knowledge or consent, and even when consumers set their privacy settings to deny access to their location. More recently in our case against Turn, we alleged the digital advertising company deceived consumers by tracking them online and through their mobile applications, even after consumers took steps to opt out of such tracking. We're grateful to the outside researchers who worked hard to identify and publicize these practices.

Second, Privacy Con provides data for our policy work, and helps identify areas where additional research is needed. For instance, we used OpenWPM, a tool developed by one of last year's presenters, that automates evaluation of privacy on websites in our recently published study about cross device tracking. In this study, OTEC staff looked to assess what information about cross device tracking is observable from the perspective of the end user, including through data flows and public disclosures. In particular, they looked at 100 popular sites on two different devices connected to the same IP address, to see what information was collected that could be used for cross device tracking.

Overall, staff detected a lot of data collection practices that could be used for cross device correlation. It was often not clear why the parties were sharing this information. But the sharing could be for cross device tracking. Or it could be for other purposes. But clearly a broad range of companies have the capacity to correlate user behavior across different devices that users own.

Staff then reviewed the privacy policies of the 100 sites which revealed, surprise, surprise, that the policy privacies were vague. In the vast majority of cases, it was unclear whether the site would share data for cross device tracking. As a result, it would be very challenging for even a very sophisticated user to determine how much cross device tracking is taking place. We think this type of research is incredibly helpful for informing industry, consumers, and policymakers about what's happening in the marketplace. And it was a tool that was presented at Privacy Con that let us do this.

Third, Privacy Con helps us to identify and develop solutions to the privacy and security challenges that we're seeing in the marketplace. For instance, this past year, we've heard about the harms that can result from IoT vulnerabilities. The hacking of vehicles that could place lives at risk, or of an insulin pump that raises significant safety concerns. We've also heard about the malicious use of the IoT botnet Mirai, that was used in DDoS attacks around the world last fall. It's never been more clear that we have to secure this software and the devices supporting our digital lives.

To further these efforts, last week we announced an IoT security challenge. We're going to be giving prize money to anyone who can create a tool to help consumers quickly identify security vulnerabilities, and push out updates to address those vulnerabilities. And we're going to give bonus points to tools that can prompt consumers to change default passwords. We think this important initiative will draw attention to IoT security problems and facilitate solutions that consumers can use.

Now, as I think about the not too distant future, where robotics, AI, and more sophisticated IoT developments reign supreme, Privacy Con will continue to help bridge the gap between the

academic, tech, and policy worlds. We'll continue to learn from this event to enhance our understanding of consumer expectations, to inform how practices in this dynamic economy align with those expectations, and how devices and data can be secured in this new landscape. Today's forum, which is going to feature discussions on IoT and big data, mobile privacy, consumer privacy expectations, online behavioral advertising, and information security, will undoubtedly provide valuable insight on these and other issues. And it will help the FTC address emerging privacy and security challenges in a complex, dynamic marketplace.

So just to close I really want to thank our panelists for sharing their expertise, and all of you for joining us as we seek to study these important issues. And I really also want to take this opportunity to thank the FCC staff who organized today's event. And in particular Kristin Cohen, Pader Magee, Justin Brookman. And Mark Eichorn. I also want acknowledge Lorrie Cranor, our chief technologist who unfortunately will be leaving us very soon. But she's been an incredible addition and asset to us at the agency, so thank you very much Lorrie, for everything that you've done. So thank, you for being here, and now I want to turn the floor over to I believe Peder Magee. Thank you.

PADER MAGEE: Good morning. Thank you very much Chairman Ramirez, and thanks to the rest of the audience for coming out to Privacy Con Two. My name is Pader Magee, I'm an attorney in the FTCs Division of Identity-- of Privacy and Identity Protection and I'm going to be moderating our first session this morning, which I think is a good segue from the chairwoman's remarks. The first session is entitled, Internet of Things and Big Data. And I'd like to ask my panelists to come on up if they would. We have five researchers presenting on four separate and very interesting projects. They're each going to have 10 minutes to discuss their work and then we'll have a discussion period.

To get that started, I'll pose a few questions and then open up for audience questions. If you have something you'd like to ask, please line up behind the microphones after all the panelists have finished presenting their research, and then we'll take your questions. So let me start out by introducing Noah Apthorpe and Dillon Reisman from Princeton's Center for Information and Technology Policy. If you'd like to go to the podium and--

[INAUDIBLE]

PADER MAGEE: Thank you.

DILLON REISMAN: So, hi everyone. We'd like to thank the FTC for having us here today. My name is Dillon Reisman, I am a researcher with the Center for Information Technology Policy at Princeton. This is Noah Apthorpe, he is a Ph.D. Candidate with the Center, and we're here presenting research with Nick-- done with Nick Feamster on smart homes, privacy vulnerabilities of an encrypted internet of things traffic. First though, we'd like to thank a few people. Thanks to our Arvin Rhiannon, Nina Taft, and two undergraduates, Alex Yu and Kyle Berger, who really contributed to our research.

But first we should really talk. What is a smart home, even? What are smart home devices? The internet of things is a grossly overloaded term. So what we're talking about here today is a smart

home. A smart home is a home in which devices-- or rather, in which traditionally analog appliances have been replaced by computers. You're probably already familiar with examples of this. Maybe some of you have a Nest thermostat for instance. Thermostats adjust the temperature in your home. A Nest thermostat is in some sense nothing more than a computer in the shape of a thermostat. It does the same thing, but it does it intelligently. It learns your preferences.

But there's also a second category of devices we're considering here today, and those are appliances that are brand new. They don't have a non-digital analog. An example of that would be an Amazon Echo. A smart personal assistant. You can ask it questions, it can answer them. These devices have something in common, or rather two things in common. They add new types of sensors into your environment of your home. So for instance, like a microphone, a camera, a motion detector, that's one thing. And second, they all-- well, at least the ones we studied and most of them, rely on cloud services for their basic functionality.

So we'll use the example of an Amazon Echo. And we're going to talk about a model of network traffic, how these devices communicate with the internet and why that's a privacy vulnerability. So let's use the Amazon Echo as an example. Let's say you ask it a question. The Amazon Echo connects to your home Wi-Fi, much like your smartphone does, so it transmits your question-- Say, Alexa, what is the weather?-- over Wi-Fi, and so radio, and it goes to your Wi-Fi router. Your router then communicates via your internet service provider across the internet backbone, where it then reaches an Amazon server eventually, where it is then-- the question is answered, and your response comes back to you. And it's excellent, right?

Well there could seem to be an obvious privacy issue with this. All those sensors, the microphones, the cameras, at some point they're transmitting some representation of your physical environment in your home over the internet-- over the internet via your internet service provider. And this can indicate things like your physical presence in your home, your sleeping habits, your media consumption habits, the questions you ask Alexa. Fortunately it's not quite so simple because a lot of the devices-- a lot of the best developed devices-- use best practices. A lot of them use encryption. So say the Amazon Echo, you ask the question. Your internet service provider can't read the question. It is encrypted by the device, sent via your internet service provider to an Amazon data center, and only Amazon can really read that question. So we're done here and we have nothing else to talk about.

Except we do have things to talk about. What our research found is that smart home internet traffic patterns, the rates of the communication that the device takes over the internet, can reveal private in-home behaviors even when that traffic is encrypted. In computer science you would call this a form of side channel attack. And Noah is going to describe how that works right now.

NOAH APTHORPE: Great. Thank you Don. So in order to investigate this side channel threat, we set up a smart home laboratory where we went out and purchased a number of commercially available devices. Hook them up to a Wi-Fi router that we had configured to record all traffic to and from the devices. And then interacted with them in a way in which many consumers would on a day to day basis. This gave us a perspective of the internet traffic that was being generated by the devices, similar to what might be seen by a last mile provider, service provider. Or potentially what might be seen by a wife by eavesdropper, say someone in the apartment next

door who is able to overhear your Wi-Fi transmissions. Once we had recordings of the device traffic, we were then able to analyze them. And I'm going to walk you through a three step process by which an observer who is able to record traffic patterns from your devices could use to infer your in-home behaviors.

So to start with, this is an example of a plot showing traffic rates over time for a particular device. On the vertical axis, we have bytes per second traffic rate, and says over about a 12 hour period. Just being able to see this traffic pattern doesn't give you that much information, because you don't know what type of device generated it, so you're not able to make any further interferences. How using additional information though, it's possible to identify the device. In this case it was from a IoT sleep monitor, one that you would put by the side of your bed and clip to your pillow case, and would track your sleeping patterns. We were able to show that you could identify this device by looking at the DNS queries that it made.

DNS is a protocol which internet devices used to map from human readable URLs to machine readable IP addresses. DNS is typically not encrypted, and sent in the clear. And DNS requests from IoT devices often have the name of the device, or at least the device manufacturer, included in the domain name. So once you have that, it's often associate forward to map a set of traffic patterns back to a specific device. Even if you don't have access to DNS queries, we showed that with some fairly straightforward supervised machine learning, you can also identify devices by looking at features including mean standard deviation send rates and autocorrelations.

So now once you are able to see that this is indeed a sleep monitor, you are able to make some further inferences. The limited behavior of the sleep monitor is what enables that. So you'll notice that in this particular traffic pattern there are three spikes that stand out as being distinct from the background. You might think that maybe those spikes correspond to something salient that the consumer is doing. And indeed we found that those were the three times where the consumer did something related to their sleeping patterns. They either went to bed, they woke up temporarily in the middle of the night, or they got out of bed in the morning. If you were a adversary, you could purchase a sleep monitor like this of your own. You could do this experiment. And you could know that when you saw traffic spikes of this nature coming from consumer traffic, that they most likely correlate to the consumers behavior in the same way.

We didn't look at just the sleep monitor, of course. We examined other devices. A security camera which had-- we examined two different states. One in which the user is actively watching the video stream, and another in which the cameras is sort of locally monitoring the video and if it detects motion it will take a snapshot and then notify the user. Here we saw that depending on which state the camera is in, there's a very distinct difference in the amount of traffic that's being sent. So a outside observer, once they've identified the camera, they'll know whether or not anyone is actually actively watching the video stream, or whether it's not currently being monitored.

Diving in deeper to the motion detection state, we were able to show that when a motion event occurs-- say, a pet moves in the living room of your home-- it will take a snapshot of the pet, upload it to the cloud, and give you notification that motion occurred. That upload event can be seen as a spike in traffic from the device, and therefore infer an outside observer that there's

something moving inside your home which up until now, previously would have been more private information. This same style of attack also continued across other devices that we examined. This is a smart outlet, which you would plug into one of your outlets. And then you could use your smartphone and the corresponding application to turn whatever appliance is plugged into this outlet on or off. You could also turn it on or off using a physical button on the outlet. It's like a light switch. But we showed that regardless of how you interacted with this device, whenever it changed power states, it would notify the cloud, and would therefore be visible in terms of its traffic patterns.

And finally here, the same was true for an intelligent personal assistant. In this case the Amazon Echo. When you ask the Echo a question, of course it needs to go query the Amazon servers to get the response. But then you end up with this trace which can indicate when you're in your house interacting with your devices. Of course, the devices that we looked at are by no means representative of the scope of IoT. And we're interested in going forward with this to see just how wide this type of side channel threat is. We're also looking to see of course, what can we do to prevent this? We're hoping there are technical privacy preservation methods we can use to address this sort of metadata threat, perhaps using something like probabilistic privacy injection to reduce the confidence with which an adversary might be able to infer consumer behaviors.

I'd like to conclude by saying that our study indicates that just using encryption alone is not adequate for privacy protection for smart home IoT devices. And instead it may be beneficial to think of this as the floor, rather than a ceiling, of something that we need to strive for. Here's a link to a paper fully describing the research that we've talked about here today, and thank you.

PEDER MAGEE: Thanks very much. Next we have Aleksandra Korolova from University of Southern California.

ALEKSANDRA KOROLOVA: Hi. It's an honor to be here today to talk to you about cross app tracking via nearby Bluetooth devices. This is joint work with [INAUDIBLE] and we're both from University of Southern California. So many of our devices are now smart. We have everything from cooking pots to fitness trackers in our homes and our offices that are Bluetooth enabled. And most of us have more than one app on our phone to control these Bluetooth enabled devices and to analyze the data that they collect about us. But what are the privacy risks that having so many smart Bluetooth enabled devices around us pose to our lives?

Today I'm going to tell you about two such-- two such privacy risks. One of them is profiling. The idea that using the information received from smart Bluetooth enabled devices, applications on your phone can learn new things about you that you may not even suspect. The other risk is tracking. And here the idea is that apps on your phone can collaborate in order to create a cookie like, or identifier like, identifier that can track you across different apps of your phone. Moreover, both of these privacy risks can happen without user knowledge or meaningful user control.

Before I explain how these risks arise, let me tell you a little bit about the Bluetooth low energy protocol that is one of the culprits of the risks. The Bluetooth low energy protocol, introduced in 2010 by the Bluetooth special interest group, is designed to enable a new set of smart devices.

Each Bluetooth device announces its presence via advertising packets transmitted through the radio channels. What these advertising packets contain varies from device to device. All of them though, contain that device's Mac address, as well as optionally things such as the name of the device, manufacturer information of the device, and the different services that this device supports.

So for example, if I'm wearing a fitness band right now, many of you who have phones can receive information that the fitness band is transmitting. How exactly can you receive this information? Well any app on your phone can request the information that is being transmitted right now in this room via the APIs provided by the operating systems. And they can request this infor-- as long as the app is running, as long as you're using it, they can request this information unlimited number of times. And since the range of the devices is quite high, you would be getting information from many, many devices in this room, or when you are at home, or when you're in your office.

How does this lead to profiling? Very simple. As I mentioned, many of the devices transmit their names. So for example, imagine I'm playing a game on my phone every evening when I'm at home. The game developer can have access to the names of all the Bluetooth enabled devices that I have in my home. And the names themselves can be revealing and useful for profiling. If I have a fitness band that they named and called Aleksandra Korlova's fitbit, then the developer can know my name. If I have a TV that is automatically named Samsung 9 series 65, then the app developer can learn something about the disposable income I have to spend on large fancy TVs. If one of the devices that gets observed very frequently is named Mamaroo, then they can learn that either I am myself a parent, or I live in a household where there's a young infant.

How do Bluetooth enabled devices create the tracking opportunities for the apps? Well, it's a little bit more complicated, but also not that complex. Imagine apps want to exchange information about their users, and they want to find a persistent identifier so that app one and app two can figure out which user is the same. What they can do is, they can, every time the app is used, they can observe what are the Bluetooth enabled device Mac addresses, manufacturer names, names, in the vicinity. And once they've collected a long list for each user of what the devices in their vicinity are, then the apps can simply match the users based on the overlapping devices. So if for user one of app one and user x or app two have a lot of devices in common, then it's very likely that they're actually one and the same person.

Are these theoretical possibilities, or is this kind of tracking and profiling actually feasible in practice? Well, this is what we decided to find out. And to do that, we recruited 70 volunteers among USC community, to install our app that would collect the data from Bluetooth enabled devices and share it with us, so that we could analyze it. And what we found, is that this kind of profiling and tracking is really truly feasible. So among the data collected, we observed over 1,000 distinct device names, many of those containing actually names of the people and other information from which you could learn things. And we also observed that the information is rich enough that two apps that they used as infrequently as once a day, could uniquely identify 60% of our 70 users.

Is this tracking happening? Are app developers already using this? It's a bit hard to answer, because the operating systems don't actually make it easy for you to analyze what is happening. And even if it was easy, you can't know if the apps are just collecting data, or making inferences based on it. But what we do know is that there is nothing currently to stop app developers from engaging in this kind of tracking and profiling. Well, the next question is, what can individuals do to prevent it? Well, not very much. They can turn off the Bluetooth on their phones, but that's about it. And if we are actually interested in taking advantage of the smart device revolution, turning off the Bluetooth on your phone is not really a functional solution.

To conclude, what we've done is we've identified a new type of privacy risks, profiling and tracking, that can happen using the nearby Bluetooth enabled devices. And we've shown that this new type of attack is feasible. And in the course of this study, we've also discovered some of the ways that the Bluetooth special interest group, the mobile operating system developers such as Apple and Google, and also the device manufacturers, can change a little bit what is being done, in order to protect the privacy of the individuals better, while not losing that promise of the Bluetooth functionality. I'm especially happy to talk about this at the FTC Privacy Con, because from my perspective, I think changes are needed in how this is handled, and maybe FTC can be the catalyst for these changes to happen. Thank you very much.

PEDER MAGEE: Thanks Alexandra, that was great. Now we're going to hear from Maria Rerecich from Consumer Reports.

MARIA RERECICH: This is-- that one there? OK. So, hi. My name is Maria Rerecich, and I'm the director of electronics testing at Consumer Reports. Today I'd like to talk to you about the evaluating of products for privacy, security, and data practices. Now I come at this work from a perspective of technical testing. I'm an electrical engineer. I've been at Consumer Reports for almost four years now. And prior to that, I worked in the semiconductor industry. So this type of research, which incorporates testing rigor into work on data privacy, is really exciting to be working on for me. [INAUDIBLE] next thing. Oops. That was me. This.

Now at Consumer Reports, we do comparative testing of products. And we do that in order to give consumers an informed choice. We test by using defined protocols and procedures which we develop from how consumers use a product, but also refer to industry benchmarks and standards where they exist. In the absence of industry benchmarks or standards, we would-- or if we feel they're not sufficiently consumer focused, we may develop our own targets and goals for performance of products. Now by generating ratings based on this type of testing, by being able to identify better and worse product performance, we have seen that companies will compete to have better products. And in some cases industry standards or regulations will change to set the bar higher.

So it's this perspective that we have of the impact that testing can have, that we bring to this process when we consider the development of a new digital standard. So why are we talking about developing a digital standard? So we're seeing in more and more of the products we test, the TVs, the refrigerators, the thermostats, the cars, a shift from hardware to software. We've been talking about it, the Internet of Things. We call them connected devices. We call them smart appliances. And this is uncharted territory for consumers. Consumers don't know what's

happening to their data in this type of environment and with these type of products. There is no easy way for the consumer to navigate the privacy and security of the new digital world. So we see a need for a consistent and accessible standard to be able to measure these products comparatively, to be able to determine what are better and worst performers for these sort of items.

Now in order to do this, we want to leverage the deep knowledge and the expert knowledge of many in our community. So for this initial effort, we're working with several well-known organizations, including Ranking Digital Rights, Disconnect, and Cyber Independent Testing Lab. Starting about six months ago, we got this group together to leverage the diverse expertise of the group, of this core group that we have, to start the process of putting together a proposal for a digital standard. So we met first to compile a draft of various criteria. We split up the work then to exercise tests for the various parts of the proposal against a few products from three different product verticals. We tried it against browsers, against some mobile apps, and against some connected devices.

And the purpose of the testing wasn't necessarily to investigate these particular products, but it was to vet what we had done, vet the proposal for sense, that it was feasible to do testing on it, and to improve based on what we had found from that course of that testing. So we're now refining the proposal, and we're working on getting ready to launch it. So before we can test of course, we need a shared definition of what is good, so we know what we're looking for. How do we define what is goodness in this space, in this digital world? So we started by structuring the work around four organizing principles: Security, privacy, governance and compliance, and ownership.

So security answers the question, is it safe? It includes topics related to encryption and security, security updates, passwords, things of that nature. Privacy is, is it private? Deals with permissions, over permissioning, and data sharing, and consumer control of their data. Governance and compliance answers, are the policies strong for consumers? It covers how well companies may protect consumers' privacy, and also freedom of expression. And ownership, is it mine? This covers right to repair, and covers things like permanence of functionality.

So for each of these four topics, we define several criteria which are anchored on consumer expectations. So this is what's interesting. We didn't start from technical requirements, we started from what the consumers would expect. So for example, instead of saying this room should be between 68 and 70 degrees, and with such and such humidity, I'm saying, the room should be comfortable. That's the consumer expectation. The room should be comfortable. Then we would dive down into that and develop indicators, which were attributes or behaviors that would achieve the criteria. So the attributes that relate to, the room would be comfortable, would be temperature and humidity.

And then we would define test procedures. And we'd take a thermometer and place it in four sections of the room at these different times and measure the temperature. But always we wanted to go back to what the consumer was expecting, what the consumer needed. So we've got the-- so here's an example. That's the first thing I want to go to. So this is an example of criteria that we had come up with.

The criterion, what does the consumer expect? The product should be protected from known vulnerabilities that presents a danger from attackers. That's what the consumer wants from their product. This criterion has several indicators that we developed. One of them, is the software secure against known bugs and types of attacks? And then a procedure, how you would test this. Well you could launch activities from the user interface and monitor communications to and from the device. So we could apply this method to an investigation we had done last year on a mobile app that is named Glow. And that app, we had done that investigation before we developed this standard, these criteria. But it shows that it applies, and it shows that the criteria is relevant.

So for example, in the case of this Glow app, the Glow app is a women's health and fertility app. It's a mobile app on your phone, and we found that an attacker-- we were monitoring traffic to and from the mobile device-- an attacker could request a connection to a user. And then that user's personal health information was transmitted to the attacker's mobile device immediately, without the user accepting, knowing about it, approving the connection in any way. So we knew that that was a problem. And it showed it did not meet the criterion, right, the product was not protected. The consumer, the user was not protected. Your data was not protected from that vulnerability. So, as a result of that investigation, which we published in July of last year, Glow modified its app to better protect consumers data.

So we showed that we had impact on being able to determine that. We were able to get the change. The app got improved. And it's a much safer app, a much more secure app now for consumers. It also, by applying that to these criteria, we show that we can in fact-- it applies. It's a good way to think about this type of a problem. So we completed a whole list of criteria, not just this one that I'm showing here. Criteria, indicators, and procedures, and we put it in a chart. And this is what it looks like right now. So, don't get scared.

First comment is, this is not meant to be readable, so don't try to read it. It's just to show the sort of depth that we have here, and how many things we've got going on this. Lot of criteria, a lot of different details to this. It also shows that we have a-- this is a work in progress. We've got some criteria that are fully defined. We have criteria, indicators, pretty thorough procedures figured out. Some we have, we know what we want to test, but we're still working on how to do that. And some are not very well defined yet. We need to understand what we mean by that. What is goodness in this space? So we've got them color coded right now, green, yellow, and red, depending on how well baked we think that they are.

There's still a lot of questions to be answered. Do we have the right criteria? Have we identified the most important information? And how can this apply to different types of products? It's not just for one type of product. Now this is a big problem. This is a very large-- and it's too large, consequential, and ever changing a problem for us to do it on our own. And we can't do this, really, without the folks in this room, the people watching, and those who hear about this later. And so we're looking for help from everybody.

When we do launch this, we want to-- we're going to launch it as an open source. We're going to launch this in-- we need your help-- I'm sorry-- to make it better. We'll be releasing it as an open source project along with our core group of partners. And we're inviting input and feedback from

consumer organizations, security researchers, academics, hobbyists, and anyone who's interested. Our goal is an openly sourced digital standard that can be used to hold manufacturers and providers accountable for how they manage consumers' privacy, security, and data.

We want to have a guide for testing on a level playing field. And we want to give consumers a way to make informed choices. It'll also be an evolving standard. And it'll need continued work. And we're asking for everybody's help when we do launch it. We'll be having it available for feedback, Inputs for everybody who wants to contribute to it. To provide any comments or feedback, please e-mail us at that e-mail address above. And that's it. Thank you.

PEDER MAGEE: Thanks very much, Maria. And last but not least, we have Alethea Lange from the Center for Democracy and Technology.

ALETHEA LANGE: Good morning, everyone. Thank you so much for the opportunity to come and speak with you today. I'm really excited to be here at Privacy Con and see so many familiar faces in the room. OK, here we go. I'm presenting on behalf of a group of researchers who did this research together last spring. So it's myself, Rena Cohen, Emily Paul, Pavel Vanegas, and Gautam Hans, who all worked on this research. And it a partnership between the Center for Democracy and Technology, and the University of California Berkeley School of Information. So I'm going to tell you a little bit about what we're going to talk about today. The paper is actually pretty extensive, so I've chosen just a small subsection of the results to share in the small time that we have.

The goal of this research wasn't to figure out how consumers feel about existing online personalization. It was to figure out how consumers feel about online personalization as a practice, generally speaking. So for that reason we used hypothetical scenarios that may or may not be happening in real life. This makes it a little less applicable necessarily to what's happening right this moment, a little less actionable at the time, unless you are a company, or a startup, or an institution who's looking to personalize some of your practices, in which case this could give you an insight into how that might be received depending on what your plans are. So with that in mind, we're going to go through the design of the study, a few results, and then a few of the key takeaways that the group developed for policymakers which we interpret broadly.

So personalization brings us a lot of good things, right? Here first and foremost, location personalization in this example, you can see that I got what I wanted in the second result with a pretty much useless query, because Google knew that I was in Washington D.C. So they gave me the information I needed, despite only giving them the search, movie showtimes, which is not very helpful. So that's good. There have been concerns raised however, that personalisation also creates opportunities for discrimination. Pretty famous example that we've all heard mentioned a bunch of times, the Wall Street Journal found that Staples was pricing staplers depending on how close you were to a competitor's store. So if stores are not dispersed equally, if people in under-served population-- in under-served areas have less options, they will also pay higher prices. So that's a concern.

And then you saw another example in the news recently with Facebook's ethnic affinity advertising. Facebook allowed advertisers to target folks based on their ethnic affinity, which

was a criteria that Facebook developed, and is not meant to be a direct measure of race. But is been to sort of be a proxy measure of what are your interests. So this was announced at-- well, it's been around since 2014, but was announced at South by Southwest as a major achievement in marketing. The press responded with some concerns, and then later pro publica, placed an ad raising more concerns for the government. It caught the attention of the Federal Housing authority and HUD. And subsequently Facebook actually altered its practice to no longer allow advertisers to target advertisements based on-- advertisements for credit, housing, or jobs, based on people's ethnic affinity. So you can see that it's just not a hypothetical example, there are real consequences in the world if people are unhappy with how you're personalizing your work.

So here's our research design. We looked at a couple of different domains, really common areas for personalization, advertising, search results, and retail pricing. We looked at different kinds of data types. So, what criteria are they using to personalize this to people? And then we looked at what the source of the data was. So this is a little bit more abstract for users, but did you provide the information? Like you tell some sources your gender. Did they guess it? Did they watch you and sort of see what you're doing? And if they did guess, if they infer the information, was it right or wrong? So using these three pieces, we assembled hypothetical scenarios that we gave to the respondents.

Here's one that actually corresponds to the Facebook ethnic-- ethnic affinity example. The ad is shown to you based on your race, which was inferred, and is accurate. So we gave people one from each section of domain, and they ended up with 18 different options that respondents saw. They saw one of these, where they were able to answer, how fair is this? They were able to give an answer. And they were also asked to give an open text response. The survey was designed through Qualtrics. It was administered through Amazon Mechanical Turk. And then you can see a little bit here about the demographics. I won't go into it too in-depth, there's some criteria that we got a pretty good measure of diversity, and some where we didn't. So, not perfect.

Normally I would give like a longer lead in to this, but here's the response for location. So this is city or town of residence, pretty non granular location criteria. You can see that for the most part, people felt like this was fair, this was above board, right? The middle line represents neutral, give or take. Above is, if you will felt that it was more than neutral, more fair than neutral. Below is if it was more unfair than neutral. Each of the colors of the bars represents the different source of the data. So provided, accurately inferred, and inaccurately inferred. And all of our slides on the results will follow this format. So your town of residence, people were generally seemed like they think this is pretty fair.

One thing to note here that will come into play when we're looking at the general takeaways, is the presence of negative results for inaccuracy. So one of the general data quality criteria is accuracy. We wanted to test whether or not that was really true. It tends to be the case, although it's not always the case, that inaccuracy hurts you. So one thing to observe as you're looking at the slides. And then here's a sample of the open text responses from folks. You can get a deeper sense of what they were thinking. Getting an ad, perfectly acceptable, maybe beneficial. Search results, maybe depends on what I'm searching for. Like weather, movies, restaurants, useful. If it's something like a book or a movie, maybe not so useful.

And then the last one I think is really the most important result from this section. People were pretty OK after they thought about it for a while. At first you think, well it's not really fair to charge people in different places different prices. And then as soon as you say that, you're like, oh, actually that happens all the time. That's pretty normal. So people kind of had to think it through a little bit, and if they could draw a parallel to something they were familiar with, it helped them interpret the information.

Gender had a bit more of a diverse response. Here again you can see inaccuracy tends to hurt people. One thing to observe, and I this is also a theme across most of the results we saw, is that between ads, search, and pricing, the stakes tend to be higher. Like, what advertisements you see, people generally don't consider that super important. What search results you see, really important. And then what price you're charged has concrete economic consequences, So it's the most important, or the most easily understood to be important. And you'll see that the responses will follow a general-- as the stakes get higher, feelings get more severe.

Here's the very nuanced response on gender from the quotes. One thing that might be a little bit interesting is to observe that little bit of a disconnect between ads that seems harmless, and then products that cost the same-- probably should cost the same regardless of gender. There's clearly a connection between what's advertised to you and what you pay for a product or service. But it's more in line with, which product or service are you getting? And there's been some research that demonstrates pretty clearly that women pay higher for the same service. That is a different thing, the pink tax. So it's interesting that people disconnected those pieces.

And then race, very negative. People do not like the idea that stuff is personalized on race across the board. Here you see very clearly that the consequences get bigger, people's feelings get stronger. This mirrors the Facebook example. A lot of what happened with the Facebook example is a gut negative reaction. People just don't like the idea. There are certainly examples you can think of, though, where it makes sense to have personalized results based on race. Maybe if you're searching for a particular hairstyle or cosmetics, or maybe if you need advertisements like for products that need to reach a really specific audience to be successful. It might actually be detrimental to say that they can't target those products on race. But those are sort of smaller use cases, and they didn't come to people's minds as quickly even though you saw in the location example, people did really have a more intuitive understanding of what happens in the real world.

So it's really interesting and it's more complicated than just this, but this gives you a good sense of folks gut reaction. Here's more context for the race. These quotes are pretty powerful. People are concerned about relevance, they're concerned about stereotypes, they are obviously concerned about discrimination. People feel it's pretty high stakes, and definitely felt was problematic. And felt they may be missing out on what they wanted. So just a few takeaways to end the conversation, and I think these would be useful for anybody who's looking to personalize or think about personalization and how folks feel about it.

High stakes domains, people feel more strongly that personalization should be more careful. They have higher reactions to what kind of personalization is happening. Personalization based on location data, however-- but again a pretty non granular version of location information-- was

pretty acceptable across all the contexts. People felt like they understood why that would happen. They saw some benefits to that, in fact, for their own behavior. The third one I think is really important, especially at the FTC, thinking about the FIPS, is that the data quality was meaningful. Accuracy did matter to folks, and you saw some mixed results on it. But generally speaking, inaccurate data was perceived with a more negative feeling.

And then personalized pricing should mirror offline practices, as I mentioned. And the last one I think is really something that's going to become just more and more important as we go forward in the next few years, is that personalization based on race is really controversial. And I think what we saw in the quotes that people responded to was it was controversial because people weren't sure it was relevant. Which is really interesting because a lot of the goal of personalization from the company point of view, is to increase relevance. But they're using sort of a different version of the word relevant, than I think a lot of people are, when they're sort of casually thinking about what's relevant to me as a person.

So I think that this actually adds a lot of really important information. And you saw how the Facebook ethnic affinity example played out in real life. That should help folks who are doing startups or any sort of data intensive analytics think about how to improve the quality of the products they're offering and how to reach the audience respectfully, and yet reach a relevant audience. Thanks.

PEDER MAGEE: Thanks. Thanks very much. And thank all of the researchers for their great work. What I'd like to do now is, we've got a little bit of time, around 20 minutes, to have a discussion. I'll start out by posing a question for each panelist, and if anyone in the audience has a question, please line at the microphones and we'll call you in turn. So why don't I start out with the first project, Noah and Dillon. Your research indicates that a passive network observer such as an ISP, can potentially analyze Internet of Things network traffic to infer sensitive data. What other entities could be in a position to do this, and would it be harder or easier for them to do it than an ISP?

DILLON REISMAN: So, we didn't really touch on this as much in our presentation, but we're kind of working on this now. We talk a lot about last mile kind of observers, but similarly wi-fi eavesdroppers can also do this sort of attack. It's slightly different, because they have less clear information. For instance, they get the traffic over radio, so they can see the actual traffic leaving the device. But if you use, say a password, on a Wi-Fi router, generally then the Wi-Fi radio is encrypted. So they can't see the contents. Which is similar to the last mile, but it also means they can't see the DNS queries, which is how we identify devices.

But we actually have discovered, and our further research has kind of shown, that through machine learning this is not really an issue. You can actually use the traffic rates to also, if you have a model trained on preexisting devices that you know about, that maybe you trained in your own lab, you can actually build a model and identify the devices without those DNS queries. So your neighbor can do this potentially, too, or anyone driving by with a radio.

PEDER MAGEE: So obfuscating the DNS query wouldn't necessarily prevent this type of tracking.

DILLON REISMAN: Not against that sort of threat, no. Not against that observer.

PEDER MAGEE: Are there certain devices that you think are more at risk than others?

NOAH APTHORPE: Yeah, that's a good question. And we definitely wanted to emphasize that the attack that we mentioned is not specific for a certain type of device, it's more of a general class. But we think that medical devices or personal activity trackers are especially worrisome, just because they perhaps get at information that many consumers would consider to be more private than perhaps when you're turning on or off a light switch. And they also typically have sensors that are able to detect medical data that the other devices that we looked at weren't. Like, your security camera probably won't know your blood pressure, but if you have a blood pressure monitor, that could potentially be revealed.

PEDER MAGEE: But that goes more to the type of data that's collected, as opposed to the device itself-- the security of the device, or the way the device is designed?

DILLON REISMAN: So, I guess what-- Noah mentioned something like-- let's say you had a device that actually knew your blood pressure. I guess we wouldn't say, our sort-- this class of observation wouldn't tell you literally what that blood pressure is. But the mere presence of that device would indicate something about the user. Perhaps they have high blood pressure, perhaps they have diabetes. It can be used to infer higher order behaviors too. So that's really what the potential threat is with different classes of device.

NOAH APTHORPE: Right, because we also didn't get into, in the talk, but we think that it might be possible to combine information from multiple devices in order to make these higher level inferences. So if you can see when a user is interacting with their sleep monitor, their television, or their kitchen appliances, maybe you can make more interesting inferences about their lifestyle choices than you would be able just looking at one particular device.

PEDER MAGEE: OK. Aleksandra, in your paper, you talk about the fact that both Apple and Google have some privacy protections in place with respect to apps running scams for nearby Bluetooth devices. But you say that both companies need to do more. What measures should these companies and others do to prevent apps from using Bluetooth signals for tracking?

ALEKSANDRA KOROLOVA: Very good question, thanks. I think what more they can do is actually change their APIs a little. So right now, all the apps can receive all the information that all the nearby Bluetooth devices are transmitted. That's not necessarily necessary. And also it doesn't have to be the case that they receive the accurate information.

So they don't need to receive the same Mac address every time. If the Mac address gets perturbed, per application level, then the attacks that I described would be harder. If the name doesn't get transmitted every time, then the attacks would be harder. That's one change they can make to the APIs. Another thing that I think they should make is, they should give users more control. So the same way that users can decide on a per application level whether this application should be able to access your location, users should be able to decide on a per application level whether this application should be able to access Bluetooth.

PEDER MAGEE: And you note that short of just turning off the Bluetooth on your mobile device, there's nothing a consumer can do right now to prevent an app that's designed this way to reach out and collect Bluetooth signals?

ALEKSANDRA KOROLOVA: On Apple no. On Android, starting with API level six, so Android 6.0, users have to grant one of the location accesses, in order for the apps to access Bluetooth. But when users grant location access, they are not expecting that with this they are also granting [INAUDIBLE] access.

PEDER MAGEE: Did any of the devices that you looked at, do they have any sort of disclosure about the concept that apps may reach out? And I'm talking about the Bluetooth devices. Did they say anything about the potential for this type of data collection?

ALEKSANDRA KOROLOVA: Not as far as I know, because there are lots of different privacy risks. And this is probably not a risk that anybody has yet paid attention to. So probably the device manufacturers don't know that this risk exists yet.

PEDER MAGEE: OK. Trying to get at least a question each. So, Maria, I know you mentioned that you're looking for input from stakeholders on how you can expand and shape the protocol. Can you be a little bit more specific and talk about what sort of feedback you might find the most helpful?

MARIA RERECICH: Yeah, sure. So when we do release it we're going to be putting it out on GitHub, so people will be able to see it there, make comments, any inputs to that. We're looking for people to be able to refine aspects of the standard, and also fill in the gaps where we have them. So coming up with test procedures. Also being able to extend it to different types of products, and be able to have it be applicable to different types of things.

PEDER MAGEE: We heard up here just now about two pretty sophisticated means of collecting information about consumers and potentially using it for tracking. I would imagine pretty difficult to detect. We have these very smart researchers up here who went out and found it. But that's probably not going to be true for everyone. Is there a role for this type of research in shaping your standard?

MARIA RERECICH: Yeah, I think definitely. Keep in mind that the standard starts from the consumer expectation of their data privacy and security, right? So something as simple as, the consumer wants their product to be protected from vulnerabilities. Well these are vulnerabilities. They're not ones that we might have predicted ahead of time, but they're are things that can be added in to this evolving standard. And the evolving standard could be, we're going to check for these sorts of different vulnerabilities and be able to check it in this way on these types of products. So it's very open to this additional research. And to be able to incorporate new things that are found, new things that are suspected. And being able to plug that in to a growing standard, to make the products more and more secure and private for consumers.

PEDER MAGEE: I know that one pressure on research is funding. Have you thought about developing-- partnering with people, or are developing it through Consumer Reports-- a bug bounty type program where you incentivize researchers to look for these types of vulnerabilities?

MARIA RERECICH: So, what's interesting is that one of the items in the standard we're working on is a company that's open to-- we think it's beneficial if a company is open to getting reports of vulnerabilities in their product. And so if they have a bug bounty program, that's a good thing for a company to have, because they feel it shows that they're open to that kind of information.

PEDER MAGEE: So that would be reflected in how you rank, or comparatively rank--

MARIA RERECICH: It could be. It could be, right. It could decide to be.

PEDER MAGEE: OK, great. Alethia, as the scope of big data gets bigger and automated personalization using machine learning becomes more sophisticated, are companies more or less likely to be able to tailor their advertising, search results, and prices, in ways that consumers perceive them as beneficial rather than unfair?

ALETHEA LANGE: I think that's entirely up to the folks who are deciding how it's tailored. I think what's really interesting from these results is that they were-- broadly speaking, people had a pretty nuanced understanding of what the trade-off was. They understood that they were getting benefits. They understood that there were things that they were losing. I think where the equivocation happened was on the concept of relevance. So the way that we define ourselves-- and I really appreciate Maria's framing of the issue-- just think about how consumers think.

Think about how people feel. It's really not that hard. They're like, go talk to some people. How do you feel? Who are you? And people aren't going to say, like, I'm a white lady between 35 and 45 who lives in a major city. Right, like, that's not how I introduce myself. It's not how I think I am. But I might be-- so I might be uncomfortable if I find out that that's the way I've been categorized, and if everything around me is tailored for people like me, but as defined by somebody else.

What you want to create is a world in which people can more define their own futures, and tell you who they are. And then they'll be more comfortable with what kind of tailoring they're getting. They have no problem seeing infinite advertisements for bicycles. I don't consider that a violation of my privacy. I like bicycles. Even if somebody inferred that I liked bicycles, I would say, yeah, I like bicycles. Like, that's OK with me. So it really just depends on how in-depth people want to go. I think you have to be careful, though, to make sure that you are not adding information onto this dossier.

Like, here's a white lady between 25 and 35, who lives in a major city, and likes bicycles. You no longer need that first part. I mean, lady might be nice, because it'd be nice to have clothes that are fit for women, if somebody would make some for bicyclists. That'd be great. But it's like, that other piece can go away. And I think that's really what we have to make sure happens is, it's not that-- my conclusion that you should draw shouldn't be, go collect more, specific information. It

should be, start thinking more clearly about what you want and what you need, and what people expect and what they like. And then only get those parts and let go of the sort of blunt instrument stereotype categories that are causing probably more trouble than they're worth.

PEDER MAGEE: But doesn't that concept of getting rid of the blunt categories sort of encourage additional data collection so that you can refine it, make it more precise? There seems to be a little bit of tension there.

ALETHEA LANGE: I think you can learn stuff-- you can drop the underlying infrastructure that we pulled forward from offline advertising, where you had to be blunt because there were only so many vectors. Where you had to go to magazines and you had to know what area people lived in. Or you had to go to what your interests were in a really broad sense. And now you can have a more personalized view, which is the goal of advertisers and companies. And I think individuals like that, so long as what the companies think is a personalized view doesn't seem offensive to them.

And offensiveness isn't just about being relevant. It may be relevant that people see ads for high blood pressure medicine, or for plus sized blouses. Those may be, in the strictest terms, relevant ads. All right, people don't like them. My dad got advertisements for tombstones when he and my grandma were planning her will. Relevant? Yes. Offensive? Yes. Things can be both relevant and offensive. And that's the category that if you're trying rid of stuff, I'd get rid of those first. And I wouldn't add a bunch of other stuff to figure out what those are. Just sit and think about it as a person for a little while, instead of just machine learning. I mean, well, whatever the computer says is relevant is probably relevant. And even if that's true, it doesn't mean it's a good thing.

PEDER MAGEE: Interesting. I see we've got an audience question, so I'd like to take that.

SPEAKER 1: I have a question to Aleksandra Korolova. The part of the [INAUDIBLE] is, you lead it to the names of the devices. So like in your example, if I have a really expensive TV, it indicates my disposable income. Can't a consumer right now just change names of devices to something much less meaningful? Because if I have to indicate my mother's maiden name, I learned not to give my mother's maiden name. I make it up. So similarly, I can make up a name for my expensive TV, or for my heart monitor, or for anything that I don't want people to recognize.

ALEKSANDRA KOROLOVA: So in some cases, yes, users can change the names of the devices. But typically the way it goes is that they change it to something more recognizable, not something less recognizable. So if I have a fitness tracker, I'm going to call it Aleksandra's Tracker, typically. I'm not going to call it Non Tracker TV. And for some of these devices, I don't think you can change the names. I don't know if all of the devices provide the opportunity to change the name. And in general, this kind of profiling is not top of mind for a user when they're naming their devices. They're naming it for their convenience. So you want to know that this Bluetooth device is a TV, when you're trying to access it. It's not much of a solution.

SPEAKER 1: Well, thank you. It could be a partial solution.

ALEKSANDRA KOROLOVA: It could be a partial solution, but it's a solution that would come at a cost of usability among these devices. And would put a lot of the burden on the individuals.

SPEAKER 1: I have a quick question to Alethea Lange. And you've worked for the Center of Democracy, as well as technology. And the democracy, in my opinion, suffers when people get segregated into like minded groups. And the advertisement that enforces stereotypes also, in my opinion, enforces segregations in the groups. Is it possible to do something about it?

ALETHEA LANGE: Yeah. You've just become my favorite person, because you've given me the opportunity to use my favorite example, which I bring up a lot. Which is, I absolutely agree with you. And I think there's a lot of benefit for democracy in the noise of seeing things that are not deemed relevant for you. And even in the context of advertising, which is fairly casual, getting a sense of what kind of products are out there, and getting a sense of how things are advertised to other people, gives us some insight into the world that other people live in. We talk a lot about how your computer in modern times is a mirror as well as a window. When you open it, you're seeing not only what you want to see, but how the world perceives you. Seeing advertisements that other groups of people are seeing is really important.

The example I often use is, maybe there's some value to democracy for men to see ads for period products, for tampons and pads. Like, maybe there's some value to that. Maybe it's something that demystifies it. Maybe it's something that makes it just a part of normal life, as opposed to something weird that women need to deal with. And that's not an easy argument to make, because there's not a lot of men-- although I'm sure there are some, single dads and whoever, who actually need to decide what brand of these products are going to be relevant to them. And you're sort of asking people to throw away money in the service of democracy.

The counterargument I would use to that, is that if you think about the way that parenting has changed, and the way that we perceive parenting in both advertisements and real life, those things have to go in lockstep. And so you'll see more men now, more dads, featured in parenting ads. More advertisements to parents and not moms. That's a really important cultural development. And advertising plays an important part in that. And I think in some ways they drive it a little bit. Like, it comes from people, and then it's sort of picked up on by folks who are deciding who sees what ads, and what the ads look like. And then it's perpetuated and becomes normalized over a period of years. So advertising is really powerful. And I absolutely agree with you that it's really important that there is some noise in that, or what would be considered waste, in the service of democracy.

SPEAKER 1: I'm happy to be a favorite person, but that just adds to this. On the [INAUDIBLE], we have half of people who are professional women in technology. But being women, if you enter the search for a book, you would be more likely to get some chick like book, than a book about technology, or important political policy decisions. Thank you.

PEDER MAGEE: Well, those were great questions. Thank you very much for asking. I want to be mindful of our time. Our panel has run its course. And I want to give people a chance to have a break. But thank you all very much. Your work is fantastic. It was very interesting. We appreciate you taking the time to come and present to us, and we really appreciate it. So thanks

so much. We're going to take a break until 10:35. Refreshments are available for purchase in the cafeteria, which is down the hall. Regrettably, neither food nor drink may be brought back into the auditorium.