FTC PrivacyCon 2018
February 28, 2018
Segment 4
Transcript

MARK EICHORN: I'm Mark Eichorn. I'm in the Division of Privacy and Identity Protection here at the FTC. And welcome to our fourth panel of the day. We're going to be discussing tools and ratings for privacy management. And following this panel, we will roll into closing remarks by the FTC's chief technologist, Neil Chilson, so please stay around.

I want to thank all of the participants today. I think it's been a great day, a really interesting day for me. I've learned a lot today. And thank you all for watching online and for being here in person.

I would like to start off with a presentation by Periwinkle Doerfler. She's a PhD student at NYU within the Center for Cybersecurity. And hopefully, we will have a demo here.

PERIWINKLE DOERFLER: Right. Is this my clicker, or do I have one of those?

MARK EICHORN: What's that?

PERIWINKLE DOERFLER: What do I press to click it?

MARK EICHORN: This one.

PERIWINKLE DOERFLER: Great. Thank you. OK. So hello, everyone. My name is Periwinkle. Yes, really.

I'm from NYU. I work with Professor Damon McCoy. That's backwards. OK, great.

The presentation I'm giving today is based on a paper that's been accepted at IEEE Privacy & Security, better known as Oakland, that was conducted with other people, as well, most of them at Cornell. We're talking about spyware but generally, mobile-- specifically, Android-- applications in the space of intimate partner violence, hereafter IPV.

To give a little bit of background, one in three women and one in six men in the US in their lifetime will experience some kind of intimate partner violence. And this is a space that gets blocked out a lot in UI/UX studies.

But the important thing to keep in mind is it's a different threat model than a lot of times security researchers look at, because the standard advice of, well, change your passwords doesn't really apply. Because if you're in an abusive or stalker kind of situation where someone has access to your passwords and you change your passwords, you may, then, see an escalation from digital stalking to a physical confrontation.

So generally, when we're doing this kind of research, we're not thinking about standard defenses. We're thinking about a situation where people are vulnerable and they're not able to employ a lot of these things.

So we're looking at consumer-facing applications that exist in the Play Stores and the open web that are readily accessible to a potential abuser, easy to install, easy to use. And there are hundreds of those, some of them more benign than others. And then there are also dozens, if not hundreds, of forums and blogs, YouTube videos, demonstrating for a prospective abuser exactly how to install, use, and hide this kind of software.

There's an important distinction, though, between things that are overtly spyware-- the stuff you see over on the left is pretty clearly-- it's nasty and tends to be nasty. It's surreptitious. It actually has a picture of a man beating his wife. So it's pretty clear that stuff's bad.

But there's also this whole other realm of what I would call dual-use applications, which are applications that have pretty much all the same capabilities as overt spywares but are branded for things like tracking your children or tracking your employees, things that are legal and legitimate. So this is a big focus of our research is understanding there's a dual-use ecosystem and how they can be-- but more importantly, how they are-- being used by abusers.

So as a demonstration, these apps, I think, are a lot more invasive than people realize, because they have the ability to do all sorts of things-- remotely activating cameras without notifications, reading and forwarding all texts, recording and forwarding all calls, tracking location, storing location. In the case where we were able to jailbreak an iPhone or root an Android, you also have the ability to read the contents of Facebook messages, WhatsApp messages, Kik, Fiverr, et cetera. It goes on and on and on.

But incredibly invasive tools. So this is why I handed you the phone. I'm going to give a demonstration, now, of an app called TrackView. I have two apps that are-- yeah, switch it over. Great. So I have a phone up here, which is the controller app, the controller phone. And then hopefully, this works.

Yeah. So if you would just pass that around. You can see that, up here, we're projecting the screen from the phone up here, which is recording from the phone down there.

And maybe somebody who's in the audience holding the phone can verify for me. That phone looks like it's off. You can't see that it's recording. It doesn't say that it's recording. Yeah.

So this is why I wanted to pass it around, because it gives you an idea of the fact that it's super creepy and super invasive, and there's no way of knowing that that's what's happening. So thank you all. In the interest of time, I'm going to switch back to my slides. I hope this is an informative look at how invasive these apps can be.

And that app is freely available on the Google Play Store. If you have an Android phone, you could go. You could download it right now. And for a mere $10 a month, you can have all the

creepy features where it doesn't say it's recording, and it doesn't have an app icon. It doesn't do anything on the target phone.

So I do this demonstration a lot when I go and talk to people, because I think you always assume that things that are that invasive and that hidden are hard to find. But they're not. It's quite literally on both the Apple App Store and the Google Play Store.

And not only are these apps super creepy, but the companies know it, and the companies advertise this use case. So if you go onto the Google Play Store and you search for "apps that can help me track my cheating wife without her knowing," you see advertisements. Similarly, if you go on Google web search, and you search, "how to track cheating husband with his cell phone," same thing. Lots of advertisements.

We searched Google continuously for a month with a corpus of about 400 search terms. We found ads on about 200 of them. Some searches showed as many as seven ads at a time. 40% of those ads were for applications.

And they are varying degrees of creepy. We acknowledge that certain apps are probably just advertising on the word track, and that's a reasonable thing if it's a "track your children" app. But some of the ad text is explicitly saying, "catch the cheater now. Tracking application." So it's definitely relevant, and these companies are very aware of this use case. They, in a lot of ways, are advertising this use case.

And even when they're not advertising explicitly this use case, they are complicit. Because what we did, then, was we contacted some of their customer support. And like, hi. Now, I'm Jessica Williams. And I have contacted your customer support, and I've said, "Hi, if I use this app to track my husband, will he know that I'm tracking him?"

You get responses like this that say, pretty much, no, that's fine. Or no, that's fine. Or no, that's fine. Out of 11 companies we contacted, eight responded. And of those eight, seven responded like this. They were entirely complicit. We made it very clear in our initial request that we intended to surreptitiously track an adult human being that we were in an intimate relationship with.

And they were very helpful. We had one company-- shout out to TeenSafe, if they're listening-- that actually admonished us and told us that was illegal, that we couldn't do that, which is an appropriate response. Alas.

The other thing that we found is that the existing anti-spyware tools are just not effective in preventing these problems, because a lot of these applications are not hard-line malware. These applications have legitimate use cases, such as tracking children, employees. Or there's applications that are designed for if you lose your phone, that you can find it again.

So VirusTotal flagged 7% of the applications as a risk. Norton caught 13% of those that were on store and 70% of those that are off store. I highlight Norton specifically, though we tried many, many antivirus and anti-spyware applications, because since we've written this paper, Norton has

actually reached out. And we're working with them towards building a tool that is not going to be generally deployed, because most people who have child-tracking apps want them, but to be able to be deployed at shelters and similar kinds of places so that people who are at risk can have their phones scanned in a little bit more fine-grain level.

But the takeaway here is that if you are in an abusive relationship, if you're trying to leave an abusive relationship, or if you've left an abusive relationship recently, these kinds of applications can pose a huge risk. Because if you're trying to leave and you know that this stuff is on your phone, you can't remove it for fear of being physically assaulted. You can't contact a therapist, you can't contact a shelter, you can't contact family and friends if you know that your abusers are able to read the contents of those messages. It clearly creates an extra barrier to leaving that kind of a relationship.

And in the case where you don't know that this is on your phone, maybe you do try to leave and then you, again, risk a physical confrontation. Or you do leave. You get as far as the shelter, and you can't figure out why your abuser keeps finding you.

This is already, kind of obviously, patently, a big problem. But the other thing that we really find, this is hurting at-risk populations even more. Because oftentimes, people feel the only recourse they have is to run their phone over with a truck.

Not in my own research, but in the research of Sunny Consolvo at Google, she's actually done studies on this. And she's talked to people. And I remember that being an actual line that she had from a survivor, which is, "I ran my phone over with my truck to make sure the software is gone."

Not everybody can afford to get a new phone. A lot of people can't. And it probably goes without saying that the people who can't afford a new phone are already more at risk, especially if they're in this kind of situation and are trying to leave it.

Where do we go from here? As a research team, we're exploring a lot of avenues, the biggest one being developing a tool that can more effectively find and remediate these kinds of softwares. We'd love to see stricter enforcement on the application stores. And we hope there are legislative or regulatory approaches. This slide is pretty bare bones, because it's a difficult question.

Certainly, if any of you have ideas for legislative, regulatory, corporate policy, technological solutions, or just even things that would help people in these situations, I'd love to hear it. Hit me in the Q&A. Come up to me afterwards. We're always looking for more collaborators to really try to solve these kinds of issues. Thanks.

MARK EICHORN: Thank you.

[APPLAUSE]

Our next speaker is Saksham Chitkara, a graduate research associate at Carnegie Mellon University. And his paper is "Does This App Really Need My Location?-- Context-Aware Privacy Management for Smartphones."

SAKSHAM CHITKARA: Hi. So as Mark mentioned, I'm Saksham. And I'd like to start by acknowledging my co-authors and my advisors, Yuvraj and Jason. So I'm going to be talking about context-aware privacy management. Let's start by looking at what smartphone management right now is on the privacy management.

So right now, you have millions of apps. You have billions of downloads. And you have more than 100,000 developers. How do these developers make money? Barring some paid apps and some in-app purchases, most of the smartphone industry right now depends on advertisements.

And advertisements entail private, personalized advertisements, which means collection of your private data. So developers usually used third-party libraries to collect your private data.

In the app permission model right now, it's a black box. The users don't really know why the app is requesting my location. So if Angry Birds requests access to your location, you don't really know why.

On top of it, it's like the number of decisions that you have to make. It's really, really large. So if you have just 80 apps, that leads to, with five permissions approximately per app, about 400 decisions. That's cognitive overload.

And one more additional thing is that the third-party libraries which developers use, developers really want to focus on what they want to build. They have limited time. So they're like, oh, Google has already done all the advertisement and personalized tracking for me. So let me just plug and play and use AdMob.

And once these developers have, they just plug and play and use AdMob. They don't really focus on the privacy concerns of these libraries. These libraries, they run with the same permissions as the app. So it doesn't really matter. The OS doesn't really check. Is the request originating from the app or the third-party library?

So how can we add the context to these permissions right now? If we add a purpose to each and every permission, the number of decisions that the user has to make would increase. It would shoot up exponentially. I don't really want to make 2,000 decisions to figure out where my private data is flowing.

So how can we leverage what we know about third-party libraries to provide a better model? Let's look at third-party libraries in more detail. We analyzed some of the top apps, the top 11,000 apps. And if you see the graph over there, about 4,000 of the 11,000 apps, they are like, if I am requesting access to this particular permission, they really need it.

If you look at the middle part of the graph, 2,000 of these apps, they are like, OK, I need this permission, but I also have a third-party library which will also access this permission. Not too bad.

But the third part of the graph, that's about 4,000 of these libraries. They don't really need your location information or your personal information, but they are just requesting these permissions to mine your data. This is a really dangerous territory. Without any use case, they are mining your data.

So we looked at third-party libraries in more detail. If you consider the top 11,000 apps, some of the libraries are really well supported. They are backed by big players-- Google has AdMob, and Crash Analytics and Flurry by Yahoo. So they are very well documented. You just plug and play and they work.

So a lot of these developers use these popular libraries. But some of the less obscure libraries, they're not so common. On top of it, if you look at all private data accesses, then they are originating from very few libraries.

So on the graph on the right, you can see that there's an inflection point about at 40 libraries. So what this means is that if there is an app requesting access to your location, chances are that, 50% of the time, it's not really required for functionality. That's a big number. And it's originating from one of those 30 libraries. So 50% of the time.

So we introduce app list library-based controls. We allow the users to make decisions on the basis of the apps or the libraries. So we are able to infer the stack trace data and figure out the context, which allows us to separate the data flows between the apps and the libraries.

What about the purpose of these particular libraries? Turns out that libraries are used for a very singular purpose. If you are installing and using AdMob in your app, you are probably displaying personalized advertisements. If you are using Flurry, then it's probably for analytics and debugging.

So we can just look at these 30 libraries to infer the purpose. And using just a simple thing, we are able to provide the true purpose of the data used.

On top of it, we used crowdsourcing to scale our collection. We only need one user in our system, which sees a particular app and a permission and the true use case, based on the stack trace data. We add it to our DB so we know that this app is requesting this particular permission for this use.

So we introduced our Protect My Privacy app. It was published on the Google Play Store and the rooted store. And we gave the users the functionality, the ability to make decisions based on these libraries, as well as these apps.

We assured the users that-- I don't know if you can see it, but there is this library, which is present in these particular apps on your phone. Which decision would you like to make? Do you want to allow the location to this particular data, or deny it, or fake it?

We send them a reminder that you recently made this decision. Do you want to change it? And then we have the UI, which allows the users to see all their decisions that they've made for the libraries.

So 1,300 real world users downloaded our apps. And turns out that if you just consider 90 decisions, then 90 decisions is nothing. 90 decisions is like five permissions per app, so 18 apps. I bet most of you have more than 18 apps on your phone.

So with just 18 apps, the number of decisions that the users have to make goes down by more than 25%. And this number will continue to go up, since half the accesses are not really needed for functionality and are only made by the 30 libraries.

What about the effectiveness? So in the previous model, if you had AdMob present in five different apps-- and let's say it was present in Angry Birds. Angry Birds was requesting access to your location. You denied it. It was present in some voice app. You denied your location to the voice app.

But it was also present in the Maps app, and you had to allow your location to the Maps app. So the end result is that the libraries end up with your private data.

In our new model, this will never happen. The data will only flow to the third-party libraries if users consent to it. It's transparent, and it gives the users the control. It will always be better than the existing model. So based on all our results, turns out that users are protected in a 25% better manner compared to the existing model.

What about the user decisions? Initially, the users were sort of unaware about all the private data leaks that are happening in their devices. Once we showed them the notifications that, oh, your data is leaking through these particular libraries from these apps-- they had no idea that this was actually happening. So they tended towards caution.

We have more results in our paper, but the overall goal right now-- the users were blocking much more decisions than they were allowing. They wanted to see a true use case of it. We wanted to analyze why the users are making these decisions.

So we did experience sampling and asked for in-app feedback. We asked the users, you recently made this particular decision. Why did you do it?

So based on the results which users said, they were like, OK, I allowed my location to this particular library, because I'm OK with it. If you're showing me ads for nearby restaurants, I want those ads. So they were OK with sharing their data and receiving personalized advertisements in return. This was only the 25% of the users who were allowing their decisions.

While denying, the users mostly said that it's about, I don't really see a use case of my data. I don't really see why my data is being used. So they tended to block more. They wanted the libraries to earn their trust before they were willing to share their private data.

So overall, we introduced this app. We had organic users. And we support our system by crowdsourcing. So if only one user sees a particular app permission, we will know it. And we can display it to the other users in the form of the notification that this app is known to access your private data for this particular use.

So based on the results from 1,300 users, the number of decisions that you have to make reduces by 25%. And your protection increases by a similar number. So thank you.

[APPLAUSE]

MARK EICHORN: Thank you. Our third panelist is Ian Douglas. He's a colleague in the Office of the Privacy Commissioner of Canada.

IAN DOUGLAS: Afternoon. So we decided we would take a look at health and medical devices and try and come up with a scoring system for them as to how they dealt with and protected privacy. So our initiative was proactive in nature. It wasn't a result of any investigations. There was a GPEN sweep. And GPEN is a group of DPAs worldwide that, every year, do a sweep to see how privacy is going through the world.

And we basically picked a bunch of devices that they had. So it was kind of a random choice. And based on the groupings that we had, we added a few more devices. So we had three in each category.

We tried to do trend analysis. And there was absolutely no enforcement actions at all. This was an educate and guidance initiative. And when we found something of interest, we would talk to the manufacturers, and we would let them know what it was we were concerned about or found. And in most cases, they would immediately make a change.

So the prescription used for this was we set it to look at data in motion. There we go. OK. Data in motion is interception of the data between the devices, or the IoT device and the app, or the app and the server. And we wanted to look at whether the information is encrypted, not encrypted, whether there was decent safeguards in place for that.

Data at rest is basically the storage of the data. And in the event it stored it on the phone or it stored it in the device itself, we wanted to know if it was encrypted. And in the event you lost the device, no one would get access to it.

So then we also check security. And that was SDKs, which was what was just discussed, the use of third-party parts that sometimes are sending information that you're not aware of. The passwords, how they were managed, how they were selected, whether or not they were stored within the applications themselves, whether we could find them easily, versions of the tools

being used, whether or not they were outdated and susceptible to attacks. And then pairing whether or not the device was always on, always listening, easy to connect to.

And then under practices, we had passwords again. And that was hard coding them or having the keys hard coded in the software so it was easy to find them. And always on, I discussed earlier.

And we tried to map them to the 10 principles of PIPEDA, and so we went with consent. Is the user able to determine whether or not the information is to be sent? Limiting collection, the number of third parties the data was shared with.

Limiting the use and disclosure. So that's basically the count of the number of elements that were collected. Safeguards is how well the device actually stored your information or transmitted it. And then individual access was the ability for the user to remove data or change the data after it was stored on the device or if you wanted to delete the information from the device or app upon selling or you didn't want it any more.

So for limiting collection, this is basically the number of data elements that were gathered. We assume that the more information they gather, the more they know about you. We also look at data types, which some of the people in session 2 this morning talked about, where a data type was crossed. So you had a data type that was identifying or a data type that was behavior. If you had two of those, that's more important than two just identifiers.

Limiting disclosure. Again, that was the number of third parties. We decided that, based on the number of observations, we saw during our studies that the number six seemed to be the most third parties that data was sent to, opposed to the amount of third parties we saw with apps. We were looking at apps before, and that was 40 or 45. So it was surprising that it was a lot lower.

Safeguards. Again, the security of the stored and transmitted data. Individual access. So that's modify and delete. And number of operational mandatory [INAUDIBLE].

So the scoring overview. We haven't actually published our report yet. So I have the copy of how we got to our scoring, which is here.

Limiting collections, so offers service by using as little personal information as possible. That means that we're trying to look at the number of elements that were collected and see if, did they just collect everything they could think of? Or did it make sense as to what they were collecting?

We also looked at the timing and the frequency of the data being collected so that we could see if they could determine, are you home, not home? You close to your house? As you get closer, some apps would flip things on, maybe, as we saw this morning in session 2.

So we then decided that the data type should be taken into account. Because originally, when we first looked at it, we thought, OK, maybe a model that was based on-- sorry. Logarithmic. Yes. OK.

So we originally thought that maybe a logarithmic equation would be better. But we realized it immediately took the user to a bad score. You can only get one or two items, and then you were badly rated.

So we thought, OK, well, we'll try and go with a different route. And we decided we would try a linear model. But linear grouped everything all together all the time. And we couldn't really say, well, this one's worse than that. It wasn't obvious.

So then when we introduced the data types, then we said, OK, so if we give 25.1 data elements, which is the population of Canada if you do 2 to the 25.1th power, that's enough to identify an individual. And we said, OK, now, if we allow 25 data points and they start crossing data types, we would start to reduce the slope. And the slope would be reduced every time you added another data type to the collection.

OK. Limiting collection. So in limiting collection, we said, OK, the more people you shared data with, then the more places your data is possibly exposed. And we figured, OK, what's a number that would be relatively good? And after looking at all of our devices, we saw that, on average, no one went over six. There was one did it.

But we went, OK, if you share one single element with six people, not too bad. If you start to share elements across data types with six, well, now you're starting to get worse off. And we found that as you cross four data types with four third parties, the score would be at zero. So again, we were reducing the slope based on the number of third parties that you shared new information with.

Safeguards. So safeguards, when we look at that, we kind of thought originally that we could just have a matrix and say, OK, if you have good passwords, that's good. If you have good encryption, that would be fine.

But we realized that it started to be somewhat subjective in nature. And we didn't want to have people criticize--

[BUZZER]

--us saying that-- oh, sorry. The things that we picked would be too obvious. So we grabbed CVSS, which is used to pick a score for vulnerabilities and hacks and things like that you find on the internet. They always get a CVSS score.

So we eventually came up with a combination of the two where we said, OK, you have to have decent encryption, you have to have ports that are open. And if they don't need to be, then OK, don't do it. And we also, in the event, found a vulnerability, we would use CVSS to score that vulnerability.

OK. Limiting access, we basically said if the user is able to go in and find a button that says, OK, delete my information, no problem. If I can reset the device to zero, that's great.

For consent, we basically took the total number of [INAUDIBLE] elements that were optional. We put them over all of the data elements that could be collected. And we used that ratio to come up with the number at the end of the day.

OK. My last line. OK. So we know the amount of data collected was quite varied. We had expected to find a lot of PI, but it wasn't. It was a lot of the developer SDKs were actually sending information back about the device.

So it would talk about the power the thing had left, how many times it had been used, stuff like that. And we realized that, OK, that shouldn't really be included in personal information.

The sharing was not as high as we expected. As I said earlier, we thought it was probably going to be something like 40 people that it would be shared with. And it wasn't. It was down around six, and a lot were only at two or three.

The established medical companies had much better safeguards than the newcomers to the game. They would make more errors, possibly in their selection of components. We had a weight scale that would turn on a access point, because the developers didn't realize that the component they had selected actually included an access point in it. And it also had a web server in it, which allowed for web configuration of the device.

It was password protected, but the password was very easily guessed. And we found another device that, when we updated the firmware on it, started sending out thousands of emails. And the email addresses were basically all of their users.

So we contacted them. They, three days later, put out a patch. It's fixed.

Scrubbing data was not always easy or possible. So we couldn't find instructions on how to do it in some cases. So if you wanted to get rid of it, you couldn't. The only way you could get rid of the data was to actually wreck the device.

And controls over what was collected was not always great. There were times when we would say no location information, and then they would, in an SDK, probably, work around it and get location information anyway. So that's about it. Thank you.

MARK EICHORN: Thank you.

[APPLAUSE]

Our next speaker, Katie McInnis, is a policy counsel with Consumers Union in the DC office.

KATIE MCINNIS: Thanks, Mark. Hi, everyone. As Mark mentioned, I'm Katie McInnis. I work as policy counsel for Consumers Union, which is the advocacy arm of Consumer Reports. And I'm very excited to share with you our first results of our digital standard today. Let's see.

All right. So we released our digital standard around this time last year. And our digital standard is a collaborative project that we completed with ranking digital rights, disconnect, and the Cyber Independent Testing Lab. The standard is designed to set expectations for manufacturers when they're developing connected products to make sure that they're keeping consumer privacy and security in mind, along with other digital rights, such as ownership and governance.

So as I said, this is the first time we have released testing results on our digital standard. And the goal of the standard and our testing under it is not only to educate consumers about the security and privacy options available to them but also to help incentivize or influence manufacturers when they're designing and producing these products themselves.

Although the digital standard is designed and can be used to analyze a large array of connected products or services, we thought smart TVs were a very convenient and obvious place to start. They are a hugely popular option for consumers. Many of the top TV models are also smart TVs. And we saw a huge number of connected TVs on the market in 2017, and that number is only expected to come up. In addition, smart TVs are able and are transmitting a huge amount of data about users and their viewing habits back to the manufacturers and the business partners.

So to conduct this test, we selected five TV models-- an LG, a Samsung, a Visio, a Sony, and a TCL Roku. Now, all of these models, except for the TCL, are recommended Consumer Reports models using our regular testing based on color clarity, ease of use, et cetera, but not assessing them on privacy and security.

The testing I'm presenting today was undergone with ranking digital rights and disconnect. In order to evaluate the TVs, we not only looked to the policies the consumer is presented with, but we also looked at the consent and data flows, along with security vulnerabilities.

So no surprise here. All the models that we tested want to watch what you're watching. So they do this using automatic content recognition, also known as ACR and other methods.

ACR works by taking little snippets of whatever you're watching, even if it's a kind of dumb product like a DVD or Blu-ray disc, in order to assess what you're watching. And then companies are collecting and using and processing this data not only to serve up targeted ads, which you may expect, but also to get at some of the more quiet companies, like Netflix, that have kept a lot of their viewing data in-house.

The TV policies that we also assessed stated that they could use the data from your viewing habits on these TVs in order to serve up ads to you on other devices that happened to be on your home network. We also found that consumers, although they can limit the data collection-- you're really undermining the functionality of the smart part of your TV. And we also found two models had concerning security vulnerabilities.

So ACR, not that new. But it has been around for a few years. It was a subject of an FTC action recently which led to a settlement with Visio for $1.5 million last year. Visio was collecting this viewing data from users, like we saw here in our TV models, but they were doing it without knowledge and consent of the user.

The FTC deemed that this was an unfair practice. And in her concurrence, Commissioner Ohlhausen said that viewing data is considered sensitive information. So you have to get consent from the user. But what does consent actually look like?

Here we have a screenshot of our Sony model that we tested. You can see that this is a consent prompt, and the user is asked to either say yes or no to a privacy policy. However, there's not a dedicated prompt letting the consumer know not only that ACR is happening but also the ability to opt in or out, although it is buried in this privacy policy.

So the question that is presented is whether or not this is effective notice and consent to the user. One good comparison to this example here is a Sears case of a few years back where the company buried the disclosure that they were using software to track users browsing data in a lengthy user license agreement. The FTC required the company to clearly and prominently disclose the types of data the software would monitor, record, or transmit.

So the question in here is, is this prominent disclosure? It's unclear. But we did find, via a reader poll that was conducted after we released our report on these findings, that 43% of respondents did not know that their smart TVs were transmitting data back to the TV manufacturers and their business partners.

In addition to this concern, many users may not realize that they can actually say no to these privacy policies. We've been conditioned, using our online services and apps, to see it as a take it or leave it situation where, if I'm going to use an app on my phone or go to a website, I have to agree to their terms of service and privacy policies in order to proceed or use their services.

That's not the case with these TVs. And for the models that we tested, users could actually say no. And yes, at some point, maybe they'd be undermining the functionality of their TV. But I don't think that users would necessarily know that they could say no to these policies and still proceed in the set up of their TV.

In fact, some of these pages even tell them that that's not possible. So here, we have the Samsung consent prompt. And as you can see, it's set up so that the user will say yes to all of these, despite the fact that the user does not have to. In addition, the screen tells the user that they have to say yes to these in order to proceed, despite the very tiny little skip button up there in the right-hand corner, which still lets them set up their TV without agreeing to these policies.

So as I mentioned, consumers can say no to some of these policies and, therefore, disagree and not allow some of this data collection. But then you're also losing, as I said, the functionality of the smart part of your new TV.

This is another consent page that we are presented with. This is the LG model. And in contrast to some of the other TVs that we tested, this is a more granular consent regime. So not only do you have here on the left hand side-- they're telling you what you're agreeing to in simple title language. But you also can turn off and on the functions as you see fit.

So this viewing data is collected for two primary reasons. One, for advertising. And then, two, to also recommend that, since you watched Game of Thrones, you may also like x, y, and z show.

One is clearly helpful to the user, and LG, here, lets you use your viewing data in order to have the smart TV recommendations that you may like this other show while also saying that you don't want them to monetize your data for advertisements. In contrast, we have the TCL model, which uses a Roku operating system, and has an all-or-nothing regime. You either have to say yes or no in order to proceed.

So for all the models we tested, there were post-set up television controls. The screens I've shown you so far are when you're actually setting up the TV for the first time. But for all these TVs, you can go in later and then readjust your settings.

However, these controls are really hard to find. So as you see here, it's not exactly intuitive. And the only one that had a dedicated privacy page is Roku, which is very helpful. But yet again, want to emphasize that if you do turn off some of these settings, you're undermining some of the functionality of your TV. And our reader poll that I had mentioned before, we found that the majority of respondents did not know how to disable these functions on the TV, which isn't surprising.

So the scopes of the controls and data collection were still kind of unclear for users, as I'm going to demonstrate in the next two slides. Next three or four, actually. This is a Samsung disclosure. And even though you can limit data collection for viewing information services and interest-based advertisements, Samsung still reserves the right to collect data for other purposes.

And more confusing still, [INAUDIBLE] shows that this data was going to a server called LogIngestionSamsungACR.com, even when viewing information services and internet-based advertisements were turned off for the user Although we did speak with a Samsung representative and they said to us that this wasn't data related to automated content recognition, it's still hard for a user to know where this data is going to.

Also, we have here-- this is the Sony TV. And you're presented with the Google privacy policy, which reserves broad data collection and usage rights. But then the user is followed up with a screen about allowing some data collection to improve the Android TV experience for diagnostic use.

So is that limiting what you previously agreed to? Is that saying that the data is collected only for the diagnostic purposes? It's very unclear.

We also discovered two security vulnerabilities, two TVs with security vulnerabilities, in our test. One was in the Roku operating system on the TCL model and also in the Samsung API, and it allowed a fairly unsophisticated user to access your TV and do kind of creepy things, like disconnect the TV from the internet, change a channel, change volume.

And while that's upsetting, I do want to emphasize that the concern here is the fact that you can actually use the data from these TVs to correlate with activity on other devices and serve more targeted ads. And I don't think that many people are buying smart TVs with that understanding.

In the future, we plan to publish more detailed findings from our smart TV test. We also plan to use the digital standard that I spoke of earlier to do assessments of other similar products and services. We also plan to look at criteria that we didn't get to look at in this test, such as interoperability and repairability.

And as a final note, thank you for letting me talk to you about our test results today. And I encourage you to go to TheDigitalStandard.org and contribute to our digital standard, which is an open source project, and give us any input you might see through our GitHub. Thank you.

MARK EICHORN: Thank you, Katie.

[APPLAUSE]

Norman Sadeh is a professor in the School of Computer Science at Carnegie Mellon University.

NORMAN SADEH: Thank you. So I'd like to talk today about a paper that was presented back in July at the Computer Vision workshop on privacy and security. The first author on this paper appears first. Unfortunately, he's not available to present, and so it's fallen upon me to try to do my very best here. The first author is Anupam Das, and there are a number of other people who've contributed to this paper, as well.

As I think we all realize to some extent, there are lots of cameras out there. Cameras have been deployed on a very broad scale. You might have come across a number of different statistics. Some people have said that on a typical day, we tend to come in front of something like 70 or more cameras, and that's probably not mentioning all the cameras that we have in our pocket.

And so as this happens, the question is to what extent people are comfortable with how much monitoring is taking place and how often data about them is being captured by these camera systems. Right. And these are more statistics that you can find. Some of those numbers are pretty amazing if you think about it. 6,000 cameras in Times Square, right? Who'd have thought? But someone apparently has counted.

And so the ways in which these cameras pop up these days come in many shapes and forms. Right, so those are cameras that are potentially deployed in the street. But as we know, you also have these glasses and other forms of gadgets that people are starting to wear. You've got, obviously, cameras in a variety of different environments that are potentially going to be capturing information about you.

I've read an article recently about deployment of these cameras, for instance, in China, where apparently now there are police officers that are wearing these cameras and doing facial recognition in real time to identify suspects. I don't know that we're going to see that happening anytime soon here, but that's certainly one of the trends that is emerging in some places.

And so what sorts of applications are behind these cameras? As I said, it could be monitoring and surveillance. But it's also lots of marketing types of scenarios that are being pursued, so scenarios where you walk in front of a store window and you've got cameras observing you, trying to look at how you're responding to different kinds of products, what products you're spending more time looking at, but actually also adjusting advertising offers that are being shown to you based on demographic information, based on your gender.

And in fact, even people experimenting today with customized ads where they're going to recognize you as a customer specifically and start offering you specials that are in the menu based on what they've identified as potentially being your preferences. So clearly, lots of different scenarios are being experimented with. I could go on with a very long list here of applications.

I'm sure that many of you are familiar with Facebook Moments. The display here shows you a few other types of scenarios, the specialized recommendations based on who you are at the restaurant. I'm sure many of us have heard about the Amazon Go store that's trying to figure out what you've been purchasing and charging you as you exit the store without requiring any interaction on your part. So many, many different kinds of scenarios, clearly.

And obviously, all this data can be mined, also, and that can lead to yet a number of additional inferences about your preferences, who you tend to go out with, what your health situation might be, and all sorts of other things that are fairly easy to identify as potentially sensitive for a variety of different people. And so when it comes to regulation in this space here in the US, as I think we all realize, we certainly don't have any regulation that cuts across the board and addresses the collection of this data.

There are a few states that have regulations when it comes to identifiable biometric identifiers. The ones that I'm aware of are Texas and Illinois. So in those states, only according to law, you have to disclose the collection of this type of data, which covers data being collected by cameras. It requires even some form of consent.

Obviously, in Europe, they've taken a different approach. In other places, they've taken different approaches. So certainly, if you're looking at GDPR, there's an expectation that these things will be disclosed and that you're going to be obtaining what is referred to as "affirmative consent," whatever that means. They're still in a process of trying to define at this stage how that's going to be interpreted. We can expect these things to evolve over time.

But what we wanted to do in our research is we wanted to better understand how people feel about these scenarios. And obviously, there are potentially different views. You might say, well, at this point, everybody knows there are cameras everywhere. So chances are they're already expecting this, and maybe we don't necessarily need to go overboard here in terms of notifying them and potentially allowing them to explicitly or requiring them to explicitly opt in or allowing them to opt out, depending on which particular approach you want to take.

And so we've done these studies in particular. In fact, one of our colleagues [INAUDIBLE] presented a paper earlier today on research that we've done across a number of IoT scenarios.

That study that you presented earlier today, in fact, included a number of scenarios that were centered around the use of cameras and facial recognition.

And so they could be deeper as far as that study is concerned, the kinds of results that you get when you ask people how they feel about many of these scenarios. And for those of you who, perhaps, were not there when this was presented, this is a vignette study where we asked a number of people to think about a variety of different IoT scenarios, including scenarios where facial recognition would be used, placing people in as realistic as scenarios we can through these vignette studies, requiring them to tell us how comfortable or uncomfortable they are with the scenarios, as well as telling us to what extent they would want to be notified about the presence of this data collection, [INAUDIBLE] of scenarios.

And the quick answer is that biometric data is very sensitive. People are not, in general, very comfortable with the presence of these cameras, and many people want to be notified. This is, in particular, true for cameras that rely on facial recognition.

So there's some level of discomfort when it comes to just passing in front of cameras. But when you tell people that these cameras are doing more than just recording footage about you, there is facial recognition going on-- and facial recognition, as you may or may not be aware of, covers a number of different scenarios these days, so the GAO actually has come up with a fairly broad definition of what facial recognition covers. It's not just identifying who you are, but it includes things such as looking at the demographic data of facial expression recognition and all sorts of other things that can be inferred with fairly high levels of accuracy these days.

As you probably realize, also, this technology has improved quite a bit, whether it's with convolutional neural network, deep net learning, et cetera. We're able, today, to achieve very high levels of accuracy across a number of different types of scenarios, not just even identifying who you are but also scene recognition. For instance, identifying what activity you're engaged in. Obviously, who you might be with. All sorts of things of that nature.

And so it's very clear that people, when you conduct these sorts of studies, are not necessarily comfortable with these practices across the board. There's a desire to be notified. And obviously, that would match your expectation if you think in terms of, for information, practices, principles, notice, and choice, right?

So the question then is, well, if there is a desire among many people-- not everyone, but among many people-- to be informed about these practices and to be given some choices, how do we do this, right? So there was a speaker earlier this afternoon who talked about mobile phones on this panel and how we might be able to give people more control through permissions on mobile phones.

Well, in IoT and within the context of cameras, unfortunately, those cameras are not residing on our mobile phones. So there's no uniform mechanism whereby we can inform people about the presence of these cameras and no uniform mechanism where we might be able to enable them to opt in or opt out of some of these practices.

And so work has revolved around trying to see to what extent we could actually develop an infrastructure that would support notice and choice in the presence of cameras and, more generally, in the context of other types of IoT scenarios. So this research is being conducted in the context of a larger project that's called the Personalized Privacy Assistant Project. And today, I'm talking specifically about camera notice and choice within that context, which, in its own right, is actually very rich and for which we're planning to make an infrastructure available within roughly a couple of months.

And so how does this work? Well, the basic idea, if you think about it, is not terribly, terribly complex. The basic idea is that there should be, in fact, a mechanism whereby we might be able to discover the presence of these cameras and whereby we could use this to notify users about what these cameras are doing.

Are they just capturing footage? Are they retaining that footage for a minute? Are they retaining it for the rest of times? Are they mining it to death? What are the kinds of things are they doing with the results? Who are they sharing this data with?

And so we've built an infrastructure that revolves around what we call information resource or IoT resource registry. These are registries that could be controlled by any number of entities. We envision registries controlled by building managers, by people who rent a room, by cities, by neighborhoods, by mall operators, by universities. In fact, we've deployed these registries both on our campus at Carnegie Mellon as well as in a building at UC Irvine.

[BUZZER]

So this has actually been piloted. And people can go about specifying the presence of cameras-- in this case, specify what data these cameras are collecting, who they share the data with, what sort of processing is being applied to that footage.

And along with this, we've developed clients that run on people's mobile devices. We call them IoT assistants. And these assistants can go about discovering nearby registries and resources registered in these registries-- cameras, in this case-- and inform the users about the cameras that are nearby and what practices these cameras are engaged in in terms of collecting their data, processing that data, and sharing it with others.

And so if I had to quickly outline how this works-- and I'm to obviously stop within a minute before I get kicked out-- but if you imagine that you're, for instance, at a theme park, and theme parks have been known to explore with these sorts of technologies, you might go about, and you might discover that, in fact, there are a few different registries out there where people have registered different sets of cameras. And you might be able to discover what these cameras are doing on your cell phone. And if they're opt in or opt out settings, you might be able to decide whether or not you want to opt into some of these practices or whether you might actually want to opt out of some of these practices.

And so the initial version of this that we're going to release within a couple of months will support just what I described right now. Moving forward, we've done a lot of work in the context

of the mobile apps. We've shown we can learn people's privacy preferences. And so we're hoping that, through the same type of learning, we might be able to help people configure security settings semi-automatically, occasionally verifying with them that what we're doing is consistent with their expectations.

So we've also done some real-time denaturing of camera footage. I don't have the time to show you that. But at the end of the slide, there is an SQR code. And so if you want to see a demo of how that would work in real time, whether you want to opt in or opt out of some of these processes that capture your footage, we are actually in a position today to get pretty close to doing this in real time. It's not perfect, but we're getting very close.

So in short, as it's time to conclude now, I wanted to point out that we talk a lot about IoT. IoT's very broad and very complex. Even if you just focus on cameras, there is a lot to be done in that space.

And we're hoping that through our work, we're going to empower people to abide by some of the expectations of laws that you find in states like Texas or Illinois or GDPR in Europe, enable people to become more aware of what data is being collected about them by these cameras, and also eventually expose to them settings such as the ability to opt in or opt out of some of these practices. Thank you very much.

And this is the QR code I was mentioning earlier. So stay tuned if you want to register for the mailing list that we have on this project website. We'll let you know when the infrastructure is available. If you'd like to deploy this in your street, in your city, or in your college, we'll be more than happy to make this available to you. Thank you.

MARK EICHORN: Thank you.

[APPLAUSE]

All right. So we welcome your questions. As always, if you have one, people will come around and grab your cards.

But I'll just start off with a general observation. I guess, as I sort of thought about this panel and tried to sum up what it's all about, it's very hard, in a way. Because we've got spyware, we've got third party libraries, we've got health products like glucose meters and so forth, we've got smart TVs, and we've got facial recognition. So in that respect, it sounds very disparate.

I think they're all linked by the fact that it's all about notice and control. And some of this is giving people more notice prior to purchase about what's going on, and some may be control out and about in the world, like the procedure Norman just talked about.

So for a particular question, I actually wanted to start with Peri. Because Peri cited some figures, and it's actually pretty stunning. I wanted to add a couple.

According to the National Center for Victims of Crime, 7 and 1/2 million people were stalked in one year in the US. 61% percent of female victims and 44% of male victims were stalked by a current or former intimate partner. An estimated 15% of women and 6% of men have been a victim of stalking during their lifetimes.

And then as far as the role of technology in all this, a little bit less research has been done about that. But in your paper, you cite interviews with victims of domestic violence, intimate partner violence. And basically, three of 15 people interviewed said that tracking software seemed to play some role in that. So that's 20%. That's a pretty significant sign that this is going on.

But I guess my question, Peri-- and it's sort of a question for everyone, as well, because one potential fix to the problem that you raise is this fact that something could be on my phone and I don't know about it and that the platforms have rules about icons being visible. So those could be better enforced.

But it raises some interesting questions about the role of platforms, in part because, as you mentioned in the paper, there's also stores that are not the official store provided by the platform that provide this software. And so even if the app store policed its apps better, they might not be able to completely solve the problem. But do you have any thoughts on that?

PERIWINKLE DOERFLER: Yeah. And that's an important point that I don't think I highlighted as well in my presentation as I could have for lack of time. We focused mostly on Android apps in the study, but iPhone apps also exist. And on both Android and iPhone, there's an official app store. And for Android more so than iPhone, you have the ability to download apps that are not on that store. Those are some of the more malicious ones.

But the one I demoed, for example, is freely available on the app store. And when you've got that paid version on the phone that was passed around, you couldn't find an icon for that app. There's no notice that it's recording. There's no notice that it's tracking. It's nothing at all. If you were the person that owned that phone, you really wouldn't know it's there.

And that's directly in violation of a lot of Google's policies that developers have to agree to when they sign up and put their apps on the app store. But the issue is that those policies aren't enforced.

In contrast, the same app exists in the Apple ecosystem, and those capabilities aren't there. In fact, you see would-be abusers get very upset in TrackView's customer support forums, because they can't find a way to hide the app icon if their victim has an iPhone. And that's because it's not possible, because iOS, as an operating system, does not allow that, whereas Android says you can't do that but then, as an operating system, doesn't enforce it.

So I think in that particular case, there's a pretty good reasoning that it can be done on an operating system level, and maybe it should. That's certainly not a universal answer. But I think there are some places where we could see operating systems make some changes. But with some things that are more nuanced, it's a little more difficult.

MARK EICHORN: Thanks. Does anybody else has comments on the role of the platform?

SAKSHAM CHITKARA: So I'd like to add one more thing that if you are looking for a very generic solution, then transparency. And as long as there is transparency and the users have the control, all the problems can be solved. If it's the IoT space or the TV space or the smartphone space, as long as users know where their data is going, why it's being used, they can make informed decisions. Once that is in place, once the infrastructure to do that is in place, I think we'd be in a better shape.

MARK EICHORN: Thanks. Norman, do you want to--

NORMAN SADEH: Maybe. So I was just going to add that it's clear there is a trend, if you look at platforms like the app stores, to scrutinize increasingly more closely what's being placed there. I think they realize that the onus is on them, right, and that clearly, many of the app developers are not terribly sophisticated. And so obviously, the reputation of the app store is critical.

But it's all about trust and usability, and trust is obviously a big part. So you're going to go to the app stores that you trust will not have malware on them. And so there's clearly more competition these days, even if you look at the main platforms to make sure that they're up to snuff as best as they can. It's never going to be perfect.

MARK EICHORN: I think that's definitely true. I mean, the interesting twist in the spyware example is I'm not the phone owner. I mean, I'm somebody who has access to somebody else's phone, and I'm actually trying to subvert the protections in [INAUDIBLE].

PERIWINKLE DOERFLER: Yeah, and I would agree with what you said about how, in an ideal world, when users are allowed to make all these decisions, then we shouldn't have anything that we flat out disallow. Because there's always somebody who's going to make every privacy tradeoff. And that's probably true.

But the issue is, I think, you make a critical assumption there of an informed user, of a user who is, first of all, using their own phone, installing their own applications. But even if they are a user who understands the privacy implications-- and I think that's something that we look at even trying to talk to victims who are removing stuff from their phone, is understanding what the implications of removing it will be. What will their abuser no longer be able to see?

So I agree that, obviously, consent and notice is a great next step. But it's very difficult in a space where you potentially have multiple parties involved with the same device.

MARK EICHORN: I had a question for Katie and Ian. So you all went to a tremendous amount of work. And Ian, you looked at 12 products, basically, and Katie, you all looked at five.

And so I'm just wondering, how does this scale over time as we-- and I think, Ian, from your presentation, it was pretty clear that, at this point, we're trying to apply these principles. But there's a fair number of judgment calls and sort of weighing how much sharing is a lot of sharing.

Or you talked about various tradeoffs between how many parties it's going to and the type of information and so forth.

Is privacy too context dependent to really sort of scale this up so that you have one type of test where it would work for both smart TVs and other types of products? And maybe, Katie, you can start on that, since you guys have been developing the digital standard. And I know it's to be applied in a lot of different contexts.

KATIE MCINNIS: Yeah, so one thing that we've been working through with the digital standard is this all-encompassing version of good, right? Across governance issues, which can get at some of harassment and abuse on online platforms, all the way down to whether or not you know which things are on on your phone or whether or not you were given the right consent and you know that the data is going x, y, and z places.

So it is really hard to decide, how do we test that out with a bunch of products? Which is one reason why we're proceeding with an abundance of caution and analytical skills and relying on our 80-plus years of testing to make sure we get this right with a bunch of case studies. And eventually, we're going to build this out so it's one part of our testing model, right? Like if we were testing refrigerators and some happened to be smart fridges, we're going to also, hopefully, down the road, include privacy and security as part of that assessment, because it necessarily has to be in the long term.

But it is first scaling it out for one test or one way of analyzing these products. I think that's incredibly hard. Because issues like repairability and interoperability, I think, are going to be more important for the quintessential farmer with his John Deere tractor and whether or not he can repair that product. That's not really going to be the case with a smart TV where many people don't actually fix their TVs in this day and age.

So it's going to be hard. And I think-- again, context specific. But again, our rating system has been context specific for years about, what are you really concerned about with your washer and dryer? And over the years, energy use has become more and more important.

So I think that we'll just rely on our years of testing to make sure that it scales. But I also encourage lots of other organizations to use our model in order to conduct their own tests. And that way, we hope that, by a group effort, it's going to influence the marketplace over time.

But it is really hard. There's so many connected devices, and privacy is so context specific. So it's a challenge. But we hope, five years out, that we'll have a great model going forward.

MARK EICHORN: Ian?

IAN DOUGLAS: So yeah, we're looking at our model now. And I guess it's going to change over time, because we're going to look at different devices. And as we do, we'll probably notice that there's other things that will come into play. And they'll make changes to what we've decided is the scoring system.

For now, we're trying to do a high, low, medium type model at the end of the day. But it's still hard to say, OK, how many is too many? And that's where it gets difficult.

MARK EICHORN: Great. OK. All right, for Peri, here's a question from the audience. Will a factory reset get rid of a hidden tracking app on a phone?

PERIWINKLE DOERFLER: That's a great question. "Sometimes" is the answer there. It depends on the nature of the application, how it's installed, whether the phone is rooted or jailbroken.

Anything that's come directly off the Android or Apple Store, probably yes. Most things that have come off of the open web, that depends on whether the phone has been, in the case of an iPhone, jailbroken, in the case of an Android, rooted. If it has been, then no, probably not.

The other category of apps that we looked at a lot that we talk about in the paper are things like Verizon Family Tracker, which come prepackaged with your OS if you buy a phone directly at the Verizon store, which you cannot get rid of, no matter what you do. So yes and no. The answer is sometimes.

MARK EICHORN: And are there any phones that you cannot put this tracking software on? I know there are phones where you can't get it off.

PERIWINKLE DOERFLER: Yeah, so it's specifically the apps that are packaged with the OS. So if you buy a new iPhone, Find My Friends is installed, and you can't uninstall it. Similarly, if you buy a new Android, Gmail. You cannot uninstall Gmail no matter what you do.

It's the same thing where there's some of this stuff layered on top of the OS that you can't get rid of. The two major operating systems are Android and iOS, and those are the ones that we've looked at, Android more so than iOS.

I don't know about Windows phones. I mean, it's possible that there's different things there. I, myself, have a BlackBerry. But BlackBerrys are Androids these days, so I can't tell you about when BlackBerry had its own operating system. Yeah.

MARK EICHORN: A question from Twitter for Norm. I guess, do tools to help users opt out include opting out from surveillance for law enforcement purposes? For example, in public streets or possibly secret operations. If so, have you considered those interests? So I guess this is considering that this will be rolled out in the future and will be pretty ubiquitous.

NORMAN SADEH: So realistically, we believe that the first step in terms of making this technology useful will be to focus on notification, right? And obviously, opting out in terms of government surveillance, I think that's an illusion. It's probably even an illusion in terms of notification in some regard, I think.

So the focus, however, is, number one, on making people aware of what's being recorded about them. So I think, under some regulatory regime, there is actually pressure to do that. In other

places, we're hoping that people will just find it's a good idea. You certainly find many places these days where people have a sign saying this room is under video surveillance, except that the sign is very tiny, and 90% of people in the room will never see it. So we'd like to, obviously, overcome that.

In terms of opting into different practices, again, I think, in some jurisdictions, there will be a need to obtain a form of consent. In others, that might take a bit longer. And so we're going to have to see how things evolve.

We're also hoping that people will-- so our registry can be managed, also, by activists. So we're hoping that some of these registries will just be populated by activists who will say, well, I've spotted these cameras. And whether or not the people who actually own these cameras are willing to admit that these cameras are there, we would like to see some people indicate that these cameras are there, potentially even speculate as to what these cameras might be doing.

And we would like these registries to acknowledge that some of the things in these registries are more credible than others. And so if the registry is managed by a building and the building says that it has a camera that does x, there's a good chance that, in fact, it does x. On the other hand, if it's a registry that's managed by activists, the manager will be speculating. So we'll try to also support differentiating between these different scenarios.

MARK EICHORN: I had a question for you as well about, have you thought about different scenarios for this? Because it seems like the same kind of idea could be used for other types of tracking, like mall tracking, where I don't normally have a lot of insight into that.

NORMAN SADEH: Right. So the infrastructure we've developed, in truth, is by and large being developed for a wide variety of IoT scenarios. I discussed cameras today, because cameras are obviously very broadly deployed and very intrusive in many different regards. And you have no way of knowing what these cameras are doing. You have no idea whether they do scene recognition or just capture your footage.

So I think that, in its own right, is very rich. But yes, we are actually looking at a number of other scenarios. And we've actually supported other scenarios beyond just cameras.

MARK EICHORN: And Saksham, I guess, with respect to your research, one question there would be, how do you see this playing out? I mean, it seems like 75% of people sort of refuse the third-party library tracking. So obviously, that really affects the ad-supported model. So how do you see that playing out in the future if this were sort of like if everybody had this on their device?

SAKSHAM CHITKARA: So there are a couple of alternates to this particular thing. One would be that you can have a data currency model, which people are looking into that. OK. If you want privacy, you should pay for it. Or you should pay for it in some form, either video data or with some digital currency or some actually currency.

So that's one way, but I don't see it happening in the near future. So the alternate would be limiting the data which is sent out from your device. So if people actually want to block a lot of it, you obviously can't break the entire third-party library model.

So you would want to limit what flows you would want to inform the user that, OK, this is happening. And you would want to have an option to opt out. Like right now, it's all or nothing. Either you give your data to the third-party libraries or not.

So I'd like to see a ecosystem where the users have more control over their data. The users should be like, OK, I want to block my data for this particular library. And then you can sort of limit it that, OK, you can block for this particular thing but not for every particular thing. As long as users know about it and users have control over it, that would be an ideal place.

MARK EICHORN: OK. All right. Let's see. I guess one final question, then we'll roll into Neil's presentation. But Katie, this is a question for you. Do you view the FTC's Visio decision as setting an industry standard requiring opt-in consent from consumers for ACR features? And I guess I'll take the opportunity with this question to note that our colleagues in New Jersey worked with us on that matter, and some of them are here today.

KATIE MCINNIS: That's right. I should have mentioned New Jersey in that decision. But 10 minutes, what a tight time.

Yeah, so I think that's up to, obviously, to some extent, the FTC. But we've seen that the companies-- a lot of their privacy policies were updated after the FTC decision, which also was around the start of the year. So maybe they were just refreshing everything. But I do see that these TVs have obviously tried to be in compliance with that. So maybe it is setting an industry standard.

In any case, I think, anytime your data is being collected, you should have some kind of notice. So the Visio decision is only effecting one company. But definitely, I think the other companies have tried to disclose and let consumers know what's going on.

And to some extent, consumers are benefiting from this practice. I definitely want some TV recommendations, and so I would benefit from the practice. I may not want it to be used for advertising, however. So it's good to know what's going on, I think, with your devices and what they're transmitting.

MARK EICHORN: With that, let's give the panel a hand. And our chief technologist is coming up to speak.

[APPLAUSE]

Thank you all.

NEIL CHILSON: Thanks, Mark. And thanks to the panelists. And thanks to all of you for being here today. Quite a day, huh? We started with Commissioner Chairman Ohlhausen's remarks this

morning. We had 20 fascinating, very detailed presentations, four panel discussions, a lunchtime poster session. And I'm sure, somewhere in there, you all squeezed in a quick read of our mobile data security update report. So great job.

So I just want to share three things with you really quickly. First, I'll offer quick thoughts. Basically, one thought per session of my takeaways of today. Then I'll give a really high-level picture of how maybe policymakers and engineers should think about solving some of the problems that we heard about today or generally in the technology space. And then I'll end with some quick thank yous.

So first, some observations. I learned so much today. I'm going to limit myself to one per session. On the first session about collection, exfiltration, and the leakage of private information, each of the panelists set out quite concerning scenarios involving a wide range of technologies-- IFTTT recipes, browser extensions, session replay scripts.

I was heartened to hear, at least from three of the panelists, that, upon publishing their research, companies took action to address some of the problems, at least in part. That, to me, must be really gratifying as a privacy researcher, and I think it shows the great value of this sort of research. So that was kind of my takeaway from session 1.

Session 2 on consumer preferences, expectations, and behaviors-- I think one thing that I took away is that there's a lot of unease by consumers when they encounter new technologies, whether that be connected toys or the mTurk system, and that even in cases such as the mTurk system, where it was pretty clear the user had full knowledge-- I mean, they participated in the system, and they didn't always know what the information was being used for, but they knew a lot about what information they were giving. And they knew something, at least, about that they were getting paid for it.

Even after they had done that, they felt some regret in many cases. And I thought that was interesting. Even when they had considered tradeoffs, they still felt unease about the use of their data.

And that actually rolled into my takeaway from panel 3. I mean, the idea there, I thought, that came through to me the most was the real need to look at real choices that consumers make when they're faced with two different options or when they're faced with tradeoffs. And how do they make that?

And I thought there was a lot of consistency, actually, in the reports, the studies, especially the two who are on the far end, between what they took away from the data, essentially, that consumers didn't seem to value this that much or that they seem to have a real disparity between what they said they cared about and then how they acted. Now, I think they had kind of different takes on what that meant. And I think that comes down to whether or not you think that the expressed preference is the real preference and that consumers are subsequently being sort of tricked into sharing their data when they don't really want to, because they said they didn't, or if you think the stated preference is sort of their tradeoff-free version of how they would make

decisions and that the revealed preference when they're actually faced with choices is the real one.

And to me, that kind of comes down to-- another way you might ask this is, who's right? Is it the consumer making the decision, or is it the expert looking at that decision and saying, hm, I don't know, they might have messed that decision up? And I think that perspective came out between the two presenters on that panel.

And then finally, on session 4, I would say my key takeaway there is that there's a lot of great work being done to give consumers new tools in this space, and I was heartened by the sort of range of types of tools. So we have consumer apps that try to directly limit the amount of data that's going out of your phone. We have government surveys that are studying these and publishing information so that the market can work better so that these companies who might have issues know that they're there, and they can address them.

We had Consumer Reports doing their great work, both educating consumers and also offering a sort of benchmark that companies can try to meet. And then we had university frameworks that are exploring comprehensive solutions to a very difficult problem of knowing when you're being observed in public spaces. And so that sort of broad range of tools being brought to bear on this concern, I found that quite interesting.

So those are kind of my takeaways from the four sessions. Turning to what we can do as engineers and as policymakers in this space, obviously, policymakers need technologists to help make informed decision making. And educating policymakers is one of the key reasons we have PrivacyCon.

But how can we translate this education about the problems into actual-- that technical expertise into actual policy solutions? That's a really tough question. But I have an answer for you.

It's The Lego Movie. Well, maybe not. But how many people here have seen The Lego Movie? For those of you who haven't, you should watch The Lego Movie. But I apologize. I might spoil it a little bit.

So it's an animated film. It's these little yellow-faced Lego characters who live in a world made out of Lego bricks. And it looks like a kid movie, but honestly, I find it one of the most compelling movies about tech policy and how we might address problems generally.

And so the main villain in that movie is President Business, and he's this guy. He's a dictator. He has a perfect design for Legoland. In fact, he has a cylindrical prison full of little Lego engineers who help him come up with these plans.

And it ends in a joke that I'm sure is missed by the average 8-year-old Lego enthusiast. President Business calls this his think tank, which I thought-- I love that joke. That's amazing.

So President Business believes that his expert Lego designs are better than everybody else's. He's actually somewhat puzzled and infuriated that people would do things differently, because his

way is so good. It's so orderly. And in fact, he imposes those designs on everybody in Legoland, or he tries to. Both through force and also through, essentially, propaganda-like advertising.

But there is this team of rebels in this movie. And there's a boring guy named Emmet, who's kind of an every man. And these rebels, some of them are expert builders. They're really good at crafting things on the go, on the fly, with the materials that are near them, to solve the specific problems that they face.

And Emmet is not an expert. He's actually really bad at this. And his designs get mocked by his fellow rebels all the time. But ultimately, they actually turn out to be useful, even if they're really ugly and weird.

And the main thing is that Emmet and his rebel friends, unlike President Business, they're really just trying to solve the problem that they're faced with the materials that they have at hand. And they apply their sort of local knowledge and their own creativity to solve those problems.

And so needless to say, President Business doesn't like these guys. They really mess up his plans, even more so than the regular residents of Legoland. And so he gets really tired of this, and he ultimately decides he's going to just superglue his designs everywhere. That's a spoiler, by the way. Sorry.

He's going to superglue everything into place so nobody can mess up his plans. And obviously, a fight ensues, and the heroes win, largely because they convince everyday Lego people that it's OK to make your own designs and try to solve your own problems with the tools that you have.

So what's this have to do with tech policy? Maybe the parallel is obvious. I mean, it is to me. But I'll walk through it a little bit.

I'm an engineer. I'm often asked to help solve policy problems. And there's a real temptation as an engineer to come up with the solution.

Surely we can do legal code the same way we do computer code, where we consider every possibility, and we account for it. But the real problem is that humans aren't bits. They're not billiard balls or Legos.

And so as individual humans interact, they create a sort of dynamic, non-linear, nondeterministic system that we call society. And it's not total chaos. There's lots of patterns-- social and ethical norms, markets, institutions, political structures. These are all patterns in society.

But these patterns are less like the sort of perfectly predictable pendulum, and they're are a lot more like eddies or whirlpools in a stream. You can identify the pattern, but it's really hard to predict or control.

And so when trying to solve problems in these environments, a sort of single approach can be a lot like trying to grab an eddy in a stream, like you do something, you cause effects, but you may not have the effect that you want. And when top-down sort of solutions do work, they often work

by freezing the stream or damming it up or, in The Lego Movie, supergluing everything. And so that solution can often be, at times, worse than the problem that we are trying to solve.

So today, we learned a lot about existing and emerging privacy problems. But I encourage all of us to think, when we're thinking about how might we address problems in the real world and in the policy world, which is sometimes like the real world, that we would consider how to address these problems in a way that maximizes experimentation, that tries to make the most of knowledge and creativity of many people rather than a few.

And these solutions may be social. We actually heard talks about-- the Consumer Reports is a good example of this sort of social effort. It's not a regulatory effort. It's a feedback loop.

Or they might be technical. We actually had a lot of presentations about technical solutions to these issues, including mentions of things like the blockchain, which might be big-picture ways to address some of these issues.

But I do think that these sorts of very widespread uses of local knowledge and local decision making can be much more robust. And more importantly, they tend to avoid freezing solutions into place, even when the world keeps trying to move on.

So let me just end with a few quick thank yous. Thanks for indulging my movie review. First, thanks again to the terrific presenters. Thank you for sharing your insightful creative and careful work with us.

Second, to the FTC team. It was a real pleasure to work with everybody to review an enormous and very terrific set of filings, submissions, and picking out the people who got to present to us today. Special thanks to our moderators today, especially Dan Salsburg, Kristen Anderson, Mark Eichorn, and Yan Lau.

The lunchtime poster session was hard work, and it was done by Jamie Hein and Tina Young and Joe Calandrino. Other key contributors were Tim Daniels, Dan Wood, Aaron Alva, Crystal Peters, the media team back there. And our paralegals, of which there is a long list. But I want to say their names. John Abe, Christine Barker, Annie Blackman, Courtney Brown, Amber Howe, And Aaron Coughman.

Can we just give everybody who put this program together a round of applause, please?

[APPLAUSE]

And finally, thanks, again, to all of you. Thanks for coming out or watching us from home, as the case may be. And we hope to see you at PrivacyCon 2019. Thank you.

[APPLAUSE]

[MUSIC PLAYING]