

FTC PrivacyCon 2018
February 28, 2018
Segment 2
Transcript

KRISTEN ANDERSON: Everyone, please take your seats. All right, we're about to get started with our second session of the day. This session will be about consumer preferences, expectations, and behaviors. We have a great panel of presenters this afternoon.

Starting with Jingjing Ren. She's a PhD candidate in computer science at Northeastern University. Then we have Kris Micinsky, who's a visiting Professor in the Computer Science Department at Haverford college. Next we have Emily Reynolds, who's a researcher at the University of Washington. Next to her is Pardis Emami-Naeini, who is a PhD student in the School of Computer Science at Carnegie Mellon University. And finally we have Yang Wang, who's an Assistant Professor in the School Of Information Studies at Syracuse University.

So just jumping right in, we will see first get started with a presentation by Jingjing Ren.

JINGJING REN: Thanks for introduction. So today I will talk about our longitudinal study in privacy leaks across Android apps. So to start, the motivation of our work is that really, of everyone in this room, we use many useful apps in day-to-day life. And we constantly reminded to update those apps. We are often told that it is for the benefit of bug fixes and performance improvement.

However, I often wonder when I update this app, is there any privacy change. So the reason I'm asking this, is that our research group helped with the production of a documentary film called Harvest that follows a regular woman named Jenni's life for one week. And the way [INAUDIBLE] her phone to detect a privacy leaks from the apps she was using.

What we find here is that there's 3,735 location leaks in one week. And most of them happen in the background. It comes mostly from the craft installed app called Joanne. And when we investigated further that's not always the case. We collected different versions of this app, and we looked at how many times it would leaked over time. And in the beginning of 2016, there's no location leak. Then later on it was found during this five minute transaction, the location was leaked for lots of times. Then the intensity reduced in recent version.

So this is just a one app. What the really big picture we are asking is how mobile privacy evolves over time. So before I answer that, I'd like to define what a privacy is in this context. We specifically focus on three aspects. First what information is shared-- mainly we focus on personal identifiable information. Then how is it being shared, whether it's stand over encrypted or unencrypted channel.

Another aspect is, who got this information. Whether it's the first party, the app developer, or it's a third party like the checking system. So why do I care about these different aspects? Well, we've identified three types of way to evade a user privacy.

First of all, the data aggregation-- a different party may be interested in users' PII. Then when you send these PII over an encrypted channel it's susceptible to the eavesdropping attack. Now, we've created the definition for privacy. How do we measure that? And still to look at how it evolves over time.

First we need a large data set. We collected 512 Android apps that covered over 7,000 unique versions all across the last eight years. And in order to measure privacy we first need to interact with those apps. It's good if we can do it manually, like a regular user do. However, it's not a very scalable for 7,000 apps. So what we did instead is to combine automated and scripted interactions. We use Monkey to generate a render event, and we went for the case data, we needed to log into the account to use most of functionality. We manually log into the app and replay those events across versions.

So after the interaction, we were able to induce privacy leaks to the internet. And we use manual intermediate proxy to intercept both encrypted and unencrypted traffic. From there, how do we detect a PII? So here we're using the system called a ReCon data [INAUDIBLE] device presented last year, which uses virtual learning algorithm to detect the PII leaks without knowing them ahead of time. And then we manually validated results.

So after this experiment we were able to collect the different parts of the attributes of all the each app. First we have the set of PII types data that leaked by the app. And each app has a different version that is corresponding to their release time that can be ordered chronologically. And for each PII leak, we distinguish whether it's done through encrypted channel versus unencrypted ones. Also it matters whether the first party gathers the information, or third party gathers this information.

So this is our different privacy aspects about each app. Now, I'll give a concrete example. The Pinterest, it's a really popular app that has to be used by millions of users. Here we have one PII leak at advertising ID. If you look at the last part the ID was leaked 12 times to different parties through encrypted channel. And we have this piece of information for all the PII types leaked by this app.

So the first thing you might notice is that the password leak. Two versions sent a password to an undisclosed third party. And the way it affected millions of users, we reported this to developers. They fix it within a month. So that's one part.

Another significant increase you will notice is that the types of PII leaks increased in more recent versions. This included gender, location, advertising ID, and so on. So not just the unique types. If you look at a specific PII leak, for example, the Android here, the frequency also changed. It was leaked two times in 2016. Then the intensity increased into over 200 times. And this is only 10 minutes into action. So this is really an upstanding find going to location tracking.

So not only that, another thing we looked at is the [INAUDIBLE] used to transmit the PII. The location half the time it was sent in an encrypted channel. That's just one app. The take away is that we see lots of variance in privacy leaks across versions. Now we have over 500 apps. What's the big picture? Our folks aggregated findings in our study.

First is the PII leaks. It can change substantially across versions. Not just a number of unique PII types, but also the frequency of leaking specific PII. Another thing we looked at is the issue [INAUDIBLE] adoption in the mobile space. For the mobile apps it's especially harder because they make harder coded protocols, and it requires extra effort to update those domains. And what do we find is that it takes years for the app to adopt [INAUDIBLE] protocol once the domain is supported.

Another thing they started partly checking. We all know it's pervasive and this is a consistent finding, however, it is also evolving over time. We see that the checking IDs, they are moving towards [INAUDIBLE] IDs, which is good. However, we do see evidence of over 100 domains that have the capability to build a permanent linkage between a unique ID they used to track you, and the basic identity related to PII, like email address, real name, gender, location, and so on.

So that's a different aspect. Does this answer the simple question, is privacy getting better or worse? Well, this really depends on how your [INAUDIBLE] privacy and what's important to you. [INAUDIBLE] paper we went into content different aspects, and I highly encourage you to go to check out. Here I only show the case that the combined [INAUDIBLE] from, the PII types, and the destination domains that received those PII. So the curve here, it goes up. The higher value means more risk, which means the combined risks actually wasn't over time.

Again, this is mainly due to more PII were leaked and more domains were contacted. So to conclude, by our definition of a privacy, we find that it [INAUDIBLE] time. However, even if this situation, we still recommended you to update the apps for security reasons. There's a clear need to continuously monitoring your mobile applications using existing data system like ReCon [INAUDIBLE] other research tools.

To this end we also develop a web interface-- a web interface that will provide a customized preferences to help understand how the privacy changed across versions. And this is also important both for the users and us and the developers. Because sometimes they are not necessary for malicious purpose, and they could fix the box based on the findings. So with that, I conclude my talk. Thanks for your attention.

[APPLAUSE]

KRISTOPHER MICINSKI: Is there a little slide advancer thing somewhere? It's this thing, perfect. All right, so I'd like to start out by saying that this was work with many of my collaborators at Maryland. So to introduce this space I want to pose this example app to you. So we have this little app that's going to come up, and it's going to do something like help you order coffee. All right, and the app is going to have a few different buttons that it shows you. And one of those buttons is going to help you locate some coffee shops nearby, something like that. Maybe share the app with your friends.

But then another button is going to help you do some voice ordering or something like that. And I think most of us can agree that if we click the Voice Order Coffee button, then we would probably reasonably expect that our microphone might be used. But maybe we would not expect

that if we click this Voice Order Coffee button that our location might also be used by the application.

And so this is the kind of setting that I really wanted to address in this work. The question was, how do users reason about these uses of various permissions and different privacy relevant inputs to an application? And is that what applications are doing? So with respect to some notion of the actual user interface, and the way that the user clicks through an application, how are apps actually accessing the data?

So to do this and study this, I do two separate studies. So one of them is an app study. And on that we measure 150 Android applications using a special tool. And that tool is going to help us study from the way that an application runs, what permissions it uses, and under what circumstances. And that's the important part is that this tool is going to help us really dig into when an app accesses some information, when and why did it occur with respect to the rest of the apps user interface.

And then I'm going to run a user study to understand, based on the results of this application study, are users expecting that the way that the apps are accessing information with respect to its user interface, does that actually match what users might expect based on that user interface? And that overarching question here are, are apps doing the right thing?

So the first part is doing this app study. And the way that this works is we have an instrumented app. So an instrumented app is an app that we have a special tool that takes and inserts some code to log relevant portions of it. So it logs things like buttons you click, pages that transition, when backgrounds event handlers occur, things like that. And then we actually give that instrumented app generated by our tool to a user. And the user installs it on a phone, then they run the application, and they generate a set of log files for what the app does.

We take those log files and we throw them through this visualizer. And the visualizer produces a graph of how the application works. So it says for example, these are the various screens, these are the buttons, and there are edges in the graph for when one of those things happens after another.

So the results of our study are shown in this graph on the right. We have the shades of orange and yellow. Those indicate foreground uses of permissions. And going down these rows, we have all the different kinds of permissions we found. So for example, in the top column we have a microphone that is only used interactively, and the orange is clicking-- after you click on something. And the yellow is page. So during the duration of the time that something was paged on the screen.

So we found that the observation here is that a lot of these resources, maybe the ones that I might think are actually maybe most sensitive, are already used pretty much interactively. We only found one place where the camera was used in the background. And that was because it was an authentication app where if you failed to enter the password correctly too many times, it would take a picture of you in the background. So apps are already using some of these resources fairly interactively.

But then some could be used more. So there are some permissions where the app just fetches the data before it ends up getting used. And then the fact that it was used, gets shown to users later for example. And then there are some that are really rarely used in an interactive way at all. Things like the power, information about the application, and unique ID of the application. And I'm not really sure why this is, honestly. Although I suspect it might be because it's hard to explain to developers the way that the apps are structured is not amenable to being able to explain them in an intuitive way. But that would be something to look into more.

And next to go along with this app study, I performed a user study. And the thing that we want to measure here is do the user expectations about when these various permissions will be used-- does that align with the patterns that we observed in these applications? So for example, if something is seen very interactively used in the applications that we study, do users also expect that that's the case?

And to do this, users watched slideshows of a whole bunch of different configurations of the way that these apps might use information. So consider that we wanted to answer this question, for example. This was one of our hypotheses, among many others. But we might ask, is a background use of a permission expected after you have had a prior foreground use of that permission?

So you might see a scenario-- or the user would see a scenario-- like this. Where they would first see an app description, and then they would see what we call a user action. And a user action is some event that happens in the application. So in this case, it's that they see the Start screen of the application, and we show them that the Voice Order Coffee button was clicked on. And then immediately they see a screen pops up that asks them for permission to access the microphone, which is what you would see now with an application.

Now we also measured scenarios in which this dialogue didn't appear, to understand if for example, the dialogue was truly necessary under that circumstance. And then we asked them questions about what resources they expect might be used on a Likert scale. From for example, do you think that your microphone is being used right now? Definitely yes or definitely not. And then things like, your location, and then a few other things.

We asked them some distractor questions as well. So that it wasn't just about security. And then we also, after this first event, had a second user action, or a second event that occurred. So in this case the user takes the app and puts it in the background and they see the home screen. And then we ask them those same questions again, maybe in a slightly different order for example. But we might want to measure now, does the user still expect that the microphone might be used when they're on the screen.

Now what we found, the first I think won't be surprising, which is that after you click on something or directly interact with it that cues you in that a resource could be used, users are much more likely to expect that that use could occur. So in our case, it was 106 times more likely to be a little bit more likely to expect that that resource would be used for a click than if it was used in the background, if you didn't get any notification for it. We also, and this is a thing that I think is most interesting to me, but makes sense if you think about it for a while, but we found

that this policy of asking permission for something the first time it's used can condition users into believing that it might be associated with that given UI element.

So for example, imagine an application that has a map screen. If you ask for permission to use the location on the screen that, for example tracks your run on the map screen. And then you stop your run. You put your phone back in your pocket. Users will be less likely to expect that the location gets used when you're not doing that thing. And I thought that was pretty interesting.

So taking away from this, what can we say? So I think that we can say that for many things that users expect, for example the camera and things like that, applications are already using many of these resources interactively. And we would just recommend that this be the case. So for example, some auditing tool, maybe like the one that we built, or other mechanism could be used to enforce that this happens. And there's been other work on that.

We also would say that interaction in some scenarios services authorization. So if you do something like click on a picture of the camera, and then you take a picture, that should be fine. You shouldn't have to ask the user for a separate permission to do that. And this way we can avoid things like habituation so that users start to ignore security warnings, because they think the operating system is not very smart about it.

And the last one is to try to make background usage predictable in the sense that we'd like to help users understand it even more. All right, so thanks a lot.

[APPLAUSE]

KRISTEN ANDERSON: Thanks very much. Next we have Emily McReynolds, and she'll be talking about connected toys.

EMILY MCREYNOLDS: Hey, everyone. Quick disclaimer-- so I spent the last five years at the Tech Policy Lab doing really interesting research like that. And I recently moved to a tech company, and I'm in their research department now. But this doesn't represent them. This is the research we did over the last two years at the Tech Policy Lab.

So Toys That Listen. It's a fun, catchy title, and I want to acknowledge my collaborators. We had three undergrad students who helped with this work. And then Maya Cakmak and Franz Roesner are professors at the University of Washington.

So in 2015 we had a bunch of smart toys coming on the market. Hello Barbie is pictured there. Mattel and its makers hoped that it would be the toy of Christmas 2015. You have Cognitoys Dino. This little dinosaur that is actually connected to IBM Watson's AI, to answer questions. And then My Friend Kayla. And these were the smart toys. And this Wall Street Journal article asks, toys are getting smarter, should we be worried? Yes.

And so just to back up a little bit and give some context. When we talk about smart toys, particularly when we talk about smart devices today, there can be all kinds of meanings. And so in the 1970s chips started getting small enough that we could put them into toys. And in Teddy

Ruxpin we had our very first smart toy. He seemed smart. He could read you stories and his mouth moved along in time. And then later we had Furby in the '90s, who learned to speak English. But these toys were not originally connected to the internet. And they certainly weren't the smart toys that we're talking about when we look at these guys today, which are connected to the internet, which include artificial intelligence, and sometimes even machine learning.

And so then we have to wonder well, what's the big deal? Why are people worried about these? We're worried because kids have them, yes. And data is being transmitted to the internet and into the cloud. And onto someone else's computer. But we also are seeing data breaches. So some of these aren't as secure as we'd like. This was the other thing that happened in 2015 that got us motivated and interested in doing this research. A hack exposed, it turns out, millions of pieces of data on kids who were using these tablets. In fact the FTC just settled with this company last month. And they hacks continue.

So there was an internet connected teddy bear whose database was just left wide open on the internet for anyone who would like access to every voicemail that any child, parent, or adult had ever left on the system. And more recently, My Friend Kayla-- fun story, we actually didn't use My Friend Kayla in this research, because it was insecure. So we knew that people couldn't possibly have the correct mental model about how they worked. But eventually My Friend Kayla was banned in Germany as a spying, espionage device, because it didn't clearly label that it was collecting data, and that it was transmitting data. So maybe if you hold your phone up people understand it might be able to record. But if I bring a doll to your house, it's not so clear that it might be recording you.

So we had some research questions such as what are the privacy expectations and concerns when parent are thinking about connected toys. How do children understand these toys? And how do parents want to control these toys? And so what we did was hold a user study. And in that user study we had Hello Barbie, and Cognitoys Dino, and we brought pairs of parents and child into our lab. And so we would have the parent and child sit down. And we'd show them the toys. And then the child would go play with the toy, and the parent would set up the Hello Barbie app so we could then ask them questions about their experience with these toys. And try to understand how they understood these toys and the implications of their internet connection and those kinds of things for privacy.

So a little bit more about Hello Barbie is helpful. This is actually the picture on the box if you were to buy a Hello Barbie. It identifies that there is a microphone speaker, and that these tricolor LED lights are embedded in the necklace. And they do light up. So when you go to talk to Hello Barbie, you can see her necklace flashing. And in order to talk to Hello Barbie and to get a response, you have to hold down the belt buckle. But nowhere on here does it say when you hold on that belt buckle, you're being recorded. Everything you say is going into this web interface. So your parents or the other person on the other side, they do have a login. They do have good authentication involved. But everything you say is being recorded and can go be listened to.

And so what we see here is that it's organizing conversation, it contains recordings of everything. Anytime you hit that belt buckle. And it also encourages sharing on social media. So you have a

nice easy link to Facebook and Twitter there. And all of that informed the questions and the responses we eventually got.

And this is Cognitoys Dino again, one of the things that was really good was his mouth lights up and you actually press his belly, which the kid's got a big kick out of, to talk to him. And so for the parents we asked what they expected to be in the toys privacy policy, how they felt about the ability to monitor what their child said to the toy, and would they share what the child said on social media. And this was among many questions, but these were some of the interesting responses we got.

And then we also asked the kids question. What would you talk about with the toy? Do you think the toy can remember what you say to it? Would you tell the toy a secret? And these questions are trying to get at what the kids privacy perspectives are. We were working with kids who were ages 6 to 10, so we didn't want to influence their perspective by asking how do you feel about your privacy. So we tried to elicit that information through these kinds of questions.

This is the email you get from the system when you get it set up. And I think it's really good in that it focuses on giving permission. So THIS isn't just your normal, generic, I accept the terms and conditions, but it actually has in bold right their, it focuses on permission giving. It also explains how they use the recordings and why they need them-- in order to share them with the parents and to improve their services and technology. And it reassures that they don't use these to contact or advertise to children.

But the parents had some important concerns. They wanted to know where the recording was going. They were concerned about who can see it or get the information. On the other side, a parent said look, we have all much stuff to go through right now. I don't have time to go through all these. It's one more pile of media, and I'm not even going to look at it. So we have to rethink how we expect people to handle this information, because you can have notice and consent, but if the control mechanism is hey, you can go in and delete the recordings. And they're not going to do that. That's a significant concern.

And had interesting responses. Because remember, there's no way to really notify them in this situation that they're being recorded unless the parent tells them. So this parent said, hey, did you know that Barbie doll, when you're all done everything that you share with her would end up on the computer so we could talk about it. Would that make it fun for you? The kid was not so enthusiastic. Finding out that they were being recorded, they thought that was actually pretty scary. It was a great moment in our study.

And then back to that sharing thing that I identified earlier. We got some great quotes from parents. I think past a certain age you need to respect your children's privacy. And the child is not consenting to the questions they ask Barbie to be shared on Facebook. So really good perspective on this sharing across social media. And we asked the kids the questions you saw earlier. And asking them if it remembered. You can see some of the responses here. Maybe, no, not quite sure. Would you tell the toy secrets? Yes, maybe, no, remember their parents in the room. And they're also talking to strangers. So there's some implications there for the answers we're getting from the kids.

And do you think your parents could hear what you said to the toy? One kid said, probably. The rest of that quote is, "if it's recording me." And we are like, yep, you got it. You figured this stuff out.

So there were an important connection made by parents and even some of the kids to other connected devices. They understood that there were potentially implications when talking about Siri or Cortana or Alexa or some of these other AI devices. And I think that leads to some really good recommendations around making sure we provide recording indicators, thinking about whether or not we actually need to store this data, how long do we need to keep it, can't we delete it after a certain amount of time, particularly if we know people are going to have the ability to review it.

And then think really carefully. Kids we're playing with these. Some of the comments we got were hey, this toy is great, but Siri does way more than this. So my kid's going to be bored in five minutes. So really we need to think about how non-toys are interacting with our child users. And so with that, I'm going to thank one of our great funders, the Consumer Privacy Rights Fund at the Rose Foundation. And there's my contact information. Thank you very much.

[APPLAUSE]

KRISTEN ANDERSON: Thanks very much, Emily. Next up we have Pardis. And she'll be talking about privacy expectations and preferences in an IOT world.

PARDIS EMAMI-NAEINI: Hi everyone. I'm Pardis, a PhD student of computer science from Carnegie Mellon University. For the next 10 minutes or so, I'll be talking about people's privacy preferences while interacting with internet of things data collection scenarios.

This is a joint project with Sruti, Martin, Hana, Lujo, Lorrie, and Norman. The internet of things is a network of sensors and smart devices. These devices are everywhere. As you look around, you may find a camera in this room. You may have an Alexa, a smart thermostat, or smart doorknob in your house. Or you may have a Fitbit with you now.

However, it's not always clear what data IOT devices are collecting. Or what they will do with that data. For example, by looking at this light bulb, you cannot tell that it is keeping track of when you enter and exit your office. And sending this information to your boss. In our research group at CMU we are designing a private assistant. This assistant will inform people about nearby data collection scenarios, and give people choices to control their privacy. Your smartphone that has our application installed on it, will talk with nearby IOT devices and answer your questions about you data practices. For example, what are these devices collecting? With whom are they sharing my data? Or for how long are they keeping my data?

To design this private assistant we need to satisfy some design goals. The first goal is to inform people about nearby data collection scenarios without overwhelming them. For this purpose, we need to know what people want to be notified about. The second objective is to give people choices to control their privacy. For example, we would like people to be able to make decisions to allow or deny data collection. But once again, we don't want to overwhelm people with

information. So we need to know what factors are important for making this decision. And the third goal is to predict people's privacy preferences so that we can automate privacy decision making. For example, If it can accurately predict that people prefer to turn off an IOT device in a particular situation, our system can just do it for them.

To help us meet these goals we ran an online user study. In our study we asked participants to imagine themselves in some hypothetical data collection scenarios. This kind of a study is often called a vignette study. In a vignette study participants are shown short descriptions of hypothetical situations and are asked some questions to elicit their attitudes toward those situations.

This is an example of a scenario a participant could see. The words inside the brackets are the factors. You're at work. This building has cameras that are recording video of the entire building. The video is shared with law enforcement to improve public safety, and they will not delete it. We instantiated the text in red in different ways. For example, by changing "work" to "home" or "library," or by changing "improve public safety" to "determine possible escape routes" or "optimize the number of staff."

We recorded participants for a 15 minute survey using the Mechanical Turk platform. Mechanical Turk is a crowdsourcing service where people are paid to perform short online tasks, like fill out surveys. We recorded 1007 Mechanical Turk participants from the United States. Each participant was shown 14 scenarios. And after each scenario we asked some questions to understand how often they would want to get notification about data collection in a scenario like this, how comfortable they are with that data collection, and if they would want to allow that data collection.

To interpret their results we built some statistical models that explain people's preferences for getting notification, comfort with data collection, and willingness to allow data collection. We found many things about people's privacy preferences. I will tell you about just a few, and I will strongly encourage you to read our paper on that.

Our results showed that in the model to explain people's preferences about getting notification, people are more likely to want to be notified when their data is being shared. They're also more interested in getting notifications when their biometrics are being collected. On the other hand, they are less likely to want to be notified when they find the purpose of data collection beneficial to them.

In the model to explain people's comfort with data collection, we found that people are more comfortable when data is being collected in public locations. They are also more comfortable when environmental data has being collected. For example, room temperature. However, they are less comfortable when data is being collected in private locations such as their houses. And they are less comfortable when their biometrics are being collected. For example their fingerprints.

And in the model to explain people's desire to allow data collection, we found that people are more likely to want to allow data collection when they find the purpose of data collection

beneficial to them. However, they are less likely to want to allow when their data is being shared. Some of the results may seem intuitive to you. In the next few slides I will explain. But just knowing any of these results is not enough to explain people's preferences.

So one question that some of you may have now is, after all these results, what are the most important factors in explaining people's privacy preferences? Is that the data type, the location of data collection, the purpose of data collection, or the retention time? Well, there's no individual factor that is perceived to be the most important one. Even if we start looking at some with specific scenarios, it is rare that a single factor is enough to explain people's preferences.

However, we found that the combinations of some of the factors explain people's privacy preferences better than any individual factor. For example, in the model to explain people's desire to get notification about data collection, we found that the interactional data type and user perceived benefit is the most important combination. For example, people are more likely to want to be notified when their fingerprints are being collected for a purpose they don't see as beneficial.

In the model to explain people's comfort with data collection, we found that when data type is combined with how realistic people think it is that a scenario is happening today, is the most important combination. For example, when video is being collected in a scenario that participants believed could be realistically happening today, they are more comfortable. And the combinations of factors that go a long way toward explaining whether people will be willing to allow data collection are-- data type combined with the location of data collection. For example, when people's fingerprints are being collected at the department store, they more likely want to deny it.

As I said at the outset, our long-term goal is to design a privacy assistant to help people control their privacy. Before we can build an effective privacy assistant, we need to better understand people's preferences about their willingness to allow data collection or how often they would want to get notification about it. In our study that I just talked about, we found that people's preferences are complex and context dependent. But that combinations of some of the factors like data type combined with the locational data collection can sometimes be helpful in explaining whether people would allow data collection, or whether and how often they would want to be notified of it. Our future privacy assistant will help people make the decisions they want to make. Thank you so much for being here today.

[APPLAUSE]

KRISTEN ANDERSON: Thanks very much, Pardis. Next up we have Yang Want, who will be talking about privacy violations in crowdwork.

PARDIS EMAMI-NAEINI: Hi, good morning, I will be talking about our recent research looking at how people experience privacy violations on various crowdwork platforms. So just to give folks who are not familiar with the crowdwork-- so I'm going to use crowdsourcing and crowdwork interchangeably. This refers to the practice say, if you have something you want

people to help with-- say a survey, or you want somebody to design a logo, you can ask people to help you with that.

And there are various platforms that allows just ordinary people to contribute to certain tasks that you want help with. I show here some famous examples-- Amazon Mechanical Turk, which the previous speaker talked about, CrowdFlower, and more broadly speaking the gig economy with Uber and Airbnb, there's just new services popping up every day allowing this kind of crowdwork.

So our research-- prior literature has shown that there are various privacy risks on these crowdsourcing or crowdwork platforms, but they don't really look at the real experience faced by, where you can [INAUDIBLE] by these crowdworkers, and that's really our focus. So our study not only looked at the kinds of concerns crowdworkers have doing work on these platforms, but also what are the specific experience they actually encountered. And we think that's very important.

We use Amazon Mechanical Turk, or MTurk in short, as a case study, mostly because MTurk is a very popular platform for this kind of crowdwork. what we did is just a baby step uncovering these issues. We did an MTurk survey. We recruited about 400 crowdworkers around the world. And MTurk has workers really around the world, not just in the US. And the main research question is that, what are the things they actually experience that they felt their privacy was violated.

For those of you who are not familiar with MTurk, this diagram basically illustrates the process. So there are two types of users on MTurk. there are task requesters. So these are people who want to get help. And then there's other type is the crowdworkers who provide their contribution to these tasks.

So as a crowd requester, you can divide up your project into smaller tasks, and then you're going to publish these tasks on the MTurk platform. The workers can select which task they will like to do. And then, after they submit their contribution to the task, the requester will have the authority to say hey, do I like this. If I like it, I'm going to pay you this amount of money that I specify in the task description.

So I would recommend you to look at our papers for more analysis. But here, for some time, I'm just going to go through a few examples looking at different types of violations that people reported. We grouped different privacy violations leveraging Daniel Solove's "Taxonomy of Privacy." So essentially there are different types of violations.

So the first type is information collection. This has to do with collecting sensitive personal information. So this is a quote from one of the respondent. I mean, if you read it, basically what happened is the task was about a study looking at people's attitude towards getting health insurance. But part of the task required them to upload a picture of their health insurance card. I'm guessing the reason why they asked crowdworkers to do that is because they want to make sure that people who are actually participating in the study have some experience was health insurance. So providing a picture of their health insurance card is a qualification.

But arguably, I mean they don't really need that information to participate in the study. And in this particular case, the participant was very upset about giving up this very sensitive information. And one thing to note is that prior literature suggests that a lot of these are crowdworkers-- they are trying to make ends meet by doing these small tasks, earning some money, oftentimes under minimum wage. So in a way, they're really vulnerable, which makes these sort of issues particularly troubling.

Another type of violation is information processing. As you can imagine, this happens after the data has been collected. So this a very interesting example. The task was about posting pictures of you doing things like smoking and other things. They didn't know, or they didn't tell the workers, they're going to use these pictures for art exhibit. So a month later, they found out and it was a negative surprise.

Another category is information dissemination. This has to do with, once the information has been collected from these tasks, they're going to further disseminate-- either sell them or share with third parties. And the Amazon user policy actually clearly says, "you as a task requester should not request personally identifiable information." and they explicitly say that you cannot collect emails and phone numbers. And low and behold, requesters still do these things. And they still request email. So maybe Amazon should beef up their enforcement of their own policies.

Next category-- invasion. This, in our context refers to things like getting spam messages or even physical stalking. This was one participant talk about over time through different tasks of the same requester, she's giving up both her pictures as well as her address. And which makes her very concerning, because once you have all this information, you compile them together, you can really track individual users.

The last category is what we call deceptive practices. This has to do with a user encounter, phishing attacks, malware, where download some sort of malicious software. So in this particular case this participant was talking about getting phishing attacks. Essentially the task was trying to trick people, disclosing their login credentials.

OK, in addition to the reports we heard from our participants, a member of our research team actually worked as a crowdworker on Amazon Mechanical Turk. And just to get a sense of the kinds of things you might encounter as a real worker. So I'm going to show you some examples. So this is one example where-- I don't know whether you can see those clearly-- basically as the crowdworkers to dig out some personal information for a particular individual. We call this human flesh search. This is interesting because this is a case where it's not the crowdworker's privacy are being violated. I mean, all of us-- any citizen's privacy could be violated if somebody wants to launch a human flesh search on these crowdsourcing platforms.

Another example-- this is the task where they ask crowdworkers to upload videos of themselves, but in order to do that, they have to download a third party software. The thing is, most crowdworkers-- they're probably not very technically savvy. So they're not really in a good position to determine whether the software is malicious or not. So this is a real security risk.

And my final example is this-- it's not from US, it's a Italian example, in fact. And this task is asked the crowdworkers to transcribe a receipt or some sort of form. So one column shows people's real names. The other column next to it is essentially the Italian social security number. This is another example where other people's privacy might be violated through crowdwork platforms.

Another interesting aspect of this example is the looming GDPR. It's come into effect in just a couple of month, and Amazon really have to deal with the issues, because they're dealing with EU citizens. So I think there are two major takeaways from this early research. One is that the kinds of privacy violations we observe from our study are not just reported from workers from a specific geographic location. Meaning that it's very common across the board, which suggests that these issues are system-wide issues. So that's one.

And second, as I talked about earlier, many of these workers are vulnerable. They are there to earn some quick money. And oftentimes they don't really think about their privacy very carefully, which put them at these kind of risks. In terms of implication for technology and policy, we would advocate that there should be tools that enable or empower crowdworkers to be more mindful about potential privacy risks in these tasks. That's not something they're currently considering, partly because the interface on MTurk doesn't really speak to the risks.

Another aspect of tool building is to help requesters. I would tend to think that most requesters are not malicious, but they also haven't really thought about the privacy implications of their request-- their tasks. So some way for them to clearly communicate who they are and what kind of information they need from the task, and how they're going to use this information, would be very valuable.

And lastly, while Amazon has these user policies, but they don't really require the requester to say anything about their data practices. We think that that's either something they should advocate as a best practice, or something that they should require the requesters to do.

So with that, I would like to thank my co-author, Huichuan Xia, Yun Huang, and Anuj Shah. And most of the hard work is done by Huichuan Xia. And then if you're interesting know more about what we did, here is the paper we just published. Thank you.

KRISTEN ANDERSON: Thank you.

[APPLAUSE]

Great, thanks very much all for a great set of presentations. We're going to get started with the discussion session. As we mentioned in the first session, we have Kristen over here, who has some question cards. They're also available outside in the hallway. If you would like a question card, she can come and hand one to you. And we'll collect them, and she can hand them up to me so we can incorporate your questions into the session as well.

So in the first session of the day, we heard about several practices that were taking place behind the scene that affected consumers' data. And during this session you guys all told us more about

what consumers are expecting, and what their preferences might be, what might affect that, and how they behave under certain circumstances. So I also invite you guys to ask questions of each other and to interact as much as you'd like.

But just starting off with a couple of questions directed to each of you individually. I'll start with Jingjing. So you found that consumers mobile privacy is less protected now than eight years ago - that there's more PII collection, leakage, more third-party sharing, slow HTTPs adoption, more cross device tracking. So given all those findings, what do you think are the most important things that consumers need to know about their use of, and their updating of apps? And what are the most effective ways for them to better inform and protect themselves?

JINGJING REN: Yes, so that's a really good question. The first step of actually protecting your privacy is really understand the apps you are using-- what information they are collecting about you. So that's why in the end of the talk I provided these web tool. Among the popular apps we've looked at, we provide all the transparency into the privacy leakage. And also moving forward you should-- consumers should use the container use system like ReCon, Lumen, other research projects that in real time monitor the mobile apps. And our long-term goal of this project is really to develop something that can pop up and let user know the privacy changed the way [INAUDIBLE] the installation of the mobile apps so users can make informed decision.

KRISTEN ANDERSON: And how easy are those tools for laypeople to access and use? Is it pretty user friendly, or does it take some technological skill?

JINGJING REN: It's really a challenge to build those tools, because we want collaboration from different party. For example, it would be ideal if we can do it in a place [INAUDIBLE], because they are the one who distribute these apps. If we can work with them to [INAUDIBLE] by this tool and make it user friendly, that will be very helpful for average users. Yeah, I think that's the main challenge right now.

KRISTEN ANDERSON: So where are they available now? Where can people go to find them if they're interested?

JINGJING REN: So the web tool provided a study was done, but it's an ongoing thing that we still need to develop.

KRISTEN ANDERSON: OK, do you have a sense of when that might become available?

JINGJING REN: Well, it's not quite.

KRISTEN ANDERSON: Still under development. OK. And then have you actually obtained feedback from consumers or others about the app versions tool that you did develop?

JINGJING REN: So that tool was released not a long time ago. We didn't really get a lot of feedback. But we do talk about that with different communities. And one comment say they always trying to understand is how do to define the privacy-- what is important, what is not

important. And there's no single answer to that. It really evolves for a user's perspective, for a developer' perspective, they want some kind of trade off between the convenience and privacy.

KRISTEN ANDERSON: It's a tough balance?

JINGJING REN: Yeah, it's tough.

KRISTEN ANDERSON: All right, and moving on to Kris. So you found that certain user behaviors could be a proxy for consent. So there's no need for some of the just-in-time notices or authorization or requests. For example, a consumer clicking on something like imports, contacts, means that she actually does consent to the app accessing contacts. What advice do you have, if any, for developers who want to get it right, who are trying to struggle to match consumers' expectations and their preferences with the developers' practices?

KRISTOPHER MICINSKI: Right, well, I think that one thing that we found in our study and have more solidly confirmed in a more recent study we've done is that users compartmentalize even if an app is using one permission like your location, but people conceptualize that usage in a bunch of different ways. So they might conceptualize the use of the location to track your run. But they might not then think about the fact that it's also being used to send some analytics data back to a server or something like that.

And so I think the thing that I would say is that for developers who want to be well-intentioned, they should be clear about the different modalities of this use of this information. And for example, if they plan to use your data to track your run, and it's going to stay on the phone, then they should tell that to users.

But if there's some other data, for example, that's going to track you and send it to an analytic server or something, then people should be conscious of that. And not just take the fact that you click on the permission icon or whatever to just use whatever information you want. And it turns out that even though you might suspect that developers might want to do this, it can be challenging, because a lot of things like analytics and add libraries-- developers will include these in their apps because they just want to be able to have these features. But they don't realize that these libraries then, behind the scenes, they just run all the time. And they have this habit of just checking for all the permissions that the app has access to. And then optimistically when an app has this information, these analytics and add libraries will use it. And developers might not-- even if they're well-intentioned-- might not even understand that that use is occurring.

KRISTEN ANDERSON: Do you have thoughts on how best to implement some of those suggestions? For example, if a company would like to use the geolocation information that they're collecting from a run for other purposes, how best to communicate that to consumers without overburdening them with notices or just bearing in a privacy policy.

KRISTOPHER MICINSKI: I don't have a good answer for that yet. Although I think that as a community, we're converging on one. I think that we have some key principles. For example, if you plan to access data persistently, building UIs that let users know that your data is being accessed over the duration of some period. So for example, if you want to share your location

with someone on Facebook, there's this nice little indicator for example, that says your location will be shared with person x for 60 minutes or something like that.

I think those things are moving in the right direction, but I really don't think we have the answer at this point. Maybe we have some shades of where it's going though.

KRISTEN ANDERSON: OK well, all this research, I think helps to inform that debate. So thanks again for your presentations. Emily, you found that neither parents nor children understood that the toys recorded what they said. And you also found that children have a need for private spaces, which I think you talk more about at length in your paper. What are the implications of those two findings together?

EMILY MCREYNOLDS: So the recommendation we ended up making was to make sure the fact that these are being recording is more evident to the kids and to the parents so that they can make these choices in a more informed manner. So what you end up having is, when you ask these questions, it's also a way of notifying them. So I think it's great to be sitting next to Kris, who's talking about what user expectations are. And how we can do these notifications in a way that people don't just continue clicking, but that they actually see what's happening.

And so that's why I think one of the things that would be helpful is if you have these kinds of toys or devices, some sort of notice that it is recording.

KRISTEN ANDERSON: And we actually got a question from the audience about your sample-- where you got your sample. And if it was from Seattle for example, do you think that people from Seattle are more or less likely to be concerned or informed about privacy and security?

EMILY MCREYNOLDS: That's interesting. One of the things we did-- the way we collected our participants was through what's called snowball sampling. So we did outreach to local parent groups. It was definitely Seattle-based. When you do an in-person user study, that's pretty much the only option for a small study like this. So it was very Seattle-based. It was interesting because there weren't a lot of people who had had experience with toys, or had say, one of these devices-- smart speaker devices-- in the home.

So I know a lot of the privacy folks in Seattle. And we intentionally didn't draw from those communities. So perhaps they were. Some of the things they didn't have time to mention were, we did have two parent who thought this kind of system would be great for monitoring what their kids said, and then being able to talk to them about it later. So it wasn't all people very worried about their privacy. Some people saw some advantages in these toys.

JINGJING REN: Here's a question. Do think that parents actually violated the privacy of the children? I mean, they have access to the recording. Are they allowed to have?

EMILY MCREYNOLDS: Or should they? I think it's a good question. So how parent deal with their children's privacy, I think isn't a new issue. You can find stories for decades of parents reading the kids journals, and kids getting very upset. So an interesting question is how is this different? And that's why I think that recording notification is really important.

And I also think having parents talk about privacy expectations with their kids. I've heard some great discussions. Parents have anecdotally told me how they talk about privacy with their kids. Not just hey, I want you to know I could have access to this information, but also you should know that these kinds of things can be recorded. These kinds of things can be collected. So be aware of that when you're using devices, toys, the internet, that kind of thing.

KRISTEN ANDERSON: And I do want to move on to Pardis, but I want to ask a follow-up question of Kris about some of the findings that you made about the indications of recordings. And Kris, what you think about the interaction with the toy where you're pressing on the belly of the Dino or the belt buckle of the Barbie to be able to activate the voice recognition features. What do you think about that interactivity and what it means for notices?

KRISTOPHER MICINSKI: Yeah, I mean, I think that one thing that I would be-- well, maybe this is a discussion we can have after a little bit later, but I would be interested to see-- the kids think that when they're doing this they're interacting with something, but I suspect that if you're an adult, you click on the button or whatever, and it's a picture of your camera. So obviously it's going to take a picture of you. But for children I just don't think that they would have the same mental model. Do you think that children are cued into the-- is it fair to think that they're going to think about their privacy in that way? If I were a kid, I would just want to have fun or something. And so I wouldn't think about my privacy probably at all. And if even somebody told me that there was something bad happening, I would just say, I don't care. I'm a kid. So what do you think that kids are thinking about that when they for example, push the button? That seems like the important part.

EMILY MCREYNOLDS: Yeah, one of the reasons I care about recording notices is because I think, once the kids know they're being recorded, they have really-- I was really impressed. I did, not a follow up study, but took a different device into a group of kids who regularly participate in studies. And they got to play with these things. And at the end we were like, so what do you think? And they were like, they're great, as long as it's not recording me. And you're like oh, so you do have a sense of what might be happening. And that you don't necessarily want that to happen, because you do care about your sense of privacy. So I would agree that they need those kinds of indicators.

KRISTEN ANDERSON: Great. Thank you. Pardis, so I understand that your work was a first step in understanding what consumers can expect and prefer. What are the next steps for research in this area that you see? How can you build upon what you've found so far?

PARDIS EMAMI-NAEINI: Well, as you just mentioned, this is the first step toward designing this privacy assistant. So the limitation that we had in our study was that participants were answering questions about hypothetical data collection scenarios by sitting behind their computers. So to make this study one step more realistic, we are now designing an experience sampling study. In this study we will ask participants to install our application, and then we will notify them at specific intervals about the hypothetical data collections that are happening in the geolocations that we are interested in.

And after notifying them about these data collections, they'll ask some questions to understand your privacy preferences in that specific situation. For example, when they enter a coffee shop we will ask how they would feel if that coffee shop was using a video camera with a face recognition system to identify its frequent customers.

KRISTEN ANDERSON: OK. And we have a question from the audience for you too. "So when looking at privacy preferences for data collection, did you account for different populations? For example, age of users that may not be as sophisticated as other users?"

PARDIS EMAMI-NAEINI: So we had demographic information as one of the factor in our models-- in all of our models. But we did not find any statistically significant result for that demographic information-- like age or where the are from or different education levels.

KRISTEN ANDERSON: OK. Getting Yang in on the fun here. Yang, you found some common privacy concerns among the crowdworkers. And one of your recommendations was that each task should provide a privacy policy. We often get complaints about people potentially not reading privacy policies, and I'm wondering how you square your recommendation that each task should have a privacy policy? How do you square that with a common refrain that nobody will read it anyway?

YANG WANG: Yeah, thanks for the question. First of all, I think this is probably cliché for this audience. We believe that notice and choice are a very important principle. And I think the real question is, how do you implement them? We're certainly not recommending having a three-page privacy policy for each task full of legal jargons. That's not exactly what we're recommending.

So for Amazon Mechanical Turk, when the requester is composing a task, they have an interface where they will say their name, a brief description of what the task is about. I think one idea is-- it's a pretty simple idea for Amazon to implement-- would be in that form they can have specific fields-- things like who you are, and what data you're going to collect, how are you going to use that data, who are you going to share that with. Even a simple thing like that, I think would provide useful information for the crowdworkers. Even if the crowdworkers don't pay attention to that. In some cases, I think having the requester go through those forms, will help them think about the potential privacy implications. So I think why we think implementing in that way would be useful.

KRISTEN ANDERSON: Those are really interesting ideas. I'm wondering if you've had any interactions with the platforms about implementing those kinds of suggestions?

YANG WANG: Yeah, so interesting you asked this. Just past week I've been talking to Amazon Mechanical Turk, and we're still talking. So I told them some of these recommendations we make. They seem very open to some of these suggestions. Especially the one I just mentioned-- a simple change in their requester interface. Hopefully they will implement that soon. So we're in discussion.

KRISTEN ANDERSON: Great. I do want to bring some more questions in from the audience. So I've got one question for the panel. And this is, "a major problem for both users and developers is notice fatigue. How do you suggest developers go about communicating critical information without overload?" I think we've touched upon this topic before, but I want to open it up to all of you to have a deeper discussion.

YANG WANG: So I guess I'll jump in first. I think as many of the research we just heard suggest that it's important to know what kind of information that people would really be concerned about. Because otherwise, if you provide these sort of things, they're going to get overwhelmed. I mean, the developers are very busy. And oftentimes privacy or security might not be where their primary attention is. So helping them understand, under what scenarios what kind of information would it be very concerning for the users will help them.

EMILY MCREYNOLDS: And I think we're facing a new challenge where-- or at least new over the last few years-- where we have these devices that you're talking to that are talking back to you, which is the research we've been doing. And obviously a 10-page privacy policy doesn't work in that situation. You aren't going to have it read out to you.

But I've heard some really interesting anecdotal ways of dealing with this. So if you ask a privacy expert, they'll say, oh, yeah when I ask one of these devices, "what is your privacy policy," it should tell me the privacy policy. Or it should refer me to that. Fine, great, go to the website. You can find our privacy policy here.

But there's really interesting data on the back end they can use instead. I heard of one company that, what they did was when told, hey, you should have a response to, "what is your privacy policy." And they came back a few months later and they had not done that. And so the person asked, well, why didn't you do that? Well, nobody ever asks that question.

So instead, we figured out what a good answer to, "hey, are you spying on me," would be. Instead of, do you protect my privacy? Are you recording me? And providing privacy notices into those questions in a way that's accessible for the user. We still have the, you set up the device, and you agree to the privacy policy, and all of those things. But those coming just-in-time notifications where they're responding to the questions people are actually asking, I think is a really helpful way to think about it.

KRISTEN ANDERSON: Does that imply that people should be-- that developers and companies should be doing more research behind the scenes to find out-- like the kind of research that you guys are doing-- find out what consumers care about, and in what contexts and work from there?

PARDIS EMAMI-NAEINI: So we've presented in our study we found that people's privacy preferences are very complex. But some factors, or combinations of some of the factors are more important. So I think my recommendation to our IT developers or policymakers is that to be more clear about their data practices. And more specifically, to be more concrete about the factors that we've found to be more important. For example, the type of data being collected, the purpose of data collection, or whether and to whom data is being shared to. These are among the factors that people are more concerned about.

YANG WANG: So if I might add-- I think one thing that the developers should be able to do is, considering user privacy as part of their development process. So it's not an add-on after the fact. So the example I give for Amazon Mechanical Turk, if you embed the privacy information at the time when the requester are making the request, that makes it easier. So imagine you can embed some sort of a privacy feedback to the developers right in their development environment. I think that would make it easy for them to adopt these suggestions.

KRISTEN ANDERSON: Great.

JINGJING REN: Yeah, I think the [INAUDIBLE] interaction is with the app developers, especially the credentials case. They also don't aware of that, because they are included as third party libraries. And they accidentally send out of the credential to these third party. So it's, to their end, it's also their incentive to know the privacy leaks. But I think there's a common challenges. They do not have a good tool to audit to even their own app. So I guess, even for the research committee, and then we should work with developers to build these type of tools. Developers are also users. So we should make it easy for them to understand the privacy implications of the apps they develop.

KRISTEN ANDERSON: We've got a couple of questions. One of them came in from Twitter and another from in the room. And I'm going to combine them into one question, and throw it up to Emily first. And then to everyone else. So for Emily. the question was, "How do you think the early experiences with connected toys will affect children as adults with IoT devices? Do you think they might be more strict with their privacy or less, as a result. Is or anything marketers can learn from this?" And then a slightly broader question is, "Should consumers assume that IoT devices are really services, that they have a limited life expectancy, or have expected ongoing costs to keep current?" so I think it's implications of IoT and privacy overall, and to what extent early experiences with connected toys might inform that.

EMILY MCREYNOLDS: So we've had a lot of conversations about how parents think about these things versus how kids think about these things. And I think there's a lot of discussion going on about say, manners. How are we going to-- if you can command a robot to do exactly what you want it to do, what are our kids learning? And I think that it's all about the discussion you have around these devices. I think you are seeing kids grow up much more connected than we ever have before. We know that, because we're in this era where we all have little computers in our pockets at all times.

But I think it's an education effort, like anything else when a kid is growing up. So they'll be more connected and more in tune with the devices than ever before. But the same conversations about privacy can still happen.

KRISTEN ANDERSON: We have one more question that was directed to everyone. And I think it would be a good lead in to our next session, which is going to be on Economics, Markets, and Experiments. And that is, "Did any of the studies that you guys conducted of consumers perception"-- sorry. "Did any of the studies measure consumer perception of the damages to them as a result of using the apps, devices, or crowdwork platforms?"

KRISTOPHER MICINSKI: No. I mean, I think this is really hard. But I have a very pessimistic take on this, which is just that there aren't economic implications for leaking privacy from IoT devices and apps. And--

KRISTEN ANDERSON: I think there are implications. Do you mean this is hard to measure?

KRISTOPHER MICINSKI: I'm skeptical. I mean, I always see people's stock bump down, but then it bumps right up back after a while. I mean, I think that if we want to act-- I mean, I think there probably should be economic implications. But I just get the feeling that the vast majority of the time, notions like privacy are so fuzzily defined that it's not-- for obvious egregious things, like leaking somebody's social security number-- sure, that's really bad. But I think the problem is for very nuanced things like maybe leaking some of your location data-- sending it to an analytics library. I just don't think that people view that as that big a deal.

And I think there definitely needs to be more understanding about what users are-- how that economically benefits the app developer. And I think there needs to be this recognition that frequently, app developers and IoT device developers-- they're disincentivized from being transparent with users. And in some sense understanding more privacy stuff can really help us in that way.

But I think it's easy, as many of us have found on the panel, to mask genuine, good things like real benefits of technology. For example, being able to use a toy. Being able to use an app. Being able to get money from a crowdsourcing platform. I think it's easy for that economic benefit to unfortunately let some of the privacy leakage-- let you justify it to yourself that you're going to do it without thinking about the broader implications of maybe losing your privacy.

EMILY MCREYNOLDS: So I would break that question down into, what are the implications for the individual, and what are the implications for a company. I think we've seen pretty clearly that if you have a data breach, or you don't meet your user's privacy expectations, that there are economic harms. Whether it's a data breach where you have to spend millions of dollars notifying the people, or reputation harm even if your stock price does or doesn't drop.

And then on an individual level, I think people spend time thinking about what the concerns are. And they have costs involved in-- how do I pursue this? How do I manage my privacy now that my, whatever number, is out in the public? And so I think if we think about it broader in terms of, maybe lost productivity or expenses, there's some really interesting and important economic concerns.

YANG WANG: So I mean, for us the answer is no. In the particular study we did, we didn't really ask that from the economic perspective. But one thing interesting about crowdworking is that there is a direct transaction to be made, so to speak. You do the work. If they accept it, you get paid. So people are-- they are making economic decision in a way. In some cases, they will get paid for less than \$1 and providing very sensitive information. We didn't study that in particular, but I would imagine researchers could run experiments on these platforms, and figure out how people are actually making these decisions that may have economic impact.

KRISTEN ANDERSON: If nobody else has anything that they want to add, I think that would actually be a great lead in to the next session, which will as I said, be on the Economics, Markets, and Experiments. And that will start back up at 1:40. So we'll be breaking for lunch now. But as a reminder, during the lunch hour, just across the halls in the conference rooms, we're hosting a student poster session. We have a great variety of presentations there. And so I highly encourage all of you to go over and check out those posters. And also interact with the researchers.

In addition, there are a number of government agencies that provide funding for privacy research, who'll be seated at lunch tables that we've also set up in the same conference rooms. You'll be able to identify them by their table tents. And we encourage you to eat and chat with them, and learn about their privacy interests, needs, and funding opportunities.

More than half the presenters at this year's Privacy Con are discussing research supported by the agencies that are going to be joining us for lunch. We're very grateful that they support the outstanding research on important and interesting topics. So please go ahead and grab some food, and take this opportunity to check out the posters and meet some of the government funders. And then we will see you all back in here for session 3, starting at 1:40 PM. Thank you very much.

[APPLAUSE]

[MUSIC PLAYING]