FTC PrivacyCon
January 14, 2016
Segment 2
Transcript

KRISTEN ANDERSON: OK everyone, please take your seats. We're about to get started with the next session.

Good morning. I'm Kristen Anderson, and I'm an attorney in the division of Privacy and Identity Protection within the FTC's Bureau of Consumer Protection. I'm here to introduce our second session of the day, which is on consumers' privacy expectations. We'll hear from six researchers in four 15-minute presentations, and then we'll conclude with about 20 minutes of discussion where we'll identify common themes and ask the presenters about their work and its implications.

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

SERGE EGELMAN: Thank you for that introduction. So this is work that I've been doing with several students recently, where we've been looking at privacy and how private information is regulated on mobile platforms.

So to give you, I guess, a brief overview, most of this work is on Android, and that's only because Android actually has a pretty intricate permission system to try and implement notice and choice. So whenever an application requests access to certain sensitive data, it's regulated by this permission system. And so when users install an application, they see a screen that informs them of all the possible types of sensitive data that that application might be requesting in the future.

And so the question was, does this actually implement effective notice and choice? So do users understand these messages about how applications could be using their data in the future? So we started this project a couple years ago by doing an online survey. We had over 300 Android users, and we just showed them screenshots of these permission screens and simply asked them if an application was granted these abilities, what might that allow the application to do? We then followed that up with a qualitative study where we had 24 people come to our laboratory, and we interviewed them about similar concepts.

And what we concluded from this was that many people were simply habituated. Since these appear every time people install applications, not only does it list what abilities and types of sensitive data that application is requesting in the future, but all the possible types that it could request, even if the application never takes advantage of that. And so people become habituated. They see lots of these requests that have lots of different data types, some of which they don't understand, and therefore they learn to ignore these because there's just so much information there.

Another problem was that people were simply unaware. Since this occurs whenever you install an application, a lot of people said that, oh, this is just part of the license agreement, and we know that we need to click through that in order to continue installing the application. So maybe this occurs at the wrong time in the process. And since it happens after the user clicks install, it could be that they're already committed to installing the application. There are various cognitive biases that relate to this. And so therefore, it's unlikely that they're actually comparison shopping based on privacy, even if they wanted to.

Another issue is that understanding whether a particular application is going to access a particular type of data really requires a good understanding of this whole permission system and what are the different types of data that are regulated by the permission systems. So understanding whether an application is requesting a data type requires understanding the whole universe of data types that are governed here.

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

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

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

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

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

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

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

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

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

So these screenshots were taken randomly whenever one of these sensitive data types was accessed so that we can ask them as a prompt, you were doing something. This is what you were viewing on the screen of your phone. It was requesting this particular type of data. Was that expected? Did you expect that application to be requesting that particular data type at this moment in time? And also, if you were given the ability to, would you have prevented that from happening?

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

So some of the problems that we found were due to incorrect mental models. So again, the goal of this is transparency. Show the user all the possible ways an application might be accessing sensitive data. Is that really working? Well, we found that in 75% of cases the application that was requesting one of these data types was completely invisible to the user. So this was mainly due to the screen being off in 60% of the cases-- so applications running, the user wasn't actually using their phone, or background services.

Another thing that we found was despite the fact that there are some privacy indicators built into the operating system-- so both Android and iOS have indicators for when GPS is accessed. This is an example of one of those indicators. It appears in the top status bar. And most people assume that the only time that GPS information is collected, this icon will appear.

And it turns out that's not true at all, and in fact the icon only appears in 0.04% of the cases where location data was accessed. And that's because every time an application requests location data, the operating system caches that for performance reasons, and also to preserve battery life. But then when another application accesses just the cached location data as opposed to querying the GPS hardware directly, this icon never appears. Similarly, applications can infer location based on cellular network data, nearby Wi-Fi hotspots. And it turns out, most applications are using those methods to infer location rather than the GPS hardware. And therefore, most of the time when location data is collected, people have no indication that that's occuring.

So having the notice and choice at the beginning when users install the application obviously doesn't work. We've tested that. The ask on first use that's currently happening isn't really working because of the different contexts in which users might be interacting with applications. So maybe we could have runtime requests all the time.

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

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

So how do we do this? How do we give users control over the things that they actually care about without overwhelming them? So we're doing some work now to try and predict the cases where applications access data where people would want to note this is occurring, whereas the other ones where applications access data that might be expected, well, obviously, we shouldn't prompt the user in those cases.

And so what we found was that expectations really did predict behavior in this case. So when we asked people if this access to personal data was expected or not and then whether they would have blocked it, there was a pretty strong correlation there. We also found that using the current model on ask on first use-- so if you look at, for each unique application and each unique data type, if you ask users the first time that application requests that data, we're going to get it right about 50% of the time, which is what's currently happening. So that's a coin flip.

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

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

So going forward, we're actually trying to implement these systems right now that can try and predict a given user's preferences based on their previous behaviors. And this is part of a pretty complex ecosystem. So we have what we're calling hard policy, which is preferences that people have explicitly stated-- so I don't want applications to be using data for x reason-- and then trying to augment that with soft policies-- so inferred preferences that systems can make about users, such as maybe looking at hundreds of thousands or millions of users, we can infer one user's preferences based on other users who are like them, like recommender systems. And also, based on the feedback from prompts-- so if we can design more efficient prompts that cater to individual user expectations, we can then use the output of those-- so what did the user actually decide-- to ensure that they see fewer prompts in the future.

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

[APPLAUSE]

KRISTEN ANDERSON: Thank you, Serge. Next we'll hear from us Ashwini Rao of Carnegie Mellon University about mismatched privacy expectations online.

ASHWINI RAO: Thank you. So yeah, my talk is about expecting the unexpected, understanding mismatched privacy expectations online. So I'll start with a motivation. So many of us on a daily basis interact with online websites. And as we interact with online websites, we may have questions such as, what types of data does this website collect about me? How does it share this data? And does it allow deletion of this data?

And to answer these questions, a user could read the website's privacy policy, which is usually a textual document in English and it discloses the data practices of the website, such as collection, shading, and deletion. However, these policies in their current form are long and difficult to read. So users usually know that.

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

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

Now, the question here is, could we generate effective privacy notices by extracting and highlighting these data practices that do not match user expectations? So the concept is simple-- a privacy notice does not have to inform you about things that you already expect or know. A privacy notice has to inform you about things that you do not expect or do not know.

So I want to make a distinction between policy and notice. A policy is usually a textual document. But a notice, which is based on the policy, is usually shorter and more usable. So here I'm showing you the privacy nutrition label, which focuses on visual formats. And so far, notices that are more effective research, our research has focused on visual formats. And our approach of extracting and highlighting mismatched expectations is complimentary to this approach. Once we identify and extract these mismatched expectations, we could present them to the user in any visual format that is effective.

I also want to say here that these privacy notices do not have to be generated or provided by the website operators themselves. These could be provided by a third party, for example through a browser plug-in. And this is something important to note.

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

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

To measure expectations, we conducted user studies. So one of the user studies that we conducted focused on the expectation in the likelihood sense. And in future, we also plan to mention expectation in the desired sense. So we presented users with different types of websites, and after the users interacted with these websites we asked them, what do you assume that the website will do? And once we elicited user expectations, we next extracted the data practices from privacy policies and then we compared these two to identify mismatches.

So in our study, we varied the website characteristics and user characteristics. So as I mentioned earlier, user expectations can vary based on these websites and user characteristics. We looked at 17 different data practices, which were split among collection, sharing, and deletion. And for collection and sharing we looked at four different types of data-- contact information, financial, health, and current location information.

So here's an example scenario. So here, the scenario is describing the collection of different types of data when the user does not have an account on the website. So you can see that we are asking the user, what is the likelihood that this website will collect your contact information? So in future, if you wanted to also measure desired expectations, we could also ask them, do you think the website should be or should not be allowed to collect this information, in addition to do you think it's likely that the website would or would not collect this information?

So we deployed this study as an online survey, and we had [INAUDIBLE] in total 16 websites. We had 240 participants that were recruited from Mechanical Turk crowdsourcing platform. So this was to elicit user expectations. The other part is to actually extract data practices from privacy policies. And to do this, we used two annotators, two experts, one in privacy domain and other the legal domain. And they manually read these policies and answered questions such as, does this policy disclose that the website collects health information?

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

Now, it's important to note that there can be different types of mismatches. Here I'm showing you two, the yes/no mismatch and a no-yes mismatch. And this is important, because the type of mismatch can impact users' privacy differently. So consider the yes/no mismatch. The website says that yes, we collect your information, but the user thinks, no, the website is not collecting my information.

So in this case, the user may go ahead and actually use the website and unknowingly give up data and lose privacy, whereas in the no-yes mismatch the website is saying, no, we do not collect information but the user thinks incorrectly that indeed, the website is collecting my information. So in this case, the user may decide not to use the website, in which case the user may lose utility but not privacy.

So some results-- so we had looked at different types of website characteristics, and we found that only website type had a statistically significant impact. And the type impacted users' expectations only for financial and health information, but not for contact or current location information. Several user characteristics also had a significant impact on what users expected. So for example, user's age impacted whether they expect websites to allow deletion of data.

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

Now, compare this with a no-yes mismatch. And this is a mismatch in sharing data practice. Users expect that websites will share their contact information for any purpose. However, websites do not do so. They only share contact information for specified and very narrow purposes. So as regards to deletion, users predominantly expected websites to allow deletion of their collected data. But websites generally do not allow that.

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

So based on the results of our study, we could come up with notices that have less amount of information than a full notice. For example, we looked at 17 data practices. A notice could show information about all 17 data practices, or we could show information about data practices where there's a mismatch between what users expect and what websites do, or actual data practices of websites. So for example here, for the Bank of America privacy notice, there were mismatches for 11 data practices out of the 17. So if you show only 11, that would be about 35% reduction in the amount of information that the user has to read and process.

We could also just show information about mismatches that are more privacy-invasive from a user standpoint. For example, I talked about the yes/no mismatch versus the no-yes. If you find that the yes/no mismatch is more invasive, we could only show information about those mismatches.

And in the case of Bank of America, it's only five data practices for which there's a yes/no mismatch. So that would be 70% percent reduction in the amount of information shown in the notice. However, the caveat here is that we do have to go ahead and test with users how effective the shorter notices would be. Yeah.

So as part of future work, we are planning to also study expectations in the desired sense and compare that with expectations in the likelihood sense, and also compare both of them to actually data practices of websites. We will also, as I mentioned, test effectiveness of notices that highlight mismatched expectations and see whether they actually reduce user burden and whether users can make better privacy decisions.

Yeah, that was all. Thank you.

[APPLAUSE]

KRISTEN ANDERSON: Thank you, Ashwini. Next we'll here from co-presenters Heather Shoenberger of the University of Oregon and Jasmine McNealy of the University of Florida. They'll be presenting on reasonable consumer standards in the digital context.

JASMINE MCNEALY: So good morning, and thank you for having us. Our project is Offline versus Online-- Reexamining the Reasonable Consumer Standard in the Digital Context. The impetus for this project is really trying to get a deeper understanding of how consumers act when online. So we know from prior literature that people, individuals, act differently supposedly offline than they do online. So we wanted to take this into a further exploration of consumers.

And we know that the reasonableness standard is a standard that's used for regulators, for example, in assessing complaints related to deception. So we wanted to find out more and explore this a bit more, so we came up with an umbrella project that used mixed methods to examine this question. One of the first things we did was start to interview. So we did qualitative interviews.

And just skip forward a little bit-- we asked our interviewees questions related to how they behave, both online and offline. And we have this quote from an interviewee where we asked questions related to their expectations related to privacy or how their information would be used, and how they attempt to control their information.

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

But more importantly than just showing photos, we asked a question about whether or not they would sign a printed contract without actually reading the terms of a contract versus whether or not they always click Yes or No to the terms and conditions of using various websites, whether it's social media or shopping or whatever the case may be. So we wanted to find out. We noted that there were some significant differences indicated with respect to their sharing behaviors, both online and offline.

So to go back a little more about our method-- so again, we used qualitative and quantitative methods. So just some breakdowns for our interviews-- we had 30 participants. We did these long-form, qualitative interviews and we're going to do more long-form, qualitative interviews as well. So we had 20 women, 10 men. We noted the average age was around 26, and then we have some racial demographic data broken down as well.

Then for our quantitative side, we did a survey. And we're going to talk a bit more about the results of the survey today. And there were 871 participants, almost equal break down between men and women. But note the age. So we had an age of 35.9-- so almost a 10-year age difference on the survey, the qualitative side-- and again, the breakdown of racial demographics.

Also important are some of the variables that we attempted to measure in our survey. These merit variables we got from prior literature, but they also emerged again when we were doing our qualitative interviews. And one of those important ones was social trust, and social trust was measured on a six-item scale. And social trust is really asking the participants how they felt about whether or not they trusted that the institutions-- the entities, brands or advertisers, the government, news media also-- would fulfill their responsibilities related to the consumer's private information. And so those are example questions on here, as well.

And then the second important variable we attempted to measure on a four-item scale was control, or how participants perceived they had control over their information. So an example question was, I can use online privacy tools to remain anonymous online. Perhaps more importantly are our main dependent variables. So we had the always click Yes. So again, we're assessing behavior, whether or not the participant always chose to click Yes related to privacy policies or terms and conditions online.

And the second one was privacy concern measured on a three-item scale about whether they thought that data companies would collect information about them that would make them feel uncomfortable. And Heather's going to come and talk about some of the relationships we found.

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

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

Then we move to the second block, and these were two variables that did come up in our survey that have also been used in numerous studies before ours. And social trust in this particular case was not a predictor, but control efficacy was. So even though they may not actually be able to control their data, the belief that they can predicted always clicking Yes. And we believe this is the result of the confidence that people have if they believe they have control, and as a result they go ahead and say, sure enough. I'm just go ahead and click Yes, but I'm confident and I trust that this is going to work out for me.

Those who had a negative-- oh, so the next block were all items that were derived from our interviews. Of course, some of them you've seen in previous studies as well, but all of them were derived from our interviews. So negative experience-- those who had had fewer negative experiences, self-explanatory, were more likely to click Yes without reading any terms of agreement, no further investigation.

Peer recommendations-- we were really hopeful that a peer recommendation would kind of be an if-then rule. If a peer recommends Snapchat to me, I will then go ahead and download it. That was not the case in our regression analysis. It wasn't significant.

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

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

And then just simple presence of a privacy policy or an icon like Trustee also predicted clicking Yes. So this was our behavior. And at the conclusion of this we thought, we're on the right track here. These cues are what is driving the motivators of actual behavior online, and we were really excited.

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

So the [INAUDIBLE] regression is in the exact same format. Higher ages and higher education-- again, no surprise-- predict privacy concerns. Lower social trust, the trust of the institutions, predicted privacy concerns, lower control efficacy, both in line with previous research. People who had suffered more negative experiences were more likely to say that they had higher privacy concern. Again, peer recommendation-- we'd had high hopes for that, but it didn't work out.

Convenience fell out of this model as a result of the two cues at the bottom. And they're definitely within the same direction as before. If the site had poor aesthetics and it was ugly and weird and it made people feel more skeptical, predicted privacy concern, and a lack of a privacy policy or a link or an icon, predicted privacy concern. And note that both of those two cues predicted both the concern and the behavior.

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

If we make really clever use of the cues-- and there may be more than the ones that we explored-- both entities who collect data-- so businesses, advertisers, the government, news media-- who use data can reduce privacy concern, which is something that they would like to do; encourage the free flow of data, another something that they would like to do and something that last night was mentioned that the Federal Trade Commission potentially may be interested in doing also; and increase trust. And on the flip side of that, consumers could rely on cues that are more uniform and meaningful, even if they don't read the privacy policies that underlie those particular cues.

So with that, I really have to move through this very quickly. So there's really a three-prong approach, and we've already begun the research to sort of decide whether or not this is the right approach. But we would like to suggest guidelines for different types of data collection and use-- something that has been echoed on the panel already-- based on the average consumer's expectation of privacy-- so there is some additional research to do-- and then delineate those types of data collection and assign a cue or heuristic to each type of data that would be endorsed by the FTC.

Here's the catch for people who are in the advertising industry who are in the business of collecting and using consumer data. They would have to adhere to those guidelines in order to use the cue on their sites, which would signify safety, increase trust, hopefully, et cetera.

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

So in conclusion, we are continuing to pinpoint consumer expectations of privacy and a way to develop these guidelines, and the resulting cues that would align with the guidelines. As Jasmine mentioned, we're continuing to collect data, both in the interview portion of the study and also in the survey, just to make sure that we have as close to a census as possible, because we are dealing with the average consumer in the United States and we want to make sure that we get everybody. Examine some additional actual variables as they arise, because while the cues that we mentioned are really good predictors, they may not be the only ones; design policies for readability understanding for consumers so that they have the opportunity to make meaningful choices if they do in fact read those.

And finally, something that we think is really important-- and I will diverge for just a second, literally just a second-- in Australia there was a really great PSA to help people avoid being hit by trains. It's call "The Dumber Ways to Die," I think. And it's gone viral and has actually resulted in lower train deaths. And it's really a silly video. You can look it up on YouTube. There's these little animal-type things dancing around and talking about dumb ways to die, and don't get hit by a train.

And essentially, we're looking to do a PSA like that, based on research in America-- I mean, that worked in Australia; it may not work here-- to allow both the consumers to understand what these icons mean, how they can use them as a way of increasing trust, and also to entice entities to go ahead and opt in to this system and adopt the guidelines the FTC has put forward in a way to align with consumer expectations. And with that, we conclude.

[APPLAUSE]

KRISTEN ANDERSON: Thank you very much, Heather and Jasmine. Our final presentation is by co-presenters Andelka Phillips of the University of Oxford and Jan Charbonneau of the University of Tasmania. I think you two went for the longest commute today, but we're glad you made the trip. Here's your clicker. And Andelka and Jan will be presenting their work on privacy in the direct-to-consumer genetic testing space.

JAN CHARBONNEAU: OK. Well first off, I'd like to thank the FTC for the opportunity to discuss our research. We're going to talk about privacy of a specific type of data, that being genetic data, the data that results from genetic testing-- so a very specific type of data. What we have to realize is genetic data is the most personal data there is out there. Not only is it a unique identifier of us individually, but because of the familiar nature of DNA it can also identify our family. So when we're talking about privacy in this context, we're talking about it in a much broader context-- not just personal, but looking at the family.

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

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

Does it matter if this happens in a direct-to-consumer genetic testing situation? Well, the first thing we have to realize is the difference between traditional genetic testing and what happens when we have genetic testing in a direct consumer setting. Traditionally, genetic testing has happened within a country's health care system.

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

When we look at direct-to-consumer genetic testing, we have to realize that at its core this is a commercial transaction that occurs in each country's marketplace-- and increasingly, in market space, because the majority of the activity is actually online. When an individual engages with DTC, they engage as a consumer. What that means is that enlivens each country's consumer protection legislation. It also enlivens some particular legal protections in contract negligence, et cetera. OK? But a very, very different situation.

What does the general public think of when they think of privacy? At the Center for Law and Genetics at the University of Tasmania, we've been looking at genetic privacy issues for the last 20 years, and in the last few years we've moved into DTC. Some of our early research in direct-to-consumer genetic testing suggested, from the Australian general public's perspective, that privacy concerns were going to be the key constraint on commercial uptake. Interestingly, this past year we found the same results when it comes to intention to biobank-- in other words, giving a genetic sample into a genetic database for nonprofit, institutional, and health-related research as opposed to commercial.

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

And that informed the research that I'm going to talk about today, which is an online panel of 3,000 respondents of 1,000 American, 1,000 Australian, and 1,000 UK respondents. We've just added 1,000 Japanese respondents, which will give us some interesting contrast. The way the sample broke down, about 10% of the people are actual consumers. And that equates to those early adopter categories. That leaves about 90% of my respondents who are the potential consumers. So we're able to look at actual versus potential consumers.

So what does privacy mean from the general public's perspective? Well, simply stated, if something's private, it's not shared. If it's shared, it's not private. OK? In a simple way, and that's how the general public looks at these things. Privacy issues arise from sharing. So privacy is all about control over sharing. Providing your permission to share means that you have control over your privacy.

So that's the way the general public looks at it. If my permission is asked, then I know what's being asked for. I have the opportunity to ask questions, but I also have the opportunity to say no, and that my "no" will be respected. So I have control over my privacy if my permission is sought.

So what do consumers think about whether or not their permission will be sought? In other words, as the previous presenters alluded to, this area of perceived control. Well, interestingly, the American respondents, 47%, thought they had perceived control. And what's interesting is, on any dimension that I analyzed on, Americans are statistically different to the other consumer groups. For the UK, it's 43%. For Australians, it's 40%. And for Japanese, it's 36%. So that's quite a difference in terms of whether or not people think their permission is going to be asked. Are they in perceived control?

If they are in perceived control, what does that mean? Well, they're more likely to purchase the DTC tests. They're more likely to participate in DTC research. And that's important, because that's permission-based. Right? They ask their permission. But do they actually realize that what they're doing is getting nonspecific, enduring consent?

They're also more likely to share broadly. They'll share with family, not friends, so there's some control. They'll share with their doctors. And that's important, because DTC companies very clearly state their results are for recreation, education, or information only. They are not a diagnosis. But as Graham [INAUDIBLE] said, it would be a very brave GP who would action a DTC test. If they go to their doctors, they're back into the traditional system. They're are also more likely to engage with online sharing communities.

But does perceived control equate to actual control? These are commercial transactions governed by contracts and privacy policies. We did some research in Australia looking at the privacy policies of the DTC companies operating there. Do they comply with our legislation? The short answer? No, they do not.

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

ANDELKA PHILLIPS: Well, I've actually been looking at the contracts and privacy policies of direct-to-consumer tests or companies that offer tests for health purposes. Now, as has been noted in the previous session and also in the previous group's work, these contracts and privacy policies appear everywhere online. Basically, any website you use, any software update you make, will be subject to terms and conditions. And they'll be presented either as terms and conditions, terms of use, terms of service, privacy statements, privacy policies, and sometimes in this context they're combined in one document. At present, these are used to gather not just the purchase of DNA tests, but also using the website and sometimes participation in any research the company is doing.

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

And because of the ubiquity of these contracts, consumers often also display an attentional blindness online. So they may just not notice these. We may not read them. We just click, "I agree." And this is really problematic in this context, and I think there really needs to be reformed. Because unlike some of the other what we've said, that consumers don't read these, I've had to read 71 contracts, and I really think there are problems here.

[LAUGHTER]

So the major privacy risks in this context arise chiefly from sharing or sale of sequenced DNA, but also from sharing or sale of other types of personal data, often health data or other data that we might normally consider to be sensitive. This is because companies are often engaging in ongoing health research, so they're collecting large amounts of personal data from consumers. There's also the risk of possible discrimination based on a person's genetic makeup.

And then there are some other risks that arise. Some of these are more future risks. So there's the possibility with the increasing use of biometrics that in the future, these genetic databases could be used for identity theft, targeted marketing-- the most obvious example at the moment is targeted marketing of drugs to particular population groups or even family groups. Also, there's the potential for discrimination in employment or insurance if this data is shared inappropriately. And more remotely, there's the risk of creating synthetic DNA.

Now, as I previously noted, these contracts are not industry specific, so often you'll encounter the same terms in these contracts that you would when you are purchasing a product or downloading a song online. And they also use really similar wording. Now, in the United Kingdom and the European Union, there is strong consumer protection legislation that deems some terms in consumer contracts to be unfair or unenforceable. And at present, some of these terms would likely be deemed unfair and unenforceable.

And this is interesting, because I know I'm at the Federal Trade Commission's conference but I've been looking at mainly American companies. This is an American industry, overwhelmingly. But these tests are sold internationally, and people's samples are being sent across borders. And so there is a need for international collaboration to protect consumers in this context.

So one of the most concerning things here is that consent will often be deemed through use or viewing of a website. And often, consent to altered terms will also be deemed through continuing to use the website. Now, as most of you are aware, it's often easy to use a website without ever looking at the terms and conditions.

So this is quite concerning, because is the other thing that's very common, and the majority of companies will include this. And 39% of companies include a clause that allows them to change the terms at any time. And only a very small percentage, around 6%, will actually-- or it might actually be 4%, sorry-- yeah, 6% will notify a person directly by email of changes. So most of the time, companies can change the terms at any time or from time to time without direct notice to the consumer. And this is important here, because it could have an impact on, well, what companies do with your data. They could change their policies on sharing, sale, or storage of data. And this can significantly impact consumers.

As Jan mentioned previously too, because this is marketed as a consumer service, companies are often including clauses that say their services are only for research, informational, or sometimes even recreational purposes. Now, in the context of health testing, I would question whether anyone orders a breast cancer risk test for recreational purposes. And stepping on, quite a few of them will also share data with law enforcement, which consumers may not be aware of. And there's often very broad sharing with potential third parties that might include affiliates.

And yes, I'm running out of time. But I really do think there's a need to improve these contracts. And following on from the previous two discussions work, I really think these contracts need to be written in a more easily understood way that would enable consumers to make informed decisions. So thank you very much.

[APPLAUSE]

KRISTEN ANDERSON: Thank you, Andelka and Jan. OK. So now it's time for our discussion session. We'll be spending about 20 minutes, which I'll be leading with my co-discussants, Alan McQuinn of the Information Technology Innovation Foundation, and Darren Stevenson of the University of Michigan and Stanford Law School. So I'll start us off. We're each going to provide some brief comments about what we've heard, and then we'll ask the presenters about their work and its implications.

So first, to me, it seems like you are all striving to answer some of the same basic questions. So what do consumers think about privacy, and why? And those include things like what are their expectations and hopes about the kind of data that's going to be collected and how it's going to be used, how much control that they have, what affects their understanding, what affects their willingness to trade privacy consciously or subconsciously, or unconsciously, for some benefit? And is that contextual? Does it vary by the trust of the firm or online effects?

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

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

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

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

Now, one of the biggest benefits that I see of Privacy Con is that it brings all of you together, the best and the brightest all working to understand the same issues, and providing us with the benefits of your learning. So we're hoping that this conference is going to facilitate you learning from and building upon each other's work. And I hope that we can continue to benefit from the insights that you've given us about how best to protect consumers' privacy, and industry can hopefully do the same.

As Chairwoman Ramirez said, it's now more than ever that we need to stay up-to-date with the latest findings on privacy and data security research in order to fulfill our mandate to protect consumers. And your efforts deepen our understanding and spur our own research in that respect. So thank you again for coming, for sharing your work and your thoughts. And with that, I'm going to turn it over to my co-discussants for their thoughts and allow them to ask the first question.

ALAN MCQUINN: Excellent. Thank you to the FTC for letting me come here today. Is it on?

KRISTEN ANDERSON: Yeah.

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

As we walk into this, there's definitely a lot of discussions over different expectations versus privacy, or people not understanding the legalese in direct-to-consumer genetic contracts. But is that a public policy problem? I'm not so sure. Let me draw an analogy. Say I'm not necessarily sure what goes into my Chipotle burrito. Sure, I'm able to pick different fillings--

SERGE EGELMAN: E. coli?

ALAN MCQUINN: --I may be able to pick different fillings, but I'm not so sure how they're sourced. So when you ask me questions about what's in my Chipotle burrito, my expectations may differ from the reality of what's in there.

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

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

[LAUGHTER]

So at the risk of stating the obvious, I think what we have here is we have evidence, empirical studies, that show that consumers have expectations. All of you in this room, you guys are not ordinary consumers because you're here at Privacy Con. But ordinary consumers, we're seeing that there are consistent, measurable expectations. I really enjoyed the studies, and I encourage you all to read them if you have not read the papers.

I think most of these papers have supported this notion of contextual integrity that's popularized by Nissenbaum and others, with the idea that preheld expectations are measurable and can be demonstrated. Two complications come to mind. So the first is the difference between expectations and preferences. It was clear in Ashwini and colleagues papers, they were really careful to define what is an expectation. What are we actually studying here?

And then to contrast that with consumers' preferences, expectations being different than preferences, which we saw in [INAUDIBLE] work with his colleague, that consumers might just be resigned so that where expectations and preferences diverge, I think this is a really fruitful area of study. So what are we measuring when we're measuring consumers' expectations? Is it what they are just resigned to give up, or is it what they would prefer? And then the papers, a few of them, kind of went back and forth on that.

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

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

So I look forward to our discussion here, and I think we can open up to questions?

KRISTEN ANDERSON: Yeah. Go ahead.

DARREN STEVENSON: Or if you have any responses to our comments.

JASMINE MCNEALY: I don't know. I like the burrito analogy. But at the same time, if Chipotle has lean steak or whatever they have, I mean, if they make representations to the consumer that it's from a certain source, then you have expectations that, hey, my beef is from a certain source. And even if we don't know exactly where it's from, we have an expectation that we should get at least a product of some, I guess, quality. Or at least, we expect that regulators would enforce the restaurant giving us a product that either won't make us sick or won't have had something done to it by a worker there, right?

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

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

So the expectation is that when the bus or the whatever sees that person in the walk, they're going to stop. Does it happen all the time? No, but I think from a public policy perspective it's putting in its proactive measures to protect people from themselves and other people at times. So I think, from the perspective of a regulatory agency that is a consumer protection agency, we'll want to do something proactively when there are signs of issues or trouble. I think it's perhaps incumbent upon a consumer protection agency to do that.

SERGE EGELMAN: So I guess the issue of expectations versus preferences-- so we've done some studies, and we have actual data to show that to some extent, this is an issue of learned helplessness. So people are just sort of resigned to the fact that all of our data is out there, regardless of whether that is actually the case. So for instance, we did a study looking at single sign-on on websites. So when you click the Use Your Facebook login to log in to this website, those sites can then request some data from your Facebook profile.

And so we wanted to see whether making that more apparent to users-- so trying to highlight what types of data might be collected by those websites from your Facebook profile-- we expected that that would have an effect on whether people use this. And we found that was not the case. And when interviewing subjects, they said, oh, well, they just assumed that Facebook is giving away all this data anyway, so I might as well get a benefit from it.

And so that's sort of the learned helplessness issue. I'm not sure there's-- I think addressing that part of it is sort of putting the cart before the horse, because I think one of the issues we need to focus on are the expectations before they're formed. Some of that might be doing a better job of public education with regard to online privacy. Other pieces might come in the form of enforcement, making that somewhat more subjective.

So yes, the law moves very slowly. Technology moves quickly. But I don't think the issue is making the policies around specific technologies. The issue here is narrowing, or closing, the information asymmetries. So while we don't expect people to read every privacy policy that they encounter, we have some expectations about what a business might be doing, as was pointed out. So regardless of what they say about what farm the beef came from, I don't expect it to have E. coli in it. And that's not something that they need to explicitly provide notice for. It just should be expected that there's no E. coli in this beef. I'll leave it that.

[LAUGHTER]

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

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

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

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

And I think people will often tend to engage with these-- and I think your work shows that-- much more differently than they would with a paper contract. So if a browse [INAUDIBLE], which is where the terms are on a hyperlink, it's akin to walking into a shop and being bound by a sign on the wall that you didn't see and walking out again. And that's really problematic. Thank you.

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

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

SERGE EGELMAN: Yeah. I think that was my slide. I think one of the bigger problems with notice and choice is that there, I guess, is unreasonableness on both sides. So there's unreasonable expectations on what the consumer should know to make an adequate choice based on the notice given to them. So it's unreasonable to expect every consumer to read every privacy policy that they encounter.

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

HEATHER SHOENBERGER: I'll also answer that very briefly. As far as using heuristics are part of the cognitive miser sort of mental model, first of all, I disagree that heuristics and using them are faulty. They're almost always correct. I mean, we rely on them all day long in various capacities. I think what we were arguing for were heuristics that were actually backed by informed and concise and true information that the FTC approves.

And so by promoting consumers and allowing them to see what these heuristics mean, promoting the cues to those consumers, it gives them a meaningful choice. So the heuristic is no longer something that risk is as much of an issue for, more as something that they can genuinely rely on as an indicator of safety.

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

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

So we came at it-- and this is an arguable point-- from the type of data collection and how it's being used, and then entities could opt in depending on how they were collecting and using that particular type of data. So there would almost be a continuum of badges or icons or cues that you could use. And then in order to use that on your site or within your materials, you have to adopt the FTC's guidelines that went with that particular icon. And we would empirically test every single element of that.

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

KRISTEN ANDERSON: We have about 20 seconds left, so I'd like to give you guys an opportunity to ask a last question.

ALAN MCQUINN: So to follow up on what Darren said with how privacy concerns kind of have morphed and changed over time, the ITIF actually released report called the "Privacy Panic Cycle" that kind of tracks this. But I wanted to see-- so several different industries up here that you've studied have changed over time. Some of them are new, like genetic testing, and Android is on its sixth release. I was just wondering if you could talk about if you've seen expectations change over time.

ANDELKA PHILLIPS: Me?

[LAUGHTER]

JAN CHARBONNEAU: I think one of the things we have to acknowledge is that we're moving into the commercialization of health, and we're moving into the monetization of health data. And so what we've found is what's existing at the moment is not industry-specific.

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

SERGE EGELMAN: So one thing that we've looked at is trying to relatively weigh different user concerns based on the technologies. And so I guess going to this issue of what policy is needed-- and policy moves slowly and technology moves fast-- while people do have very nuanced privacy preferences and expectations, at the same time there are some things that people will think of as universally bad or universally unconcerning.

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

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

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

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

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

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

[APPLAUSE]

[MUSIC PLAYING]