FTC Putting Disclosures to the Test Workshop
September 15, 2016
Segment 4
Transcript

JOE CALANDRINO: Hello, there. I just want to give everybody a moment to get back to their seats. So thank you all for joining us for a final panel of the day on the future of disclosures. I'm Joe Calandrino, and I'm the Research Director with our Office of Technology Research and Investigation here.

And I'm really excited about our speakers for this panel. They're Serge Egelman, who's with the University of California Berkeley and the International Computer Science Institute. We have Tamar Krishnamurti, from the Department of Engineering and Public Policy at Carnegie Mellon. And we have Florian Schaub from the School of Information at the University of Michigan. So rather than cutting into their time, I just will let it move on to them.

SERGE EGLEMAN: Thank you. I'm Serge Egleman. I direct the Usable Security and Privacy Group at the International Computer Science Institute. Broadly speaking, we look at human factors in privacy and security. So how people use interfaces and make decisions about privacy and security online.

So a lot of recent work that I've been doing has been looking at privacy on mobile devices. So when you have third-party apps running on mobile platforms, they have access to a lot of sensitive user data. And so we've been looking at whether users are aware of how apps are using their data, and then giving them better mechanisms so that they can more easily control how data flows out of their devices.

So one recent study that we did about a year ago is just looking at how frequently data was accessed by various third-party applications. And so to do this, we modified the Android operating system so that users could run whatever third-party apps they normally have on their devices, but the operating system would then keep a log of whenever certain sensitive data types were accessed. And then we can look at that afterwards.

And then we also collected a lot of contextual data, too. So what are the things they were doing on the device when sensitive data was accessed. And the purpose of this was so that we could try and make inferences about whether they were likely to understand when data was leaving their device.

So what we found was that, in terms of these requests that were occurring while running applications, there were about 200 requests per hour. A lot of this came down to location. So about 11,000 times a day people were having location data leaving very device. But there are several other different data types.

And so the reason for looking at how frequently this happens during run time is it used to be under Android that there were these install time prompts. So whenever you install an application, the application lists all of the abilities that it might be granted in the future. This isn't necessarily

to say that the application would always make use of the data that it has access to, but if it were to access certain data, such as location, your contacts, your photo library, that would be disclosed at install time.

And we did a study about five years ago where we found that by and large most people didn't really notice these screens. During interviews they said, oh, that's just the license agreement. I need to click through that to go ahead and install the application. And so to their credit, Google changed how Android does this. And so in addition to these install time warnings, they're runtime prompts, but users only see these the first time an application requests a particular data type.

And so one of the questions we wanted to ask is depending on how often an application requests access to data at runtime whether having the user get prompted more than once is appropriate, because certainly there are cases where the first time an application asks for data, it's completely intuitive why it wants that data, such as you click a button for finding things near you and it asks for location. You know, there's context that allows the user to understand that that's probably an appropriate use of location data.

But then once that's granted and the user doesn't see another disclosure again, the application might be using it for completely inappropriate uses. And so our finding here, over 200 requests per hour, suggests that clearly prompting every single time data is requested is totally infeasible. But by the same study, we also found that 80% of our participants said that they wanted to block at least one request. And actually they wanted to block over a third of the requests in total.

So there must be a better way. One of the things that we found was that their expectations about whether a request for sensitive data was appropriate was simply based on whether they expected of the application-- well, sorry. Let me rephrase that. Whether they said an application should be denied access to the data was entirely based on their expectations for why the data might be requested. Which suggests that if you could have a system that infers how the data would be used, you can then make more reasonable decisions about when to automatically grant or deny access to sensitive data.

Another predictive feature was visibility. So we found more often than not applications were requesting data when the users weren't even using them. So applications running in the background. Or even when the screen was turned off. Many applications access things like location and other sensitive data types.

Now, don't get me wrong. There are often many use cases where this is perfectly legitimate. For instance, various apps that might need access to real time location data while they're in the background. And so this brought up this notion of contractual integrity. So Helen Nissenbaum has written extensively about this, that basically privacy violations are about expectations and violating norms rather than simply whether users' preferences are violated.

You know, often your preferences aren't really the same thing as your expectations. It's reasonable in many cases for preferences to be violated, and users to be OK with that because there are other considerations. But instead, maybe if we look at expectations and have systems

that are designed based on what are reasonable expectations for when applications should have access to data, that might be more acceptable to people.

And so to do that, we need to design intelligent systems that only prompt users when the data's likely to be used for an unexpected purpose. But obviously this is a really difficult problem to automatically infer when access to data is likely to be expected or not. That relies both on figuring out what the user's expectations are as well as what the purpose of a given data request is by the system.

And so ideally what we want to do is automatically allow access to sensitive user data when the user's likely expect it, deny it when the user is likely to not expect it or we know that this is definitely to violate their expectations. We know what the purpose. But then only prompt the user when it's absolutely necessary, when we can't infer expectations, or maybe because we can't infer how the data might be used, the purpose.

And the reason for this is habituation. So if we prompt the user in cases where they don't need to be prompted, we know what their decision is likely to be. It's simply going to habituate them to swatting away these prompts because it's not really imparting any new information on them.

So how do we do this? Most recently we did another field study. We instrumented Android phones, and we decided to collect some behavioral data, too, about their privacy behaviors. So what apps they use, whether they play with security settings, that sort of thing to see things that were reasonably correlated with privacy preferences. And then we gave them these periodic prompts.

So we did experience sampling. They would see one of these prompts, which gives the name of an application and a resource, sensitive data type, that that application had just accessed. And we asked them whether or not they would have allowed this or denied it if given the choice.

Because we didn't want to overwhelm them, we did probabilistic sampling so that they saw at least one of these-- well, no more than one of these a day. I think on average they saw about five or six of these prompts a week. And we did sampling to optimize for a breadth of different applications and different data types.

And so from this, we had 133 participants who were using these as their primary phones for about a month. This resulted in 176 million events. And by "events," I mean applications accessing sensitive data. And from these 133 participants, five or six prompts a week, we collected over 4,000 responses to the prompts which we used as a dependent variable.

The behavioral data that we collected, I'll just go over this briefly, but we collected some behaviors such as audio preferences. So maybe if you use your speakerphone a lot you're less concerned about privacy, especially when you're doing it in public places. If you manually lock your screen when you put the phone down as opposed to putting it down with the screen on, that might say something about your privacy preferences. And so on.

And so from this with all of these responses to the prompts about whether they would have allowed or denied, we could model what's actually happening currently in the field with the ask-on-first-use model. So again, the ask-on-first-use, the first time an application tries to access data, the user is prompted. But they never see a prompt again.

With the data that we collected, we can now model what would their expectations be in the future if they were to be prompted again so that we can measure the accuracy of the status quo. And what we found is that the status quo, ask-on-first-use, seems to work about 80% of the time. So almost 20% of the time if they were prompted in the future for subsequent requests, they would have denied some of those. During the course of this month with ask-on-first-use, that resulted in them seeing about 12 prompts.

And so if we use machine learning and train a model just based on these behavioral features that I mentioned on the previous slide, we could do away with prompts entirely and customize this on a per-user basis based on very behaviors, which would result in an error rate that's equivalent to the status quo right now. We could also do a lot better if we have twice as many prompts to really gather their privacy expectations. We find that the error rates reduce to about 4%. Or if we limit the system so that we don't prompt any more than the current status quo, we get an error rate that's about two-thirds of the current error rate.

And so there are a lot of open questions if we're probabilistically deciding when to ask permission for access to sensitive data in the interest of not habituating people to expected data access. There are a lot of legal issues. This could also be applied to other domains where there's less user interaction. Maybe wearables, the IOT space.

And I'll leave it at that.

[APPLAUSE]

JOE CALANDRINO: Thank you, Serge. Tamar's up next.

TAMAR KRISHNAMURTI: Thank you. Very happy to be here today to talk about some work I did with my research team at Carnegie Mellon. This is in a domain that we haven't yet touched upon today, but we're focusing on creating a patient-centered approach to informed consent for participating in clinical trials.

So in theory, providing informed consent to enrolling in a clinical trial should be pretty simple. There are benefits, there are risks, and people have preferences for those things. And if they receive the benefit and risk information, they should be able to, based on their preferences, make a decision about enrolling in a clinical trial.

But we wouldn't all be here today if these disclosure situations were simple. With clinical trial situations, consenting situations, there are many different motivations. Physicians are motivated to heal their patients. Pharmaceutical companies are motivated to evaluate their products. And even the family and friends of patients have their own information priorities in terms of informed consent and subsequent enrollment decisions for clinical trials.

If patients could receive unbiased information in the form of a consent form, which they do receive when they're considering enrolling in trials or medical treatments, as you can see in the cartoon, often that information is so dense and so technical and in such a long paper document that they just experience cognitive overload and tend to kind of flick through it and not really read it.

So what is it that patients need to know in order to make an informed decision? Well, there are many existing standards-- international standards-- for good clinical practice that have just determined what information is critical for a patient to view in order to make an informed decision. These are things like your participation in a study is voluntary. It won't affect your usual medical care. And partly due to liability reasons, that information, when it gets translated into a paper document, tends to get translated into this very long document causing this problem.

Now, with new technologies, there are new options for delivering this kind of information. It doesn't necessarily have to come in a big stack of papers anymore. We have things like telemedicine and web-hosted videos to provide informed consent information. And that may change not just the format, but the type of information that's delivered.

And so with all this background, our research team were curious about how we could construct an informed consent document that really got at the heart of what patients care about. Sidestepping all of the liability, but really seeing what is it that patients want to know when they're sitting down to make an informed decision about a medical treatment? We also wanted to see if that consent form that we generated from patient preferences also met normative guidelines, because these guidelines are there for a reason. We want to make sure that patients' preferences would match those.

And then given that there are all these new delivery formats for informed consent documentation, can we take this patient-derived patient preference consent form and deliver it in a different media than a stack of papers? And then if we do successfully create a patient-centered consent form, will it perform at least as well as our traditional documents that we always use when recruiting patients for and enrolling patients in clinical trials?

So to do this work, we partnered with ICON. They're an organization that provide clinical trial services to pharmaceutical industries and medical device industries. And they gave us one of their lengthy 17-page long consent forms for an injectable asthma treatment.

We worked with an MTurk sample of self-reported asthma patients, and then we randomly assigned them to just one of four chunks of this long document. We took this long document, we broke it into four pieces, randomly assigned people to look at them. We also-- and there's been a lot of discussion about MTurk reliability today-- we embedded attention checks to make sure that they were really on task with what we were asking them to do.

And this is a screen shot of the exercise. So everyone had the same introductory text for context, and then they saw an excerpt. And each sentence was highlighted, and they could select the sentence-- as many sentences as they wished-- that they felt were pertinent to making an

informed decision about enrolling in a trial. They then were able to rate those sentences on how important they were to their decision making.

We then took all of that data, and we developed a method for generating a consent form. And what we did was we took the long consent form, we had coders code each sentence in that form into categories, and then we could automatically generate a consent form based on the preferences of a majority of our sample. But we were also able to analyze that consent form and see how did it differ categorically in terms of what kinds of information patients were interested in from that long traditional consent form.

And just as a note, we did this with a sizable majority. I think we took about two-thirds of our patient population and looked and generated a consent form to see what they preferred. You could use this method to generate a consent form that's focused on your target audience. So you may want to generate a consent form that's focused on women's needs. Or you may want to generate a consent form that's more pertinent to seniors.

And as you can see, patient preferences resulted in a consent form that was sizably smaller. So there was a considerable reduction in length. And there was also a change in the focus of the information they're interested in.

So in this short consent form, there's a lot more space dedicated to immediate risks. Things that patients will undergo in the trial. Things that patients might experience in the trial. And then a lot less space dedicated to things that may not be such an immediate need-- perceived immediate need. Things like privacy assurances.

We wanted to make sure that what patients wanted to see were also the things that experts felt that they ought to see. And it was a really good fit. There were just a couple of things that were different. Things like they didn't need to see the address of the person running the study, but because those things were part of the normative criteria, we popped them back into the short consent form.

And then because we wanted to see how does this translate to these other new media that are being explored for delivering informed consent from clinical trials, we made a video that was scripted on the short consent form so that the difference between them would just be the imagery of a video, but not the information that they'd be receiving.

And then we did a lab-based evaluation. We brought 76 asthma patients into our lab. We randomly assigned them to see either that patient-derived short consent form, the video that we'd scripted on that form, or that long traditional form that's used currently in studies.

And here are the results. So the blue is the short paper, the red is this short video, and the green is that traditional consent form. We have developed a knowledge test that was based on consultation with a legal team at the organization we were working with as well as their physician team. And we wanted to make sure that their knowledge that they got from the consent form was appropriate based on the desires of the people administering the consent.

And you can see that there was no lost knowledge. Everybody came out of that consenting situation with the same knowledge. If anything, the shorter forms tend to have a little bit more, but it's not statistically significantly different.

We also wanted to make sure that our shortened format and our video format didn't make people biased in one direction or another towards perceiving risks or benefits in the information that they were viewing. And there are no statistically significant differences in perceived risks or benefits. This is data for one question. We asked a series of questions, and the data all looks very similar.

And then lastly, we did achieve our goal of creating something that's much more engaging for patients. So the patient-centered versions were much more engaging than the traditional version.

So the takeaways are that we were able to, just by asking consumers, asking patients what it is they needed to know, create something that was able to communicate that information to them in a way that didn't lose their attention or have them skimming a document. And we didn't find any differences in the critical decision factors that they would need in order to make an informed choice about their treatment.

But there are some open questions. We've talked a lot about MTurk reliability, so I won't talk about it much more, but it is important in these kinds of situations to also do your research with a clinical population from an actual clinical setting. Because being in the moment of providing informed consent when you're actually considering undergoing a clinical trial could have some factors that are different. Maybe more emotional factors, more psychological pressure in that moment, and it would be good to be able to replicate that with another sample.

Also-- and these are open questions for research-- we can apply our approach to developing consent forms for different demographic groups, but what we really need to learn more about is how diseases differ on certain dimensions and possibly creating informed consent forms that are also tailored to disease needs as well as to patient demographic needs. OK, thank you.

[APPLAUSE]

JOE CALANDRINO: Florian, you're up.

FLORIAN SCHAUB: Hello. So I'm Florian Schaub from the University of Michigan. And I'm going to be talking about work that was largely completed when I was at Carnegie Mellon as a post-doc over the last two years. And the area I'm focusing on are privacy policies.

And we've been talking about privacy notices a lot today, as well as other forms of notices, and one of the things that is an issue with privacy policies is that they're just very long. They're complex. They're vague. We learned about that.

And the effect of that is that the information people are looking for, and Tamara talked about that as well, gets hidden, gets really lost in a sea of information. So what we're looking at in our research is how can we reduce inflammation overload and enable consumers to make more

informed privacy decisions. And one way we're looking at this is by simplifying disclosures based on expectations. And that's also similar to what Serge has been talking about, but it's a different context.

So we conducted a study where we looked at-- where we started with a privacy policy of a well-known fitness wearable manufacturer. And you can see this privacy policy's very long. It's three and a half thousand words. However, to their credit, they actually have a user-friendly version, which is more accessible but still quite long.

So the first step before our research was that we analyzed these policies and extracted data practices, all data practices, and came up with a compact disclosure format. And the goal here was also to be able to present this on smartphones.

In the next step, we were interested in determining the consumer's privacy expectations and awareness of data practices. So we conducted an online survey using Amazon Mechanical Turk in which we asked participants to look at a specific fitness wearable. So they were sent to the website. They were asked to research this product for two minutes. And they could come back to the survey and were then asked to rate the likelihood of certain data collection sharing practices occurring.

And we mixed actual practices with false practices. So here are some things that this device actually collects. So, for example, how far you walked. But also other things, like your shoe size or your perspiration rate, which are not collected. And we asked them how likely do you think it is that this device collects this information?

And what we found is that-- good news, first of all-- people are actually quite aware, have relatively good expectations of what's going on. So this is the percentage of correct expectations that something was occurring. These are the practices that are actually happening.

But we can also see these shorter bars. There are some practices that are less clear. So for example, the data retention policy. Collection of location information was something people were not aware of.

So we looked at how can we simplify or how can we use this to simplify the notices? And the first step, we removed data practices from our notice that more than 85% of our participants already expected. And what's kind of interesting is these are exactly the practices that are kind of associated with the fitness wearable device anyway. So collection of steps, collection of distance, these are things that are already expected because that's why you're checking out this device probably.

And the result of this is that our notice becomes a lot shorter already, and especially in the collection section you can see that we're now only highlighting these less-expected practices. So we were interested if we could take this a bit further. So as a second threshold, we chose 70%-- as you can see, there's kind of a drop-off-- and removed those practices. And that really only leaves practices in the notice that are unexpected by the large majority of people.

However, because this is a shorter notice, we also wanted to make sure that people are aware that these are not all the collection practices or all sharing practices that are happening. So in each of the sections of our notice we added a note, for further practices, please check the privacy website accordingly.

And then we tested these compact disclosures in an online survey. This was a larger survey, but it followed basically the same design as our baseline survey where we illustrated the expectations with the exception that after looking at the specific fitness wearable, participants would then see one of our notices and after that were asked to rate the likelihood. And we also had a control condition which was exactly the same as the baseline, so they didn't see a notice in that condition.

And what we find is that participants who saw the notice had significantly higher awareness of practices. So that's good. If you actually show people a privacy notice, they get something out of it. So that's nice.

In addition, we also see that we didn't find a significant difference between the medium and full disclosures. So it means when we reduce these most expected practices, it at least doesn't negatively affect the awareness. So it means there is some potential in removing highly expected practices.

However, we saw a significant drop in awareness for the short notice, for the short disclosure. And this suggests that maybe removing-- that there must be a sweet spot somewhere between 70% and 85% and we might have removed too many practices at this point.

An interesting side effect or side results is that we didn't observe any difference in the time spent on disclosure. So it didn't matter if they looked at the short notice, the medium notice, the long notice. Participants spent 20 to 60 seconds on this notice, which means if you have a longer notice, they're not going to read the information that is in there, right? They're going to read less. And to me, that suggests that notices should probably be even shorter to actually be effective.

So this is an interesting-- this is an approach to emphasize unexpected or surprising practices in privacy notices. We can further use this also to contextualize the information by eliciting these expectations for different types of services, user activities, or user goals. For example, in a recent study we looked at-- we found significant differences between expectations for banking websites and financial websites in terms of data practices.

But we can also think about personalizing information based on user characteristics. So in the same study I just mentioned, we also found that older adults had a better idea of some of the data practices that were going on. So that's something we can also use to then target specific demographics.

And we've been doing this in the mobile space where we actually analyzed the privacy settings of our study participants, and then learned privacy profiles from these settings. And based on these profiles, we can then-- and when we see a new user, we can ask them three to five questions, very general questions, about their privacy preferences and use that to assign them to

one of these profiles. And based on this profile and the apps they already have installed on the phone, we can give them recommendations on how they might want to restrict what those apps can do or can't do.

And in studies, we find that this is really effective. So about 80% of our recommendations are accepted. And only a small percentage, about 5%, are later changed back in a two-week study.

And what we're looking at right now is extending this approach of personalized privacy systems to the intent of things. So the idea here is that you maybe have your mobile phone which can act as your personalized privacy system and can aggregate disclosures and controls from multiple IOT systems and devices. So rather than having to have a notice on your smart thermostat, on your smart TV, and different things, you can have your mobile phone. And when you walk, for example, in the Constitution Building, it can tell you what data might be collected by sensors about you here.

And this can be used to then provide context of our privacy decision support. These systems can learn your preferences over time and can personalize recommendations and adapt to you. Serge was alluding to that, that some people-- we can train these models and can show them more or less prompts.

However, one thing that's really key if we want to scale this up is that we also not just have usable disclosures but also machine-readable disclosures. Right now we have long text documents, and in Joy Reidenberg talked this morning about how vague these documents are, which means they're really hard not just for humans to analyze, but also for machines to analyze. So machine-readable privacy disclosures can really help facilitate this idea of a more personalized approach to privacy disclosure and control.

So in summary, emphasizing unexpected or surprising practices in the disclosures is a really helpful approach and promising approach. We can leverage this to adept disclosures to specific contexts. And we can also personalize discloses to controls.

And there are lots of research questions that need to be answered here, but there's also really a need from the regulatory side as well to think about how can we support this and how can we provide legal security and safety for companies who want to explore this, right? So there are liability questions we need to discuss and figure out.

And in terms of methodology, so we find these online studies really effective to, on the one side, elicit expectations, but also to test different disclosure variants. However, it's also a key to really conduct lab and field studies. So once the disclosures are developed, conduct lab and field studies under real conditions because your MTurk population might be different from who you're actually targeting with your notice.

And with that, I am done and hand it back to Joe.

[APPLAUSE]

TAMAR KRISHNAMURTI: All right.

JOE CALANDRINO: Well, then thank you to all of our panelists. I think I'm going to start off, actually, with a question that relates to something that Florian mentioned near the end. So if we're doing things like applying user studies or machine learning or other methods to sort of constrain the details that are being presented to consumers, or to make decisions on a consumer's behalf, what happens when these approaches are incorrect? And how do we deal with thorny issues like liability? Or what should a regulator's role be here?

FLORIAN SCHAUB: Yeah, so I think one approach there is to be careful how these recommendations are framed and presented to users. So rather than saying, this is it, communicate the uncertainty or issues. I think that's a useful approach.

SERGE EGLEMAN: I mean, one thing that we're thinking about is whenever you have a system that's making decisions automatically on the user's behalf, there needs to be some sort of audit mechanism so that the user could go back and see what decisions were made on his or her behalf and then have the opportunity to correct those. So, I mean, that's an equally important problem, and also a pretty difficult one.

So if you give feedback to the user about the decisions that were made and how those decisions were automatically made, how do you scope that? I mean, if we're using statistical models that involve potentially hundreds of imports, what do you present to the user to tell them how this decision was made? And then also, how to change it in the future.

TAMAR KRISHNAMURTI: I would say echoing having some regulation. I think having a regulatory framework to always go back to as long as that regulatory framework isn't actually preventing advancement of technology is really critical. And then just constant user testing so that you're catching these mistakes. This is this auditing idea. Constant user testing, constantly bring it back to the consumer that you're working in their best interests for and see what they say.

FLORIAN SCHAUB: Yeah. And maybe to add to that, there's also this aspect that privacy preferences change over time, right? So this is another big challenge when you start automating this. You really need the feedback channel with the user. And I think one different problem that's really hard to solve is what do you do when your baseline information is wrong? Like, when the context information you're collecting is just wrong and different things like that, then that becomes really challenging.

JOE CALANDRINO: Yeah, absolutely. This is actually a point that Serge brought up connected to certain things like how much transparency there should be regarding what's going on behind this. What level of transparency should consumers expect regarding the approaches that are being taken and exactly how the information that's being presented to them is decided upon? And how can we avoid the possibility of things like deceptive omissions where things are intentionally possibly left out?

TAMAR KRISHNAMURTI: I would say that transparency is an ethical imperative in these situations. And it's not always realistic, because there can be things like IP issues where you've

developed a method, an approach, for creating these disclosures or an algorithm that you don't want to share with your consumer, but they should have some sense of how you've developed what you're doing and how you're targeting them.

And then also, talking from the informed consent for clinical trials perspective, they should always be able to access the full amount of information. So if you've developed some method that shortens their disclosure, say, and that you've disclosed to them that that's what you're doing, then they should always only be a step or two away from being able to see the full range of information that they might want to access.

FLORIAN SCHAUB: Yeah, I see these methods we're developing, they're really meant as a support system. So they're not meant to replace privacy policies. I think privacy policies play an important role in terms of regulation to demonstrate legal compliance, regulatory compliance, and for regulators to actually look at this and enforce compliance. But they're just not suitable for users.

And what we're trying to do, I think, and our research here is to kind of add a layer on top of that. But this layer on top does not prevent-- shouldn't prevent the consumer from going deeper and actually looking at the privacy policy. Or, in these more automated aspects, go back to your privacy settings and see what has changed. And maybe also calibrate the level of automation.

SERGE EGLEMAN: I mean, that said, I think that the current approach is entirely-- I mean, it's leaving us in kind of an untenable situation. I mean, even the experts can't agree more often than not what a privacy policy actually says.

FLORIAN SCHAUB: Yeah.

SERGE EGLEMAN: And so given that, if there's zero disclosure right now, having anything would be an improvement.

JOE CALANDRINO: Well, so one point that actually came up in some of those answers relates to having multiple versions of various disclosures and how can multiple versions of disclosures be leveraged? And even things possibly like machine-readable disclosures. How do you combine having different versions? And what are the best practices as it relates to having multiple versions of disclosures?

FLORIAN SCHAUB: Well, I think these best practices we need to develop. I think the challenge-- so one aspect there is what can we do from an engineering and user interaction side. The other one is what can we do from the legal side. And one of the big challenges is that privacy policies are often overseen by a privacy team that is maybe very cautious about how they want to present this information. So they might be very reluctant to have multiple versions of a notice because it means more risk of different versions being out of sync, being inconsistent.

So that's maybe something technology can also help with to really ensure that there is tracking. And machine-readable variants of, let's call them data practice descriptions, can really help. They could even be derived from a system at some point.

SERGE EGLEMAN: Yeah, I think that having machine-readable privacy policies would go a long way, because then it has-- if it has the full record of everything that's going on, you can have a version for the end user. I mean, some sort of parser that pulls out the things that are particularly important to that particular user. Whereas, regulators could then pull it out and get the complete version of the privacy policy as well.

But I think one of the problems that we see, as Florian pointed out, is that it's not necessarily the people who are writing the policies are completely aware of all of the practices as well. So when I was a grad student, we did some studies with P3P, which was a machine-readable privacy policy format, and we found certainly disconnects between the obvious practices that we could observe and what was stated in policy. And I'm willing to bet that more often than not it wasn't that they were being malicious, that they were intentionally incorrectly disclosing the practices, I suspect what was really happening is that the people who were writing the privacy policies weren't the same people who were developing the systems that were actually capturing and using data.

And so there's this disconnect. And, yeah, if we have systems that could then maybe do automatic audits of what code is doing and how data is being used and then update privacy policies accordingly, that would go a long way.

JOE CALANDRINO: But even for things like machine-readable versions, I mean, we have people in this room that have done work related to the fact that machine-readable version still-- there's some subjectivity there in terms of how you present data to users. How do you deal with issues like that? For anybody, actually, on the panel.

FLORIAN SCHAUB: Do you want to take this one?

TAMAR KRISHNAMURTI: No, go ahead.

FLORIAN SCHAUB: Yeah, so that's also an interesting question. Where lies the liability, basically? So when we talk about personalized privacy assistance, all of a sudden we take the presentation of the practices away from the company or the manufacturer of the device and we present it in a standardized format on a mobile phone, for example. But I think from a consumer perspective, that's actually a benefit because you can actually compare different practices and you have less of these framing effects that companies like to exploit.

So we talked about framing effects earlier in terms of how they can harm consumers. That's one of these aspects. So I think there are legal challenges, but there are also real consumer benefits to consider.

SERGE EGLEMAN: Well, I mean, in terms of what to pull out of the policy and how to present that to the user, I think that's actually another area where machine-readable policies offer an advantage because that way people make very different decisions, people have different preferences. One user might really care about one particular data sharing practice where another might care about another, and if you have a personalized privacy agent that parses the policies for you, it could then bring to your attention the things that you're actually concerned with.

But, yeah, of course that brings up liability issues. I mean, the companies that are publishing the policies, that might be an advantage because they can just disclose everything. But then the liability is shifted to whoever's writing the privacy agent and parsing the policies because it might decide to show you the wrong-- something at the expense of hiding something else that you cared more about. And then you might hold whoever wrote the privacy policy, parser, responsible.

FLORIAN SCHAUB: Absolutely.

TAMAR KRISHNAMURTI: I just want to say that I think this is an area where privacy and medical decision making probably differ quite a bit from each other. And I think that patients would probably find a machine-readable version of their consent a pretty scary thing to accept. But I think, in that instance at least, and maybe for privacy assurances too, having some kind of standardization of delivery would make a big difference so that we don't have so many differences and that we're not subject to framing effects and differences in delivery.

JOE CALANDRINO: I see. All right, so switching gears slightly, especially for your study, Florian, and your study, Tamar, you had sort of cutoffs that you had where you had to choose what-- or how much you were shortening a policy or what exactly the cutoff point was. How do you determine what that cutoff point is? It seemed as though it was kind of clear based upon your data maybe where you would choose to go, but it might not always be quite so clear. How do you figure that out in practice?

TAMAR KRISHNAMURTI: Well, in our case it was somewhat arbitrary. We chose a majority. But I thought about that a lot when we did, and I think that as long as preferences are stable-- and this also harks back to Florian's work-- I think if preferences are stable, then you go until that starts to change.

So for example, with medical consent forms there may be a lot of excessive information, things like very low-risk mild side effect information that wouldn't in any way affect the choice that you would ultimately make. So that's information that you may be able to access later, but it's not going to change your decision. Your preference would be stable.

So that's information that can get cut out. So that's in terms of shortening it. And then in terms of the cutoff of your population, how do you decide what the preference cutoff is for people, you want to capture the majority of people's preference being stable with the information that you've given them.

FLORIAN SCHAUB: Yeah, I mean, finding cutoffs is always challenging, I think. It should be data-driven. Like, it should definitely not be arbitrary. You shouldn't say, like, oh, 80 sounds good.

But in general, I think it's more what do we actually want to achieve with that? And in our case, it was we wanted to surface practices that aren't expected. So I think that's the better way of approaching it and thinking about, OK, what makes practices unexpected? Leveraging these elicitation studies, for example, to get a better understanding of that.

And Serge mentioned Helen Nissenbaum's concept of contextual integrity that actually provides a good framework for doing this kind of work by better understanding what is maybe not expected because it's outside of the context which you expect to use, for example in this case, a fitness tracker for. And then I think in the end, it's not necessarily about omitting information but rather surfacing these unexpected practices when it matters. So rather than having all of it in one policy, we can also show contextualized notices that really inform about one practice when this practice becomes important.

JOE CALANDRINO: Interesting. So I'm kind of curious about how your work can be applied to disclosures beyond privacy policies or medical informed consent, what your thoughts are even just going to new technologies like IOT or to new areas sort of beyond the areas that you're focused on in your own studies.

SERGE EGLEMAN: So I guess my lab's doing some other semi-related work on personalization and disclosures. So I've been collaborating for a couple of years now with a psychologist, and we've been looking at tailoring security messaging to different individual differences. So for instance, people who make dependent decision making, who look to others, they might be more persuaded by messages that say things like, 90% of other users made this safe choice, versus the same message given to people who test low on dependence might actually be dissuaded from making the safe choice. Wanting to be independent, they're not going to do what other people did. So we've been doing some work to try and infer different traits and also look at how different messaging for various security systems resonates with different types of people.

JOE CALANDRINO: Well, moving on, actually, so in some cases not only do preferences change over time, but actually policies themselves change over time. Like privacy policies might change over the course of several years. How does your work apply to things like policies themselves that might be changing over time?

FLORIAN SCHAUB: Yeah, so in related research, we analyzed privacy policies with machine learning and crowdsourcing, and one of the challenges is to keep this data up-to-date. But I think from a-- so we can manage that. We can [INAUDIBLE] policies and can figure out when they were updated and can then use that to redo our analyses.

But I think from a consumer perspective, the question is more what should an update to the policy mean to the consumer. I think right now if you go and read privacy policies, it says, oh, this policy may be updated any time, and you will be informed about the update on this page. No one clicks on the privacy policies in the first place. No one's going to check if the policy's been updated.

And then once in a while you get an email or you get a pop up message saying, our terms and services have changed. No one reads that. So I think the question is more maybe we need to hold the industry to a higher standard in terms of how they need to inform consumers rather than just saying, hey, we have updated something.

And one thing I think that could be really powerful there is to focus more on the consent. So if I consent to a particular version of the privacy policy and you want to change the privacy policy,

then maybe you need to ask me for consent for that particular privacy policy again. Or better, even the practices in the policy.

JOE CALANDRINO: Sure. All right, so we have only about a minute left here or so. Since we're talking sort of about the future of disclosures, I'm kind of curious about what sorts of movement you've seen from industry or by various groups to adopt some of the lessons that you've learned. I know that some of what you've done is still new enough that it might not have made its way quite so far out yet, but I'm curious what you have seen.

TAMAR KRISHNAMURTI: Yeah, well, we've partnered with the industry to do the work that I presented today. And so we've already seen the results of our work informing their approach to consenting patients for their clinical trials. But I can see the work move-- I think it takes a long time for these kinds of things to make their way into industry, but I can see that there's a movement towards personalization of medicine in general. And I can imagine that that will translate to consenting patients as well.

SERGE EGLEMAN: So our initial work about five years ago when we first started looking at access to sensitive data on mobile devices, as I mentioned in my talk, Android was just doing these install time warnings, which most people just sort of skipped over. And so since we published that work, they've shifted over to having a subset of runtime warnings for the, I guess, more sensitive permissions. And obviously that's not a perfect solution, but they've been making headway.

Apple has also expanded the number of disclosures that they have and iOS devices and also more geared towards the data types that people are more likely to be sensitive about. You know, like the access to photos or contacts and so forth.

JOE CALANDRINO: Excellent. All right, well, with that I think we will wrap up. And thank you again to our panelists and to everyone today.

[APPLAUSE]

And I'll pass it off to you. Yeah.

JESSICA RICH: Good afternoon. I'm Jessica Rich. I'm the Director of the FTC's Beaureau of Consumer Protection. Well, we've had a really interesting day-- this last panel was amazing. They were all amazing-- discussing the testing and evaluation of consumer disclosures. And I'm just going to do a really quick wrap up.

First, I always want to thank the team for organizing this amazing event. The group represents many different parts of the agency because, as you can imagine given what the FTC does, everything that we're into involves-- so many of the things we're into involves disclosure. So we really have everyone working on this. Of course, there's Lorrie Cranor, wherever she went, the FTC's Chief Technologist who lead this effort. She's taking pictures of me.

Mike Ostheimer from the Division of Advertising Practices. Hampton Newsome from the Division of Enforcement. Ryan [INAUDIBLE] from the Division of Privacy and Identity Protection. Joe Calandrino from our Office of Technology Research and Investigation. Jan Pappalardo and Laura Hoskin, over there, from the Bureau of Economics.

Nate [INAUDIBLE], Chris [INAUDIBLE], and Nat Wood from our Division of Consumer and Business Education. And also, of course, our event planning and web teams, the press office, and, of course, we'd like to thank you for coming, our panelists and our audience and those watching or listening on webcast.

So as you know, the FTC has a very long history of encouraging effective disclosures. Disclosures can provide important information consumers need to make purchasing decisions or avoid injury. They can qualify other information in order to prevent deception.

And I say "can," that disclosures can do these things, because if consumers don't see the disclosures, or the disclosures are confusing, or if there are so many disclosures coming at consumers they overwhelm them, they aren't much good. I think we've been talking about that all day. In today's economy, with the dizzying number of ways consumers interact with businesses and others online, offline, and through the multiple devices that are all around us, ensuring that disclosures are effective is as important as it's ever been. And testing is one key way to do that.

Today's event examined the many different elements at play in evaluating disclosures. We heard about how consumers process disclosures. We heard the different methods used to evaluate disclosures. We heard about studies of whether and when people notice, read, or pay attention to disclosures, and studies evaluating the impact that disclosures have on consumers' decision making and behavior.

We also heard about new approaches to disclosures. For example, systems that make decisions automatically based on consumers' past decisions and current context, reducing the number of times consumers have to make choices. And that's a pretty heavy concept, if you think about it. And, of course, machine-readable policies, which are not a new idea, but still have not been implemented the way it was once hoped and seemed to be as promising as ever as a way of managing this quandary we're in with disclosures.

Finally, our panels read some key questions to consider when testing disclosures, including what evaluation criteria and testing methods are most appropriate for a particular organization, what methods help you understand how consumers understand disclosures in practice, and how can your test results to help improve a disclosure.

In terms of next steps, we've posted a list of the resources that have been cited by panelists on our website. And we're soon going to post all of the slides and videos, too. So if there was anything of particular interest, you can look back over it.

Also, Lorrie Cranor and the team do plan to do a readout on the workshop and potentially make some recommendations and commentary on the FTC's Tech Blog. So please look for that. We'll

also obviously continue to encourage the testing of disclosures and pay attention to the results of the testing by industry academics and others.

And we're going to be reviewing the record of this workshop, and there may be other activities that grow out of it. We'll figure that out and go from there.

Most immediately, however, I want to let you know that some of the FTC staff and speakers will be gathering for some food and drinks after we're done. And any of you who would like to join them are invited. Are there-- there's instructions about where they're going and how to get there and how you can join them right outside when you walk out. So please don't hesitate.

And then finally, I want to put in a brief plug for our upcoming events and workshops, the FTC's upcoming events and workshops. In the next few months we will be hosting fall privacy events on drones and smart TVs. We'll be hosting this year's Privacy Con. The second year we'll be doing Privacy Con with more research.

We'll be co-hosting a roundtable on consumer perceptions about organic claims, holding a workshop on changing consumer demographics, and hosting a forum on crowdfunding and peer-to-peer payments, just to name a few of the events we've got planned for the upcoming next few months. Additional information is on our website. So please join afterwards, and thanks for coming today.

[APPLAUSE]

[JAZZ MUSIC PLAYING]

[MUSIC PLAYING]