

FTC PrivacyCon 2017
January 12, 2017
Segment 3
Transcript

SPEAKER 1: We're going to get started with session 3 on consumer privacy expectations. We have another great set of panelists here. And we're going to start with Jens Grossklags and Yu Pu who will be presenting their joint work.

JENS GROSSKLAGS: Thank you very much, and welcome to the last session before lunch. My name is Jens Grossklags, and this presentation is joint work with my co-author Yu Pu who will present the result section of our talk. The theme of our talk is predominantly focused on while not perhaps another trend, but something that has become more predominant that is that individuals are put in the position not only make an adoption decision for applications and to share their own data, but in the process of that, often they also share data about their friends, their family members, and acquaintances. And frequently these sharing decisions and also involve very personal and sensitive data about these affected individuals.

We refer to this phenomenon as interdependen

collected about the individual. So personal privacy,

to the existing work, we found that factors such as disposition to value privacy, as well as perceived control of personal data affects how people concerned about their own privacy.

So what is novel in our research is that we also explain the factors that drives people's concern towards other's data. In specifically we found that when individuals have higher privacy knowledge, is they would be more likely to place higher value on friends' data. In addition, we also found that the degree to which people care about others, which is measured by other regarding preference here, also affects how people concerned about others privacy.

In addition, we also investigate the relationship between privacy concern and the privacy valuation. Particularly, we found that when individuals have high concern n rs-10(edi)-2(t)-2(iv)-10(a)4(e)4()JT

Would that make you any more willing to take the dare? I would. It's still really creepy though.

So that's the takeaway. People can be uncomfortable, think that something is creepy, even though

In the rest of this talk, I'm going to build up to this conceptual model of how people decide whether or not they have privacy concerns. People run into a disclosure decision, and right away they get a gut feeling. The assessment of their intuitive concern.

If it's not creepy, then they don't think about it anymore. They just move on to their disclosure decision. If it is creepy, then that's when they assess considered concern, at which point they decide whether or not they're bothered by the intrusion.

of M Together, it largely eliminated privacy concern, even when people said they would have been concerned in other situations. But to an extent, participants trust in U of M is rational, but it could go beyond rationality.

A lot of our participants were unaware that M Together was gathering any of their data, which suggests that the extension never registered as creepy. So they didn't think about it. Never engaged in considered concern.

Others were a little more cautious. As you can see in subject 8 who says I was just flipping through. Yay. Whatever, install. And then when I went and looked back, I was like, wow. They must be collecting something in my computer.

So I guess I was maybe hesitant. I feel like that's not their motive to collect personal information from me, especially when it's coming from professors from the University. They're trustworthy people.

So here, he does get that gut feeling of creepiness once he knows what the extension is, but uses trust during his assessment of the intrusion is not problematic. And when we map that onto our conceptual model, he hits a red flag in intuitive concern, thinks it's a little bit creepy. But once he thinks about it, he decides he's not bothered.

So when we put this all together in the full conceptual model, we have the properties of intuitive concern mapping generally to system one, the properties of considered concern mapping generally to system two. And social presence, low marginal risk, and trust all having different impacts on the two types of concern. So I want to leave you with two i

concerns. That we separate out the expressions of intuitive creepiness. Just because they say they're concerned doesn't necessarily mean that privacy advocates should be concerned unless it's considered concern.

Rather, I think privacy policies should encourage congruence. Low considered concern? Don't require lots of privacy disclosures or things that are going to make consumers feel uncomfortable. Conversely, if there's high considered concern, then there should be high intuitive

the two CEOs of the companies are the best friends under the sun, it's just because there is

And another important implication is that the trackers right now we looked at a bunch of popular sites and the trackers and their privacy policy, many of them don't have a privacy policy. Even if they do, they do not say clearly what types of information they're collecting. So this makes it very difficult for ordinary consumers to make decisions. They don't know what information you're collecting about me. So they're suggesting is that either there should be an industry best practices, or there should be legislation that basically require trackers to say very clearly what types of information they're collecting or what types of information they're not collecting.

So I would end with that. And I thank my co-authors, Yaxing Yao, who is a doctoral student of mine, and [INAUDIBLE] who was a visit and researcher from the University of Rome. And if you're interested in more details about the mental models, please refer to our paper this year at CC. Thank you.

[APPLAUSE]

SPEAKER 1: All right. So we have one more paper in the session. And this is going to be presented by Mahmood Sharif from Carnegie Mellon University.

MAHMOOD SHARIF: Hello everyone. Today, I'm going to talk about a lab study exploring users in-context preferences for online tracking. This is joint work with my colleagues at Carnegie Mellon and Qualcomm.

I'm going to start by showing you a simple example of tracking with cookies. Cookies are those small tokens that can uniquely identify your web browser. They are used for many things including tracking.

The way they work is that when the user loads a website from web server, the web server might ask the web browser to store a certain piece of information. For example, that your ID is 1, 2, 3, 4. Later, when you load another page from the same domain, your web browser is going to send any cookie it has for that website or domain to the web server.

This enables the web server to identify you, and to serve you with custom content. For example, custom news articles based on what you looked at in the past. This is first party tracking.

Now I'm going to talk about third party tracking. Third party requests are those request that the user did not explicitly make. First parties might have content from many different domains.

For example, CNN might have ads from some advertiser. So when you load CNN, your web browser will send any cookie it has for that advertiser along with a request. This allows the advertiser to identify you and to serve you with targeted ads. This advertiser might be available on several first parties which enables it to learn more about you than any individual first party.

Experts have varied opinions regarding online tracking. Proponents say that online tracking allows targeted ads and customized content, something that both the industry and the users may find value in. They also say that the revenues that companies derive out of online tracking can enable them to provide free services to users. On the other hand, opponents say that online

tracking creates privacy concerns. That third parties can use it to build detail profiles about users, and that this can happen without users' knowledge.

We just saw what experts think. Now I'm going to talk about what users think. It's there it's important to understand

Now I'm going to talk about both the outcomes and the situational factors. Here are some of the

We believe that with more detailed and precise understanding of users' preferences, it's possible to build better tools to control online tracking. And to that end, and this is very preliminary work, we want to see if machine learning can be used to build better tools. On a very high level, we use machine learning to predict whether participants are comfortable or uncomfortable for specific page visits based on the situational factors that we found matter to them. This way we can build tools that can block tracking if users are predicted to be uncomfortable, and allow it otherwise.

Now I'm going to present the prediction results. So on this graph, the x-axis shows the percentage of bad tracking allowed, and the y-axis shows the percentage of good tracking allowed. Ideally, we want to block all of the bad tracking and allow only the good tracking.

Here's how we do. So in the blue area, we do fairly well. We block the majority of bad tracking and we allow some of the good tracking. While this is nowhere near ideal, this is only a first look on how machine learning can help us in this space. We believe that with more data we can do much better and design better tools.

So to summarize, we explored the users' preferences in the content of their own browsing history, and found which outcomes matter to them, and which situational factors affected their comfort. We evaluate current tools and found that they don't adequately address users' needs. And we show that there is some hope for automated preference enforcement. Thank you.

[APPLAUSE]

SPEAKER 1: OK, we're going to have a brief discussion before lunch. If any of you have questions, please come to the table before lunch. Thank you.

CHANDA PHALEN: I think that educating people on what exactly is happening is really important. Because at that disclosure decision, requiring a lot of reflection is going to backfire a lot. They won't even necessarily decide I'm bothered by this.

They'll just decide this is too hard of an decision. I'm going to think about it later. And then they never think about it. So either they'll be giving up the benefits of that privacy disclosure decision, or they end up doing something wildly wrong with their privacy. Something that would have bothered them if they'd been able to think about it.

JENS GROSSKLAGS: Attention of users is certainly one of the most scarce resources. So we have to support them in being able to do the decisions that actually matter. And I think what Chandra presented is one example of this broader space of bounded rationality research and the various kinds of burdens that we are facing in practice.

So we need to get rid of those kind of decisions that are impossible to do for users or are actually burdensome from an economic or psychological cost perspective. And identifying those is naturally very difficult, and may require also the involvement of technology. But the truth lead them to interventions, baseline regulation that frees us to focus on the things that really matter.

YU PU: So I want to say first of all education is quite important. Then I think of the mechanism and policies is also very important and to help individuals to make well-informed decisions. To apply to our case, we think it is very important to inform a user whether the sharing is anonymous, and or whether the data collected is important.

SPEAKER 1: All right. Let's take a question from the audience over here.

KRISTIN WALKER: Hi. I'm Kristin Walker. I'm a marketing professor from Cal State Northridge. And I love that you guys were talking about educating obviously. And I like that you're talking about decision making.

I wanted to ask a more macro question about expectations, because everybody wants information. I mean, it's pretty clear from what you say, consumers want information. Marketers obviously want information, policy makers want to make sure that that information is handled correctly.

But who do you think is right now shaping consumers' expectation about privacy? Is it consumers shaping their expectation? Is that technology shaping that expectation?
S6(e1.>>BDC -4.05p2(i)-(?
cdfnt t-6(r)3(e)udief2(a)aatib2(a)4(t(a)4(

YANG WANG: I guess I'll take that first shot. Great question. I think that your question about who is really shaping our influencing people's privacy expectations, I think is a confluence of many entities, stakeholders, certainly the markets where the industry has a pretty strong influence on us.

But I think increasingly, the news media, government, they also play a role. And also, not to mention the civil rights organizations that try to blow a whistle when these industries are doing

this certainty asymmetry, I guess you could call it. They have a very solid benefit that they're going to get.

Whatever they're signing up for, they want it. That's why they're signing up for it. They have this thing that they want right on the other side of giving up some of their privacy against this like maybe something bad will happen sometime in the future when the bad person gets-- There's so much there that when they have that carrot so close to them, I think they'll discount those potential losses a little bit more. Because they don't know what they are.

SPEAKER 2: This is a very dangerous comment, this is Washington, I would say that today a lot of people in the political parties are very, very concerned about their privacy. And maybe they were not as much concerned several months ago. And the point is that when you have a really bad example in front of your eyes, your psychology becomes also somewhat different.

JENS GROSSKLAGS: So clearly, behavioral economics and psychology is a great inspiration for all of us here on the panel. And we have used it to various degrees in our own work. I think the art and science of privacy research is to try to find out which of these behavioral concepts truly matter in the context of privacy, but also how to translate them to the specific domain of privacy.

So let's, for example, use the endowment effect. So endowment effect has been shown for durable goods prominently. So the first question then arises, how do we translate it to information goods? However, information goods are not really privacy.

So because once privacy is revealed, once you've given away your information, it's kind of gone. So how do you position that in the context of the endowment effect? And I don't want to elaborate too long on that, but it's just one example of where I think a lot of work, a lot of careful experimental work, but also theoretical reasoning is needed to make these concepts truly useful and impactful in practice.

YANG WANG: So I would just quickly that there is a big privacy project called nudging, which was done by Alessandro, Lorrie, and Norman Sadeh. I was on the team. So we were looking at a building design to hope to nudge people towards making appropriate privacy decisions by mitigating

experience and

don't have enough information. And when the confidence in the decisions are low, and try to pick safe defaults in those cases, like for example blocking online tracking, that case.

YANG WANG: Yeah, I think I agree with what you said. But I would caution that when you're making these automatic decisions, it can be very dangerous. Because if you think about the ethics of doing this, so why are we the researchers better positioned to make these decisions for the users.

You're taking part of their agency away. You can argue that they probably are not in the best position to make these decisions. But nevertheless, I think it's important to be cognizant about these agency issues.

run a statistical analysis to infer or how much utilities they put on each levels of an app. And therefore we calculate the monetary value of these privacy.

JENS GROSSKLAGS: So one tiny feature is that one of the attributes is actually money. And having the relative importance of the different dimensions of what an app could be composed of can then be translated with the help of this one monetary dimension in the monetary domain for all these attributes. So in absence of having some price as one of the dimension, you could now directly translate it into the monetary dimension. But since we have price as part of the attributes, we can translate also then personal privacy, interdependent privacy, and various other kind of factors that we checked for also into the monetary domain.

ELLE WINEMAN: Thank you.

SPEAKER 1: Another audience question.

MARK WEINSTEIN: Thank you. My name is Mark Weinstein. I'm the founder of MeWe, which is a privacy-centric social network that gives people an alternative to Facebook. Tim Berners-Lee is with us. The founder of the web, he's our adviser.

I'm concerned when I hear this panel talk about all this automation. You're now going to automate my privacy decisions, based on my preferences. I think Yang, I think you spoke eloquently about the dangers of that. Shouldn't we be presenting people with factual data about what's happening for every app for everything they're using with their privacy, and then let them select.

The earlier panel, they talked about we need noise, we need critical thinking. We just saw an election where fake news-- we know that everything at Facebook is algorithmically filtered for what our preferences are supposed to be. And we know that our opinions change based on information, based on factual information.

So I'm wondering whether the panel thinks shouldn't we be giving better information rather than trying to automate the preferences, the privacy preferences, and the cookie preferences of somebody? Shouldn't we actually be giving them better information? Regulating that, and then having them make a selection really every time they join an app?

MAHMOOD SHARIF: I think that is correct to some extent. We want to give users to make their own decisions. At the same time, it might be infeasible.

Because a lot of decisions they might need to make are just too many. And they will spend their time just making decisions whether they want to be tracked or not. So we should try to understand in a3Tw 24.5l6(c)6(t to)2(l6(c)7)6(s3t)-2-2(m)-2(>BDC 0.002 Tc -3,d2-2(m)eeMCID ,d2-2(m)pr

We're the

For example, we investigate whether the results indicated that people want a preferred to have an app that is free-- that costs maybe \$2 then choose a app that's free while the rest of others of the app are the same. So if the participants indicated that they would choose the costly apps, in that case we would think that he did not do very careful. He is not very careful when filling out our surveys. And hence we will maybe fail to their data.

JENS GROSSKLAGS: I think more generally one really has to use a portfolio of approaches to encourage participation to motivate people to contribute to the studies in a meaningful way, to filter out at the same time those that do not participate meaningfully, or may even be using automated tools to participate in studies. And so for that, we really need a portfolio. So you indicated one of them, which is essentially an economic filter. So if people make economically very unreasonable choices, then we could assume that they did not pay a reasonable attention, because people would prefer not to pay for something rather than to pay for something.

