Janice: Hi everyone, good morning or afternoon, or evening based on wherever you are. Thanks for tuning in for today's ET Speaker Series talk. I'm really happy to be able to present Blase Ur, who is one of my academic siblings. He is a professor at University of Chicago, and he is also a recipient of a Mozilla Research Grant, and I'm really excited to have him talk about his research that's really relevant to a lot of things that we're working on and doing here at Mozilla. So here you go Blase.

Blase Ur: Great.

Janice: We're going to be taking questions in #ET to if you have any questions, feel free to either direct message me on Slack or post in the ET channel.

Blase Ur: Great. Thanks Janice, thanks everyone for having me here, and today I'm going to talk about basically two separate streams of work, both of which were enabled by receiving the Mozilla Research Award about a year ago.

Blase Ur: Okay. So I'll start off talking about our work on privacy regarding third party online tracking. So I'll first present work that we published at the CHI Conference last year, followed by some work in progress. So happy to have this discussed internally, but the work in progress please don't discuss externally. Then after that I'll talk about some of our work on misconceptions about private browsing and how browsers can contribute to this.

Blase Ur: So let's start with the online tracking work. But actually first I'd like to introduce a little bit about my group at the University of Chicago. So the University of Chicago Computer Science Department has grown tremendously in the last few years. So I joined under two and a half years ago. Established a group in the usable security and privacy, and since then we've added four other faculty members in computer security for example. So we've grown tremendously and we have lots of great students coming out.

Blase Ur: So some of the students have been working on understanding privacy with online tracking. So this is work that my students Clair and Ben led and presented at CHI. It was a user study where we looked at how we could unpack receptions of data driven inferences. Those that underlay online targeting and personalization. So as all of us here know, advertisers use web activity to make inferences. So your browser behavior is algorithmically turned into things you might like. You might be interested in dogs for example. Then advertisers can use those inferences, inferences about interests, inferences about demographics to target advertising to you. So we were curious about how the method of targeting impacts privacy attitudes. That is how, on what basis is this advertisement being targeted to you and how do people feel about that. We were also interested in how the particular interests that are inferred, as well as the accuracy thereof impacts privacy attitudes.

Blase Ur: So to this end, we conducted this online user study, and the set up was a simulated ad. So participants signed up for this user study. They saw an ad, so for example this is an ad

for dog beds, and when they hovered over, they saw an explanation for why they got this ad. So perhaps the advertisers inferred from your online activity you're likely a woman between 18 and 24, thinks women between 18 and 24 are likely to be interested in buying this product. These explanations are things we varied by condition in the study. And particularly the different explanations that we showed reflected these prevalent methods. About half of the explanations that we showed, we said, "Advertiser has inferred that visitors to this site." So that is, you can think of this as more broad visitors or even closer to contextual advertising.

Blase Ur: The other half was the advertiser has inferred that you, in specific, might have this interest that follows. So for both visitors and for you we showed inferences about demographics, and we were also curious about how this sense of misprofiling or what's known as privacy distortion might impact attitudes.

Blase Ur: So we also showed wrong demographics, and we learned this at the beginning of the user study. We solicited the participant's demographics so we could choose accurate as well as inaccurate demographics. We did the same for things they were actually interested in, so we would show them explanations that, you're interested in this particular thing, for instance dogs, and this is why you're seeing this ad. We showed this in this interest condition for interests where the interest in the product being advertised were sensibly related.

Blase Ur: We were also curious how these not obvious connections, things like a dog to a dog bed in an obvious connection, an interest guitars to we might want to sell you a dog bed, that might be not obvious. But that's the power of big data. Perhaps some advertiser threw data at the ... Threw data into a blender and came out and said, "Oh well guitarists really tend to buy dog beds. We don't know why, but that's what's happening." So we wanted to test these unrelated interests as well. So that is, to recap, each participant saw this inferred demographic, wrong demographic, interest or unrelated interest, and half the participants saw one of these and it was to all visitors of this site have this inference, the other half was that you have this inference.

Blase Ur: We also tested two control conditions. The first was just the advertiser decided to purchase an ad on the site, right, whereas the other was that the advertisers' computer algorithms have determined that this ad would be effective. So this is essentially the magic of the algorithm. So these are our two controls. After they saw this advertisement, interacted with it, hovered over it, we asked a series of Likert scale response questions. Is this a useful reason for targeting? Was this notification informative? Would you like to know when this type of targeting is used? Are you comfortable with companies targeting for this reason. Do you consider it fair to target for this reason? Fair to collect the information needed in the first place? And are you annoyed by this type of targeting?

Blase Ur: So yeah, we collected these responses for each participant, and we went to compare participants in the different conditions, that those who were assigned to see explanations of all visitors as opposed to explanations or just you, you're the one that's

being targeted. We wanted to compare this control, this advertiser decided to buy this ad versus the magical algorithm has selected it. We wanted to compare, the algorithm has selected it as opposed to it's been targeted to this specific interest of yours, and then of course we wanted to compare interest versus unrelated interests and demographics versus wrong demographics to see whether this sensible relationship or plausible relationship as well as the accuracy of the demographic inferencing made a difference in attitudes.

Blase Ur: So now I'd like to take you through some of our results. So in general we found that participants were more comfortable with generalized targeting. So comparing all of the conditions where they were told that all visitors to this site as opposed to you. The visitors to this site considered this more fair to target. They also reported being more comfortable with this targeting. Then next pair of comparisons, we saw that the algorithmically targeted ads were considered more useful but less fair. So that is it was more fair to target when just an advertiser decided to purchase an ad on the site rather than an algorithm chose it for you, whereas of course the algorithm choosing the ad participants considered more useful.

Blase Ur: So this is getting at this tension in a paper I'd written read a couple years ago with my advisor at CMU and a number of other collaborators. We called it smart, useful, scary, creepy, were ultimately people's perceptions of online targeted advertising, they've said well targeting is smart and useful, but it's also scary and creepy, and they had trouble really understanding how to make sense of this. This work actually helped us further a little bit unpack this, and this is kind of our continuing trajectory as like how do we actually make sense of smart, useful, scary, creepy and help people move from they saying, "Well I don't know how to balance all of these." To helping them really feel empowered in terms of privacy while still gaining utility from their online interactions.

Blase Ur: So as you might expect, notifications invoking your own activity were considered more informative. So just comparing the magic algorithm to the advertiser think you're interested in dogs, and therefore likely to buy this product. Participants considered the latter more informative, and similarly, interests that were closely related to the product being advertised were considered more useful and more useful is a reason to target, as well as the notifications were considered more informative. That is participants considered it more informative that you're seeing an ad for dog beds because we think you like dogs as opposed to oh, you're seeing this ad for dog beds because we think you like guitars.

Blase Ur: That was our first study. So we were seeing some of these tensions behind kind of these methods for targeting, and on the second study we did as follow-up to this, we were really interested in how the particular inferences and the accuracy of those inferences impact privacy attitudes. As advertisers are targeting ads, for instance if they're using Google's AdWords, they have a lot of afinity audiences to choose from. It's on the order of about 2,000 different audiences that they can choose from when targeting the ad, and we were curious, how do people feel about these, targeting based on all of these different, for all of these different reasons, for all of these different interests. Right, and

so as it mentions, on the order of 2,000 for Google AdWords, if you're looking at kind of competing platforms sometimes you have a lot more. When you have series of interests that are algorithmically generated you can get into tens of thousands on some advertising platforms.

Blase Ur: Occasionally these are revealed to users but at a very high level. Actually some recent work from others has shows that these disclosures tend towards the more general, not the specific, and sometimes are in that way arguably inaccurate. So you'll see sometimes like, "Oh, this advertiser is targeting for this particular reason." And we've seen this a lot in recent efforts for political transparency. As I mentioned, there had been some fascinating measurement studies about are you seeing the real reason, and the answer is often no. We were curious, if you knew essentially a real reason for this hypothetical ad, how would you feel about these particular interest categories?

Blase Ur: Right, so this study we had 237 respondents on the Mechanical Turk crowd working platform, and in this case we tested 160 Google AdWords categories out of the about 2,000. We chose to do this subset so we could get more data, because of course the Mozilla Research Award only went so far. I would've loved to have gotten data for all 2,000. So what we did is we would show them this category in this hierarchy. Home and garden, gardening and landscape. You're receiving this ad because the interest in that. So we were interested in ... If the participant assumed that they were seeing an ad targeted them for this reason, first of all, how accurate would that inference be to them. We're interested in this because it's a potential confounding factor that we assumed based on the prior literature on privacy distortion that if the interest inferred is inaccurate or doesn't apply to you, you might feel differently, and we'll come back to that in the results. We were also interested in their comfort with this inference being made as well as their perceived usefulness of this inference being used for a personalization.

Blase Ur: So we found out on this first point, with the increase in interest participants were 10 times more likely to be more comfortable with personalization. So as the more the person was interested in this, they were 10 times more likely to be more comfortable with personalization, and they were three times as likely to find personalization more useful. So that is the sense of if it's actually relevant, if this is accurate about you, then you're going to be more comfortable and find it more useful, right.

Blase Ur: And what we were really interested in though, so that kind of confirmed this prior literature on privacy distortion. We were really interested in how these different categories differ from each other. One of the leading parts. One of the kind of aspects of the current online target that influence our thinking here was the way personalized advertising is handled is companies like Google have a small number of prohibited categories like personal hardships, alcohol, gambling, health, religion, sexual interests. Advertisers can't target on those prohibited categories. Right, and of course this is a small number, and there are a couple more than this, but out of those 2,000, it's just shy of 2,000 that you actually can't target on.

Blase Ur: So based on the prohibition of certain categories and the fact that companies allow the other categories, we were expecting, if we were to graph comfort, and again, we did this in a moment, or I'll show you what we did in a moment with this. If you were to graph comfort, responses to, I would be comfortable with the company personalizing my web experience based on an inference about my level of interest in topic. From strongly agree all the way on the right to strongly disagree on the left. So all the way towards the right would be very comfortable, all the way toward the left would be very uncomfortable, right. We would expect the distribution would look like something like this. That there would be the small number of categories on the left that participants would be very uncomfortable with perhaps these prohibited categories, kind of sex, religion, alcohol, gambling, and that most other things would be people would be pretty comfortable. So this is what we expected, but it's not what we found.

Blase Ur: So there were some instances where it might kind of match this. So for instance, participants strongly disagreed that they would be comfortable with personalization based on interest in dating and personals, or smoking and smoking cessation. On the opposite end of the spectrum, participants reported that they would feel very comfortable with targeting based on interests for instance in computer hardware. But it was not this essentially bimodal distribution, we saw actually something resembled a far more essentially a normal distribution. So most of these categories were kind of in the middle, and there were a couple of things participants were very comfortable with targeting based on, and some stuff they were not as comfortable based on. So this essentially there's a very different distribution of comfort.

Blase Ur: So let make take you through some of the interesting one. So some of the things we saw that were particularly interesting were that closely related categories could evince very different perceptions about comfort with targeted advertising. So for instance participants were generally pretty comfortable with targeting based on an interest in Christmas because yeah, who doesn't like Christmas? But closely related ... Christmas is traditionally a Christian holiday. Targeting based on Christianity, participants felt pretty uncomfortable with relative to these other categories. A similar case where we say two closely related categories differ quite a bit in perceptions of comfort were video games, computer and video games where participants generally felt very comfortable, and shooter games which were much more in the middle. In hindsight, on reflection, you can imagine why. Shooter games have a stigmatization around them. You think post Columbine, how shooter games were discussed in the media, but yeah, these very slightly different kind of in semantics categories could lead to actually pretty different perceptions related to comfort and privacy.

Blase Ur: So overall we saw actually quite a broad spectrum of comfort across categories. Right. Just to wrap with this first part, this study, we looked at this how this method of targeting impacts privacy attitudes, and we found a number of differences in attitudes based on the mechanism of targeting, that is is it based on interests or demographics, is it you or everyone who is being profiled, you or rather all visitors to this site. We also

saw this kind of very different distribution than one might have anticipated in comfort with personalization based on these different topics that AdWords uses.

Blase Ur: So how can we kind of move forward? When we did this work last year, we were interested in could advertising that works provide greater transparency. So if you look at the dashboards that are available on platforms like Facebook or Google you'll get some view into what has been inferred about you in terms of interest. That maybe it's been inferred that you like dogs. But could we do better? They don't tell you why they think you like dogs. They give you some limited aspects to correct this. What would it look like to say, this advertiser, "We think you are interested in dogs because of these things that you've done online." Would that actually help people understand what's going on better, have this transparency? And this is foreshadowing actually the next section of my talk.

Blase Ur: That's kind of from the user facing side. From the advertiser facing side we were wondering do advertisers tend to know on how this targeting impacts kind of perceptions of sensitivity and comfort. So perhaps if you ... If an advertiser buying an ad and they could then say, "Oh okay." They saw in general you can target on this, but users tend to be very uncomfortable with people targeting based on that, that might make them think twice about targeting based on these categories that are not prohibited, but perhaps arise, arouse privacy concern. So broadly, can we provide more transparency to advertisers about how targeting based on certain categories might be perceived, even if it's allowed that people might feel discomfort with it.

Blase Ur: So next I'm going to actually take you through some of the work we've been doing recently, and this is again work in submission where we said, well can we actually think more about transparency online, and perhaps can we actually try providing more fine grain transparency to users without cooperation from the advertising companies. So and again, we were saying, can we move from these kind of dashboards of here is a list of categories that have been inferred about you to explanations of really what's going on, what do they really know about you. Because this doesn't tell you that, on the bottom left, Google's Ads settings it's not telling you here is everything Google's tracked about you, or here is everything they know about you. It might just say we think you like air travel, we think you like clothing, and coffee, and tea. It's not very specific, and it's not just longitudinal data, and it's not connecting the dots. It's not saying, "Oh we got. You did these things, and now we concluded the following."

Blase Ur: So if we look at the spectrum of third party privacy tools available either as browser extensions or in browsers themselves. The model that tends to be most prevalent is this tracker based model. It's exemplified by tools like Ghostery, Disconnect, and many others of on the page you're at right now, and this is the example on the left, the page you're at right now, these are the trackers that are tracking you. So you get, it's scary. We've done a number of studies previously where we looked at what people understand, and what they understand is wow, there's lots of trackers, I never knew there were these many companies. There's so many people tracking me, I've never heard of most of them, but sometimes it's really confusing because Google is my email provider, why are they tracking me? "What does Google have to do with advertising?" Is

a quote that's come up a couple of times in user studies that I've done. Right, and so they get the sense of there is tracking going on, but not a great sense of what's being tracked or how these conclusions are made.

Blase Ur: So if we look at your Lightbeam tool we see these kind of, these connections across sites and third party trackers, but again, there's this implicit longitudinal data in there, but there's not the you did this, and this, and this, and this, and this and that lead to this. So we said, "Well can we build such a thing?" So we, my students recently built what we're calling our tracking transparency browser extension. Now I'll take you through the design of this and then report some preliminary results from a field study we did with this tool. So tracking transparency manifests as a browser extension, as you're browsing these pages, you can access this pop-up which will tell you the kind of typical things, "On this page there are five trackers." And we give this example tracker. But really what we're trying to do is provide this longitudinal data and connect the dots for people.

Blase Ur: So if the user clicks on show me more about what the trackers know, it takes them to our full screen kind of browser extension dashboard where we have a modeled after things we saw in the while on the top some static text about essentially what tracking and interest inferencing is. Then we show participants longitudinal data in this interactive plug-in. Trackers that are tracking them, interests that potentially could've been inferred, and we'll talk about this in more detail in a moment because this is actually very different than what you normally in third party privacy tools. Then kind of more, more about their history.

Blase Ur: So let me take you through some of the different sections of the tools. So typical who is tracking you, Google, Amazon, Chartbeat, Twitter, Facebook, Optimizely, Criteo, so on and so forth. So we can provide participants this kind of longitudinal sense of which trackers have tracked them, what fraction of the pages that they're visiting are being tracked by these trackers. Again, this is somewhat typical of a third party tracking privacy tool. But we said well, this is good because you get a little bit of a more longitudinal sense of this, but what do they really know about me? So we said, well one thing we could show them is what pages they visited on which they were tracked by these different trackers, but really what's really interesting here is what has been inferred about them potentially. So we said, "Well, it would be great if there were easily accessible to privacy advocates, easy to use APIs for asking Google precisely what do you know and why do you know it?" Of course advertising companies are not incentivized right now to provide such things. The kind of exception is now that there's some regulatory pressure related to political ads, there's a little bit more transparency about that but not a huge amount still. With non-political ads you get very little.

Blase Ur: So we said, "Well, can we make our own that's essentially a plausible representation of what might be happening?" So we said essentially if we could build our own in the browser, all on the client's side, nothing going to our server kind of inferencing infrastructure, what might it look like? So we started with the approximately 2,000 Google Ad categories from AdWords. These are the 2,000 ish different buckets that Google AdWords can put you into, and we said, "Well, how do we actually say which

pages might suggest an interest in that topic?" So we connected each of these programmatically to different pages on Wikipedia, and then we used, we tried out all the different well known topic modeling techniques. We said if we were to use like LDA, TextRank, use neural networks, TF-IDF, so on and so forth to basically associate, to identify keywords that are highly associated with these particular topics, then we could actually have essentially this approach where we can just pretty much do keyword matching.

Blase Ur: So we did essentially ... The way it works is we, I have pre associated each Google AdWords category with a series of keywords, and then a user goes to a webpage. In the background our browser extension extracts the keywords and then basically associates each of those keywords, or that set of keywords and that is overall the page, with zero or more of these ad interest categories. We ran some comparative studies on this, which I won't have the time to talk about, and we found that this kind of shadow system, this ... Well, the study we did was we showed participants screenshots of web pages and then essentially what different categories our algorithm would put these pages into, and we tried the different algorithms, and we tried a control of random assignment, and we found that most of the time, although far from all of the time, we could come up with an interest category that seemed plausible to our MTurk users. Whereas of course in our random assignment it was very rarely did it come up with something plausible, which we would've expected.

Blase Ur: So we were doing much better than this random assignment baseline, although it's not perfect, although advertising inferencing in the real world is also far from perfect. So we could associate each page with these categories and we said, "But what could we do with this?" We could show participants based on your browsing history, we think that advertiser, we think that advertisers might think you're interested in home improvement or all these other categories. We display this as a sunburst, and then you can click on this and find out more about which pages did our algorithm associate with potential interest in home improvement. For a very quick snippet of the results of our study looking at these inferencing algorithms comparatively, essentially neural networks and TF-IDF for this particular task seem to perform the best, that is quite a bit better than LDA, and because of the computational overhead even the neural networks were slightly better than TF-IDF. TF-IDF is much more lightweight in a browser, so we used TF-IDF to do the interest assignment.

Blase Ur: So then, now you can look at these advertisers, see essentially sites where this, sorry, the tracker. See sites where this tracker has tracked you, and we actually can show you the whole history going back since you installed the extension. We can show you what are the domains on which this tracker has tracked you, and then we can also show, and this is very unique to what we're doing, topics that, or interests that seem associated with what this particular tracker has tracked. You can click on all of these and get more information about across trackers across sites. Where is this interest coming up in what you're doing online?

Blase Ur: We also show this kind of history of when you're tracked and these kinds of these patterns in your browsing. So overall we're providing this longitudinal in interest level in site beyond what's typical. This is something that we as researchers were able to do because we as researchers can be wrong about these interest inferences in a way that perhaps Mozilla or others can't, right? So we said, "Well what would happen if participants had this tool? How would that change their understanding and attitudes?" So we recently ran this 425 participant Mechanical Truk study where essentially participants signed up, did a pre-survey, downloaded our extension and then used it for a week. So over this week it just ran in the background basically building up all these visualizations. At any point they wanted they could actually go click on the extension and see what was there.

Blase Ur: After one week we asked participants to actually look at, to click on the extension, look through it and then after spending a few minutes looking through it, answer a post-survey where we elicited basically their knowledge about online tracking and their attitudes. So these 425 respondents, we randomly assigned them to use one of six variants of the extension. We had our fully featured one, which I just showed you. We had a variant where we took away this interest layer to say does this interest layer actually seem to make a difference. We basically also, using our same visual layout, made a version that essentially was Ghostery, just showing you on each page a little pop-up of which trackers are on this page. We forked Lightbeam and we had a variant that was basically Lightbeam rebranded with our visual aesthetic, and we had a control condition which was just the static text.

Blase Ur: So overall what we found is that these, compared to these control especially, pretty much all the tools gave participants some better sense of what's going on, but particularly our fully featured extension with this view of essentially what interest could possibly have been inferred. Again, this are our best guesses based on our client side inferencing. It gave participants a better sense of what advertisers might think about them and what exactly is going on. So in the abstract, a lot of participants did have some sense. With this control condition it says, advertisers profile you, they make assumptions about your interests, and participants did have a sense of that. But seeing these particular personalized examples did increase that somewhat.

Blase Ur: So we also saw that these, using these tools, our fully featured tracking transparency tool particularly, led to basically increases in intent to take privacy preserving actions. So what this series of stack bar plots is showing you is on essentially each horizontal row is one of the six variants of our tool. So up top we have, let me just [inaudible 00:36:51] this pen, that we have the essentially the static explanations, so that is no actual data, just a static explanation of what's going on. Here we have essentially just a controlled condition which is just your kind of browsing behaviors and those types of tracking, no sense of advertising. Here the connections is essentially our version of Lightbeam. Trackers is our version of Ghostery. This company is tracking you right now, but no longitudinal sense. And so then we have kind of this longitudinal sense of here is the

trackers that have tracked you, and essentially the pages they've tracked you on, but no sense of what interest they might have inferred.

Blase Ur: So this final version here is which interest, everything we've described above, well minus the Ghostery part, minus the Lightbeam part, plus the interests. So what we see here is that, seeing this interest layer leads to some improvement, although less overwhelming than we kind of initially hypothesized in kind of the idea that participants are going to for instance use private browsing. The shades of blue are the much more like, sorry, more likely to do so. The darkest shade of blue is much more likely to do so, and this is participants' self reported intentions. So they're interested in seeking out more info about what's going on using private browsing, broadly just taking this privacy protective actions. Interestingly you'll also see essentially this, with all these tools that are showing you data about what's going on, self reported less interest in clicking on ads.

Blase Ur: So overall this is showing you kind of where transparency could be taking people, and this is something we're continuing to work on this extension anticipating that we'll be open sourcing it once this work is published. Yeah, super interested to talk either offline or even now about essentially where we're going with this. Just a whole broad area of how do we provide people essentially more data driven information about what's going on in this online tracking ecosystem. How do we empower people through transparency is and will continue to be a very big interest of mine and my students. So I think maybe what we'll do is actually spend about 10 minutes going through the other part of the talk, or actually maybe I would be happy to take questions now if there is any about this series of the work.

Janice: Sure. Actually we'll go back to some questions from before. You showed ... We'll go back to some of the questions that we have been building up in the channel.

Blase Ur: Great.

Janice: So for the first study you spoke about you presented a distribution, and does that distribution of comfort change based on the accuracy of the interest category?

Blase Ur: Right, yeah. So yeah, that's a good point. I don't have at hand the distributions split out by basically accurate or inaccurate, and if I recall correctly it was roughly the same shape of the distribution, although it was shifted. It was actually the more accurate it was kind of the more right shifted it was toward the more comfortable.

Janice: Okay. And then the next question is, if computer and video games and shooter games elicit such different responses, is this due to the choice of label of the category which would be irrelevant? And if the category were presented to the user based on the collection of examples in that category, then would this be likely to change the results?

Blase Ur: Yeah, that's a very interesting idea. I don't have any data on that. Yeah, that's a fantastic idea. As an empiricist, that makes me want to go try some of these things.

Janice: Right.

Blase Ur: Yeah.

Janice: Okay, and then to the current questions or current topic of the transparency report. People were really excited to see the word cloud, just wanted to share that, I think that's really useful. Then let's see. And then my questions are, based on this information and being able who see who is tracking you and how long they've been tracking you, what kind of behavioral changes would you want people to make?

Blase Ur: Right, yeah. So I mean, I think the ... I don't want to be heavy handed. Well actually, first I'll comment on the word cloud. So we've kind of in developing this extension and I didn't really mention this as I went through, we worked with some kind of members in our department to specialize in InfoViz, who for what it's worth hated the idea of a word cloud, so this, the sunburst came from our InfoViz people, although we did a series of basically library intercept pilot testing where we went to one of the Chicago Public Libraries, had people try out a version of the tool with synthetic data, and we actually found that the word cloud was actually really helpful for a lot of people, so that's how the word cloud was in there and we kind of agreed that it provides this interesting sense of like what profile am I leaking. To address the question of what should people do about this. I think we really look at it like that there is not, we don't want people say here is this information and you must now do foo or you must now do bar.

Blase Ur: We're interested in really empowering users to make this kind of educated trade-off about, well you're giving this information out about your interests and you're receiving arguably some benefit for this in that more relevant ads, which is this is a big trade-off. We kind of made this observation going into the work that a lot of people don't really know what's going on. When we've done prior work where we had people come into the lad, use Ghostery on a computer and try to get us some sense of what people think is going on in terms of tracking their interests, making these inferences about that, and people were not very good at that. So we basically have been looking at it as by providing more transparency, we want people to make a more educated decision, and whether that is they want to take ... They want to block third party tracking or they want to say if this is what's being ultimately tracked, then actually I want this to be happening, I'd rather more relevant ads. We want to just kind of empower that.

Blase Ur: But of course there are these open questions about what exactly is being tracked, what exactly is being retained, and I think one of our goals more broadly with the work is to try and force this out of kind of the advertising ecosystem about would it be accurate to say to a user, "We have this big data base." Sorry, I'm on Widows, I have the live demo of the extension on Ubuntu, but I don't have it up. But you can basically scroll through on the bottom these pages where you can basically see every page visit that's tracked and actually go back to that page and say, "Oh, you were on SmugMug and the inference was this." And you can go back to this particular page. If this is the sense of, like if we want to represent to people that you're in this database, there is thousands and thousands of

rows in this database, that you did this, you did this, you did this, you did this, that's potentially going to have a very different impact on the way people perceive kind of the cost-benefit analysis of tracking as opposed to oh, this raw data is thrown out instantly, turned into interests.

Blase Ur: If people could see even these fine grain interests because if you look at again, this Google dashboard. This dashboard is a lot different than the dashboard you would see if you were an advertiser choosing an ad, but if you're an advertiser you have these much more fine grained categories. So here it might just say games, whereas the advertiser might say, "Oh, we really want people who like shooter games." So this is kind of hearkening back to that first work. So I think through the creation of these different visualizations we want to further dig into what impact do these different visualizations have and provide this transparency to empower these trade-off. And that was a long rambling answer because this is something I feel pretty strongly about.

Janice: I guess we have a couple of other questions.

Speaker 4: [inaudible 00:45:13] In the room.

Janice: Are there any questions in the room?

Jonathan: I have a question.

Janice: Okay [inaudible 00:45:19].

Jonathan: Thanks. The kind of ... Oh, okay. I'm Jonathan, I work on machine learning at Pocket. To kind of get back to the difference between the dashboard that the advertisers see and the dashboard that I guess the people that are subjected to advertising sees. One of the things about when we were trying to do like explain recommendations and things like, you tend to keep these things at kind of a high level for no other reason than it kind of like, you don't have to kind of shrug your shoulders and say, "Kind of is." Right?

Blase Ur: Yup.

Jonathan: Do you think that there's really, like there's a trade-off between this specificity here, and do you think that is actually more useful to kind of get back into these kind of finer grain things like shooter games versus you like games type level? Thanks.

Blase Ur: Right yeah, so the kind of, this question broadly about specificity is really interesting, and something that our team, we spend a lot of time debating actually. I guess one point I was a little bit implicit in what you're mentioning is actually the more specific you are in explanations, the more of a privacy risk there actually is to the user on a shared computer, and so this is something we spent a lot of time debating, because if you have a shared computer and then you now have these kind of explanations that are invoking kind of past behaviors, now you actually have kind of the opposite of private browsing, you have essentially this history that isn't cleared by clear history. There's actually some

problems to be worked out there with kind of finer grain specificity. I think that to answer your question about the specificity in terms of what you tell the user, this gets back into this question about transparency, which is if advertisers are targeting you based on your interest in one particular subcategory, and that's what you are in the database, is it fair to tell people that oh, they think you like games, people will make assumptions based on that.

Blase Ur: So yeah, this is something we're still unpacking and yeah, we did this preliminary field study, we want to keep doing work and trying to unpack these issues, but yeah, specificity is a huge open question for us.

Janice: Do we have any question to any of the other comments? Okay, I assume that's a no. So we have another question around one of the other slides you showed. I'm assuming you were not able to get data around ad clicking.

Blase Ur: Sorry.

Janice: I'm assuming you were not able to collect, that your extension does not also collect clicks on ad.

Blase Ur: We did not collect clicks on ads in this version.

Janice: All right. And do you feel like if you could do that, that the self-reported behavior and the actual behavior would be in line?

Blase Ur: Right, yeah. So this ... The questions of intent versus actual behaviors comes up over and over again in security and privacy research, and so yeah. I mean, in a future iteration of the study we would like to actually measure what people do, and it's ... We kind of looked at this as a sort of proof of concept of like let's build a thing, which actually I as an optimist was like, "Oh, this'll take like two months." And a year and a half later my students were like, "Okay, finally we have a version that we're happy with." Right, and so yeah, but this is it's, we hope to look at the kind of actual actions soon. And then, sorry, what was the other part of this question?

Janice: That's fine. And then my other question is from the plug-in and the kind of visualizations that you show, would it be useful or would it be possible to also have a button that says go an delete all of this from my Google history, or go an delete all of this from my Facebook history?

Blase Ur: Right. Yeah, yeah, so right now our extension doesn't provide any actionable actions, it's just is providing transparency. So we had debated do we actually want to, when we eventually open-source this, do we want to have a button of take these actions based on what I've seen, versus kind of directing you to here is what you can do in your browser, here are other tools you can download. I think we've been concentrated on the transparency part, although the more tightly a couple you make, here is the transparency, and by the way, if you want to take some action, here is the one click

solution. I think that that would be a nice thing, but yeah, that wasn't our immediate focus.

Janice: Do you think that's going to be something you're going to look into building or adding to the extension?

Blase Ur: Yeah, I'd like to.

Janice: Okay.

Blase Ur: Yeah, well I mean, I don't feel strongly about whether we build it or whether we direct people to saying, "Here is exactly what you do in Firefox and here are the button clicks." But yeah, as long as it's minimizing the number of button clicks to achieve your aim for the user.

Janice: Right, and another question from online. Were you able to research a model where the user volunteers their own interests and affinities, and whether that helped people prefer sharing their own type of, volunteering their own categories versus having these other ones inferred about them?

Blase Ur: Right, yeah. So actually at the end of the study we had this kind of hypothetical world where we said to people like you have these trade-offs. You don't have to do anything but your browsing is tracked and inferences are made about you, whereas an alternative world could be your browsing is not tracked but you have to tell advertisers what you're interested in. So we had some preliminary very hypothetical data that there might be some interest in that, or yeah, we looked at all of these kinds of different trade-offs. We looked at another trade-off which was like the typical free with advertising, paid versus no advertising. Yeah, I chose not to report that here. It's in our paper submission, although because it's so hypothetical, whereas a lot of the rest of the stuff we have is much more grounded in the actual usage of a tool. Yeah, I think it requires further investigation, but there seems to be some interest.

Janice: Right. Did you get an idea of whether people would be more comfortable if the tracking was done say client side rather than in the cloud?

Blase Ur: Yeah, that's ... I've been following that, the literature on how you do this with interest. It's not something we asked about specifically.

Janice: Right. Think people don't, have a hard time understanding what that even is or what that means?

Blase Ur: Yeah, yeah. We have a related project in our group in IOT of trying to help people understand edge computing and kind of ML on the edge and what means for privacy versus sending data to a centralized server. So this is I think a very open question, not just in our group, but broadly in the community.

Speaker 6: Can I ask a question that's a slightly sort of meta question to your overall framing?

Blase Ur: Yep.

Speaker 6: When we try and understand why people are doing this kind of tracking, why advertisers are doing this tracking, it's because we have this fundamental thing that the web brain now is funded by advertising, right?

Blase Ur: Mm-hmm (affirmative), yep.

Speaker 6: And if we have any hope of, all of these seem like very important interventions, but they're all, and please don't take this the wrong way, they're sort of the Band-Aid, and some of them are very good Band-Aids and very big Band-Aids, but they're Band-Aids, right?

Blase Ur: Yep.

Speaker 6: And I wonder, do you see any hope towards an alternate, some kind of alternate model of funding the internet that does not rely on advertising or do you think we're stuck with this for the next 20 years?

Blase Ur: Yeah, I mean, that's a great meta question, and I ... It is something towards the end of our study when we asked about these kind of hypothetical worlds, that is one of the ones we considered because I as someone who cares a lot about privacy and kind of user empowerment online, would like to see were there such a frictionless model, would people actually want it? And if you ask these kind of hypothetical questions to people, they're like, "Oh yeah, that seems good." But I think until something is built and actually evaluated in the real world. Again, as an empiricist, I hold out hope for that and I think that is a promising direction and I wish more people were working say, could we actually move towards this world.

Speaker 6: Yeah, I mean, I look forward to seeing. We've got it as one of the research questions for the Mozilla Research Grants this time around. I look forward to see what comes in. There is this flip side that there is this bizarre democracy enabled by the inefficiencies of the advertising, right? In that if I am in a slum [inaudible 00:53:59] and I've got an Android phone, I'm using the same Gmail client that Janice is using like sitting in mountain view, right?

Blase Ur: Yeah.

Speaker 6: And there's something really interesting about that bizarre democracy enabled by these inefficiencies, so.

Blase Ur: Yeah.

Speaker 6: I don't know, I don't know how you trade that off.

Blase Ur: Yeah, yeah. We'll talk about that more later, but yeah, I like the ... [Joe fish king 00:54:21] bizarre democracy. Characterization. Great, yeah. So, yeah I think these are all great questions. I want to actually just very briefly kind of present this kind of other work that we did looking at private browsing. I'll give the five minute version because I don't want to go over my slot. Actually I have ... We presented this work previously to a number of people on the Firefox team in a kind of voice, a video call last year, and I have been very excited to see kind of Firefox rolling out with new explanations of what private browsing mode is, which is the subject of this work, and we haven't yet tested these, but it's on my to-do list of things with students this summer actually to see if this new rollout makes a difference using our methodology. I assume that all of you have done some testing there too.

Blase Ur: So this is work that my students Miranda in UC presented at The Web Conference, previously known as WWW last year, called Your Secrets Are Safe: How Browsers' Explanations Impact Misconceptions About Private Browsing Mode. Broadly as we know, users have misconceptions about private browsing, that they think it stops tracking, they think it stops people from being able to geo locate them, they think it stops malware from being able to infect their computer. And our broad motivation in this work was to say, well what role do browser disclosures have in communicating what private browsing mode, what private mode actually does.

Blase Ur: So we define disclosures as the things that you see when you open a new window in private mode. What was really inspiring this work was that these disclosures are really different. The you've gone incognito with this little table of what it won't save and will save, versus essentially in Brave what it looks like versus how Opera conveys this, versus the old version of Firefox and what it conveys. Right. So these are all actually very different, not just in graphical layout but in content, and also sometimes in actual features, like what's bundled with private mode. Firefox you have this tracking protection that's sort of bundled with private browsing mode. We had also observed that the mobile disclosures vary even more, so it's a smaller screen real estate. You'll have essentially things ranging just from mostly this icon and a one sentence explanation to what led to the title of our paper, your secrets are safe on Opera Mini's version.

Blase Ur: So what we did is this online survey again with participants on Amazon's Mechanical Turk. It's about a 30 minute study, and what we did is each person saw one of 13 disclosures about what private browsing mode did on this hypothetical browser we made Onyx, and we did this to remove the branding effects, like what people thought about Google, or Mozilla, or their prior experiences. We wanted to remove that confound. So we redesigned all the modes to use the same text as Brave, the old and the new versions of Chrome which actually changed while we were pilot testing the study. Edge, Firefox, Opera, Safari, and then the mobile versions. So we showed each person one of these disclosures or a control we made up to basically be completely meaningless. So Onyx's private browsing mode protects your privacy and keeps you safe

as you browse the Internet. It was carefully designed by Onyx's engineers to let you stay incognito as you browse. So no actual information is being conveyed here other than a bunch of buzz words.

Blase Ur: So what we did is they saw this disclosure and then they saw 20 different scenarios, and for each of these scenarios they either saw what we call a distinguishing question where we would say, "You're using Onyx's standard browsing mode, what happens in this scenario?" Like someone is trying to track your location, can they do it? And then we'd ask the same question for private browsing mode. And so this is what let us say what do people think happens in standard mode and is it actually different in private mode? We were interested in also in some of these as comparative questions where it didn't make sense to say does this happen in this mode or this mode, it was more of a, a volume question. Does this happen more, the same, or less in private mode than in standard mode? Right, and so broadly I'll just skip to the punchline. We found that some things were not misconceptions. The vast majority of participants realized that photos aren't cached in private mode, history is not saved in the menu, files they download stay downloaded on their computer even if they download it in private mode.

Blase Ur: However, participants did overestimate a number of things. About half thought that searches could not be associated with them even when they were logged in during private mode. So if they were logged into a site, they thought the site couldn't know what they were doing. So this as people who are designing these tools seems like, "Well, this doesn't make sense." But this was an important misconception. Similarly a quarter of participants felt that their IP address could not be collected in private browsing mode. 40% that geolocation could not be estimated, and this is across all disclosures, and I'll talk about the differences in disclosures in a moment. So there was also a quarter of people feeling that private browsing mode made you more safe from malware, that sorts of things.

Blase Ur: One of the particularly concerning ones was the belief that private browsing mode protected you from tracking. Whereas in standard more, a small fraction of participants thought that their ISP, their employer or the government would be prevented from tracking. So the missing part is that the vast majority of people felt that their ISP, government, and employer could all track them while they're browsing. In private modes the kind of, the salmon colored, a lot of participants did feel that private browsing mode prevented them from tracking by their ISP, government, or employer, and that's generally not true. So individual disclosures varied a little bit. So we were interested in do these disclosures from these different browsers, desktop or mobile versions, do better than our meaninglessly vague control? And the answer in most cases was no. We didn't see a statistically significant difference in how many misconceptions they had across these 20 scenarios based on which of these disclosures they saw.

Blase Ur: The instances where we saw participants do better than the control were essentially these versions of Chrome. Again, this is the old version of Firefox, and now that Firefox has this new version which I guess sample size me I like quite a lot. We're interested in testing it there. Yep. Well actually I'll-

Janice: Can you test it? Can you run it through?

Blase Ur: Yeah, we'd like to, yeah. So we haven't done this yet, but we, it's on my to-do list for the summer, and if this is of interest to people in the room or remotely, we would be happy to do that actually.

Janice: I think the answer will be yes please.

Blase Ur: Okay, yeah. Great. Yeah, I mean basically I'll leave with one meta commentary which is that language does matter a lot. I noticed in the new Firefox version that the, it's a lot more clear about what tracking protection. One of our qualitative results from this study had shown that kind of the mix of tracking protection, participants would be like, "Oh, I'm protected from tracking by the government." No, that's not what tracking protection is, but if you don't understand this ecosystem of third party behavioral advertising, you can see how they would get there. But yeah, language matters quite a lot, so we picked on Opera for your secrets are safe, and so I guess we'll, I'll still pick a little bit on a Firefox Focus where if you go in the kind of your app store, it still says browse like no one is watching is the second sentence, and perhaps that's we don't want to encourage people to browse like no one is watching because people are watching. Yeah, I appreciate the allusion, the little pun, but yeah, and I guess we'll kind of end here.

Blase Ur: I have maybe time for one final question, and I'll be around actually for the next few hours and I'd be happy to chat with anyone who wants to chat about what we're doing, what you all are curious about, what we could do to help you answer questions you're curious about, or if you want to give us hard problems and we'll try our best, that would be lovely. I have a big group of students in Chicago who are all very interested in these things and pushing privacy and security forward on the web end in IOT and otherwise.

Janice: I don't think we have any other questions, so thank you very much.

Blase Ur: Thanks. Thank you very much for having me, and thanks again for the Mozilla Research Award.

Blase Ur: On the tracking transparency plug-in. He's the one who is doing the [inaudible 01:04:45], and then but ... I have some other students who are taking it over because this is something I like a lot. Hey.