FTC PrivacyCon
January 14, 2016
Segment 3
Transcript

-Next, we're pleased to have commissioner Julie Brill provide a few remarks. Commissioner Brill has long been a zealous advocate for consumers, their privacy, and the security of their data. And we're thrilled to have her here today. Commissioner Brill.

JULIE BRILL: Thank you, Kristen. And thank you everybody who's here as well as all of you out in TV land. Lunch may be over, but the feast of scholarship will continue. It's reall bou

want to communicate between services, you may be forced to use tools that only a few select

This combination of research capability and capacity for action also describes, just coincidentally, the design of the FTC itself. So naturally, we are a ready audience for research the sheds light on the challenges we confront in enforcement and policy development. And I hope the institutions that many of our presenters call home will be lasting platforms for robust exchange of ideas with the public and private sectors for many years to come.

So with that, let's hear what you have. Thank you very much.

[APPLAUSE]

And Dan-- Dan will introduce the next panelist.

DAN SALSBURG: Thank you, Commissioner Brill. Could be the next panel come on up?

Our first session today really looked at what kind of data is being collected about consumers. Our second panel looked at what consumers expect is happening with that data. And now in this session, we're going to look at what actually is happening with the data.

I'm really pleased to have with me researchers who are going to present three outstanding research presentations. And we're then going to discuss them. So why don't we get things started with a presentation from Michael Tschantz and Anupam Datta? Michael is from Berkeley and Anupam is from Carnegie Mellon. They're going to lead things off with a presentation titled Automated Experiments on Ad Privacy Settings.

MICHAEL TSCHANTZ: Thank you. I am Michael Tschantz, and this is going to be a joint presentation with Anupam Datta. We're going to talk about Ad Fisher-- a system for looking at online trackers and determining what information they are using about people to select the ads they show to people. There are two things I want you to take away from this talk.

First, it is possible to do this with scientific rigor, despite not having access to the internals of the system. And second, we can find certain flows of information, but we can't figure out why they happened. So let's get started by motivating the problem. Here's a web page. It's the Times of India. I find it an interesting e

understands that people have concerns like this, so they and other companies have provided things like the ad privacy settings.

Here's a screenshot of my ad privacy settings. It shows various information inferred about me. Google got my age correct, but got my gender wrong. Google also allows you to go in and edit this information. So if I cared, I could go in there and provide my correct gender. Google doesn't give us a whole lot of information about exactly how this thing is working, however.

So what we have is a situation where we have our web browsing behavior going into an ad ecosystem at one end. You have various things like ad settings sitting in the middle, providing sort of a window into how that ad ecosystem works, providing inferences they create and allowing you to put edits in. And then we see advertisements coming out the other end. But we would like to understand the flows of information in the system better than they currently make clear from their privacy policies and descriptions of how these systems work. And this is a difficult task because the system is opaque. We don't know what's going on in that ad ecosystem. Google and other online behavioral trackers won't share its source code with us. We can't do the traditional forms of program analysis. So we designed Ad Fisher, a system that allows us to run experiments on these kinds of opaque ad ecosystems.

Let me run through quickly how Ad Fisher works. Ad Fisher creates a bunch of fresh Firefox browser instances which simulate users. So these could be simulating people who browse various websites. It randomly assigns them to either a control or an experimental group. These two groups of simulated users will display different behaviors on the internet.

They then interact with the internet in various ways. And we collect measurements about how advertisers change their behavior towards the simulated uses. These measurements go into a test of statistical significance, which reports whether there's a statistically significant systematic difference between the experimental and the control group. If so, we know that whatever information describes the difference between these two groups and how they behave towards the ad ecosystem is information being used by the ad ecosystem to select ads.

So this is our main contribution. We brought the rigor of experimental science to these online black box experiments in such a way that allows us to detect causal effects, which are equivalent to flows of information with the theorem reproved. It does it with statistical significance, without making questionable assumptions about how Google operates. This is important because Google is an extremely complex system. Pretty much any assumption you make about how it operates might not hold. Or perhaps it even holds at one moment in time, but not later when you're running your experiment. And we provide a high degree of automation.

So now I'm going to give you an example of one of the findings we discovered with our system. In this experiment, what we did was we fired up our simulated users, and we had half of them set the gender bit to be male, and the other half to female on the Google ad settings page. We then had them all browse websites related to finding jobs.

We then collected the ads shown to them at the Times of India. And we found a statistically significant difference in the ads shown to the male and female groups. And this, in and of itself,

So what are some possibilities here? Which entity could be responsible? So one possibility is that Google's programmers intentionally programmed their targeting system to be discriminatory in this way. We considered that to be highly unlikely, but nevertheless, it's not something we can rule out because we don't have enough visibility or access into the system that they use internally.

Another possibility is that the advertiser, the specific advertiser in this case the Barrett group that was advertising for this career coaching service, might have indicated when they submitted their bid for the ad that Google should show this ad more to male users than to female users. And Google may have honored that request.

A third possibility is that perhaps the Barrett Group indicated that the ad should be shown to high earners. In fact, in response to the questions from journalists at Pittsburgh Post Gadget, the Barrett Group actually said they were targeting users who are over the age of 45, and who earn more than $100,000, because they thought that would be an appropriate group to target, people who want to go one level up and go to the 200k plus jobs. Now it could be these high earners are much more strongly correlated. There's a stronger correlation with the male gender than the female gender. And Google may have inferred that and then decided they should send more impressions of this ad to male users than to female users.

Yet another possibility is that other advertisers might be targeting the female demographic more. And there's some evidence that the female demographic is targeted more by advertisers because they make more purchasing decisions. And those other ads may have come with higher bid amounts, which took up the slots for the female users. And the males just got the ad from this particular service because they were the leftover untargeted. There were just more slots available for the male users.

Yet another possibility, and this would be the case of machine learning introducing discrimination, is that Google's internal systems may have observed that more male users are clicking on this particular ad than female users. And since machine learning systems learn from these kinds of observations, and they're trying to optimize the clickthrough rate, they may have started serving more impressions of this ad to the male users.

So all of these are hypothetical scenarios because we don't have enough visibility into the system

related email. It's pretty hard to tell, right? Nothing in the ad really tells you anything about how it's actually targeted. What about ad 2? It's about a hotel. What does this one target?

AUDIENCE: Homosexuals.

ROXANA GEAMBASU: I'm sorry?

AUDIENCE: Homosexuals.

ROXANA GEAMBASU: You got it right. That's exactly right, the homosexuality related email. Again, it's still pretty hard to tell. And it's not just the targeting of ads on Gmail that's hard to discern. Everything is obscure on the web. For example, data brokers apparently are using-- can tell when you're sick or depressed and actually apparently sell this information.

Or some credit companies, for example, are apparently trying to use Facebook information to decide whether or not to give out a loan. You may have heard of these things from the media, just like I did. But do you know that whether these things are actually happening, to what degree and how those things affect you? I bet not. People don't know too much about these things.

Welcome to the data driven web. Web services and third parties collect huge amounts of information about us. Every location, every site, every site that we visit, every click that we make and so on. And they leverage all of this information for all sorts of purposes. Some are in line with our interests. For example, we all love our Netflix recommendations or Pandora recommendations.

But other uses may not be so beneficial for us. And the problem is that we have absolutely no visibility into what happens with our data in this huge complex web data ecosystem. Who has access to what data? For what purposes are they using it? Is this good or bad for us? How do their uses affect us? And it's not just the end users that don't know how to answer these questions.

But society as a whole has a hard time answering these questions. And I believe the FTC does as well, from my communications with them. And that's very dangerous because obscurity and lack of oversight can lead to abuses, either intentional or not.

So in my group at Columbia we are developing new kinds of tools, which we call transparency infrastructures that shed light into this dark, data driven web. Our goal is to build really large-scale infrastructures that can go on the web and track the flow of information and reveal it.

So on one hand, we can increase users' awareness of what happens with their data online. And on the other hand, empower privacy watchdogs, such as the Federal Trade Commission, to audit what web services are doing with users' data and keep them accountable for their actions.

And over the past several years, we've been building a number of these transparency infrastructures. And we're continuing to do so now. And in this talk, I'll tell you about just one of these infrastructures in the remaining time, the latest, essentially public domain, transparency

infrastructure that we've built. Before I do that, I want to acknowledge my students and collaborators without whom I wouldn't be standing here telling you about these systems.

doing the targeting. And so this is kind of a different perspective. I'm going to give a different perspective on this problem.

You're all well aware of the kinds of issues that come up with a lot of data driven applications. You've probably heard of the study that was done about detecting differences in prices from Staples' online store based on where you live. This turned out to have some kind of correlation

separate data to actually validate these things, measure their effect sizes, and to check if these things are harming a large segment of a population. Is it very significant and so on. This is where there's a lot of technical machinery coming from machine learning. At the end-- actually, a lot of the work here is to make these findings consumable by the application developer. So something that's interpretable and that they can actually use to help them debug their application.

Let me give you an example. We've actually applied this tool to a couple applications, real applications that are data driven applications. One of them, the first one I want to tell you about, is this health care application. This is actually something that was produced by one of these machine learning contests-- er, data science contests. Some company, in this case, Heritage Health company, ran this competition where they provided some data about patients going to hospitals, a description of the patient records, how many times they've been to the hospital before, what were their symptoms, and things like this.

The task was to use this information to predict whether or not --whether or not or how many times the patient would visit the hospital in the following year. So there's a readmission rate prediction. What we did we do-- we looked at the winning entry to this competition. It was a pretty good entry. It was a certain application that was able to correctly predict with some pretty high accuracy, I think around 85% accuracy, whether or not the patient would be readmitted to a hospital in the following year.

This was the data driven application. It takes these kinds of inputs-- age, gender, number of times they've been to the hospital, and so on. And then it tries to predict whether they'll be readmitted to the hospital. What did we find by applying FairTest here? We found that there are some specific contexts where there's association between the age of the patient and how bad the predictions were, the rate of error or the size of the error in the prediction.

This was a contextual association that we discovered. It was not for the entire population, but for some well defined segment of the population. I think it was something like male patients who have been to the hospital-- who had been to the ER at least twice in the past year. But within this sub-population there was a really strong effect and a really strong association between age and the error in the prediction. This is an interesting finding.

We think that this is important in this social sense because this is something that could potentially lead to actual harms. If this application was actually going to be used for insurance purposes to adjust your insurance premiums and so on. So these are associations that can really have some impact on the users of the system.

I want to tell you about another application. This is not a real application. It's sort of a historical application. But I thought that would illustrate a different capability of FairTest. So this is a very well known data set. You can think of it as graduate school admissions application. It takes people who apply to Berkeley graduate school, and then decides whether to admit them or not. This is a well known data set from the 70s.

If you don't know about this data set, what happened was that they discovered there was gender bias in the admission rate at Berkeley. Men were being admitted at higher rates than women.

Indeed, FairTest can be used to discover this kind of association. We can try to explain where this association comes from. And indeed, this paper by Bickel et al in 1975 discovered that once you condition which department the applicant wants to get into, then the effect either goes away or the impact reverses. Women in specific departments would be admitted at higher rates than men.

What we want to do is illustrate how FairTest can be used to help a developer debug their system and try to explain what's going on, what's going wrong in their system. There's this other capability in FairTest for doing this. We call it providing some kind of explanatory variables. And this will really make this a real system, a real tool for developers to use to debug their applications.

Let me just make a few closing remarks. We also apply FairTest in a couple of other applications. You can read about in our preprint, which is available on the web. I already mentioned this other feature of explanatory variables. There's another big issue out there in data analysis, which is that of adaptive data analysis, where you want to be able to reuse a data set many times. This is something that we're starting to look at and integrate in FairTest. This is open source software that can be used by developers right now.

[BEEPING]

Really, what we're trying to advocate here is that we need to empower developers with better statistical trainings, better statistical tools to make these data driven applications more fair, more socially conscious, and so on. We think that FairTest is a good way to start here. Thank you.

[APPLAUSE]

DAN SALSBURG: Joining me on the stage now are discussors James Cooper of George Mason University Law School and Deirdre Mulligan of UC Berkeley. So we just heard three presentations about tools that are designed to shed some light on how data is collected from consumers, how it results in them receiving targeted ads, web content, or it can result in discrimination. Let me turn first to James and Deirdre. What are the common themes you see running through these three presentations?

DIERDRE MULLIGAN: I teach at the School of Information at Berkeley, and I spend

experiments with fake accounts that are assigned fake input sets or inputs. That results in some targeting. We are seeing, all of us, some targeting.

But it's not necessarily true that it's the realistic kind of targeting that real users would actually see. We may be losing a lot of the targeting that real users see. We may actually have targeting that real users never see, and so on. I think that's a big problem. I think we need research in designing tools that leverage direct data from real users to achieve some of the goals we have in our system, transparency goals we have in our systems.

That said, because I've been working and have invested so much in scalability, building scalable systems that can take many inputs. But the millions, you know, not the size that real users

[BEEPING]

ROXANA GEAMBASU: It's OK.

DAN SALSBURG: Well, with that, we will wrap up the session. So thank you all so much. The cafeteria will be open during this break. So you can get coffee without standing in a long, long line. We'll be back in about 15 minutes.

[APPLAUSE]

[MUSIC PLAYING]