FEDERAL TRADE COMMISSION

PRIVACYCON

TUESDAY, JULY 21, 2020

9:00 A.M.

VIRTUAL EVENT

FEDERAL TRADE COMMISSION

I N D E X

P R O C E E D I N G S

WELCOME REMARKS BY ELISA JILLSON

MS. JILLSON: Good morning. On behalf of my colleagues at the Federal Trade Commission, I'm happy to welcome you to our fifth annual PrivacyCon. My name is Elisa Jillson. I'm an attorney in the Division of Privacy and Identity Protection. My co-organizer for today's event is Jamie Hine, a senior attorney in the same division.

Before we get started with our program, I need to review a few administrative details. We're happy to welcome you via the webcast. We will make the webcast and the other workshop materials available online to create a lasting record for everyone interested in these issues. That will include links to the research discussed and, in a few weeks, a written transcript of today's event.

As you may know, PrivacyCon is typically an in-person event. If there are technological issues with this webcast, we will work to address them promptly and we ask in advance for your patience if any such issues arise.

We will be leaving time at the end of each panel to take questions from the audience. You can email your questions to PrivacyCon@FTC.gov. If you

would like to ask a question by Twitter, you can tweet your question using @FTC and #PrivacyCon20. Please understand that we may not be able to get to all of the questions.

Lastly, I wanted to thank all of the researchers and the panelists for their participation in today's event. We are very grateful for your work in this important area.

This program would not be possible without the great work done by many of our FTC colleagues. We would like to thank our colleagues that assisted us in reviewing all of the research submissions, including Monique Einhorn and Patrick McAlvanah. We would also like to thank those moderating panels today, including Ellen Connelly, Phoebe Rouge, Daniel Wood, and Lerone Banks.

Finally, this conference would not be possible without the help of Kristal Peters, Aryssa Henderson, James Murray, and Bruce Jennings; paralegals, Leah Singleton and Alex Iglesias; June Chang from our Division of Consumer and Business Education; Somethea Mam from the FTC media team; Juliana Henderson and Nicole Drayton in our Office of Public Affairs; and Shawn Whitaker at Open Exchange. Thank you all.

It is now my honor to welcome the Director of the Bureau of Consumer Protection at the Federal Trade Commission, Andrew Smith.

OPENING REMARKS BY DIRECTOR ANDREW SMITH

MR. SMITH: Thank you, Elisa.

Welcome to PrivacyCon 2020 and thank you all for being here virtually. This is the fifth year that we've held PrivacyCon, which brings together researchers from around the country and around the world to discuss cutting-edge issues related to consumer privacy and security. I know that you will all miss the opportunity to see each other face to face, but the most important feature of PrivacyCon remains the same, the spotlight on topnotch research from a distinguished group of academics, physicians, economists, and other practitioners.

Over the past few years, PrivacyCon has been critical in keeping the FTC and other stakeholders up to date on emerging technologies and related data and privacy security risks. PrivacyCon informs all of the work that we do here at the FTC, whether it be enforcement, business or consumer education or rulemaking and policy efforts.

In light of that influence, I'll start with a few words about what the FTC has been doing to protect consumers' privacy since the last PrivacyCon. Vigorous enforcement is at the heart of what the FTC does. And in the past year, we've brought privacy and

security cases under the Fair Credit Reporting Act, the Children's Online Privacy Protection Act, the Gramm-Leach-Bliley Safeguards Rule, and our own FTC Act.

Shortly after last year's PrivacyCon, we announced settlements with Facebook, Equifax, and YouTube last year that shattered prior records for civil penalties or consumer redress for privacy and security violations. These settlements also required important structural changes with respect to how these companies treat consumers' or children's information.

More recently, we brought a trio of cases against operators of mobile apps that failed to protect the privacy of children's information or misled consumers about compliance with children's privacy laws, the stalking app, Retina-X, the Swiss mobile gaming company Miniclip, and the kid's app purveyor HyperBeard.

In recent data security cases, like Tapplock and InfoTrax, we've put a stop to misrepresentations about smart lock security and also to the failure to safeguard consumers' sensitive personal information.

In privacy cases like Unrollme and Mount Diablo, we've challenged companies that made empty promises about keeping sensitive information, like

financial information and emails, away from prying eyes.

We have also focused on educating business and consumers about data-related risks. For example, in recent months, we've issued guidance to businesses on how to develop coronavirus-related technologies that take privacy into account. We've offered advice on secure cloud computing and tips for using artificial intelligence and algorithms.

For consumers, we've put out guidance on how to safely use videoconferencing services and how to protect children's privacy while doing remote learning.

Rather than talking about past accomplishments, today's conversation needs to be focused on what the FTC should be doing going forward. Panelists today will discuss technologies ranging from mobile health and disaster apps to interconnected devices, such as smart speakers and cameras, to online ad delivery systems. Economists will report on their studies abroad to gauge the effects of privacy legislation in Europe. And researchers will describe mechanisms for consumer choice and how consumers protect themselves from identity theft.

The papers presented today will highlight

technological developments that could be a boon to consumers, but that also present risks to privacy, security, and, in at least one instance, equal opportunity.

One final note before I turn the discussion over to our first panel: In our call for research papers, we specifically asked for research on mobile health apps, and the first panel of the day will be devoted to that important topic. Why health apps? Industry reports show that consumers are increasingly using a variety of health-related apps, including fitness trackers, mood journals, smoking cessation or addiction aids, heart rate or sleep monitors, fertility trackers, diet guides, and more. Use of contact-tracing apps during the COVID-19 pandemic could add a whole new dimension to that trend.

Earlier this year, the Department of Health and Human Services issued rules that will make it easier for consumers to access medical records through the app of their choice. This expanded access to health information could be an enormous benefit to consumers. But as we all know, wherever data flows increase, the opportunity for data compromise increases as well.

We, here at the FTC, have been active on

health privacy issues, with cases like Practice Fusion, PaymentsMD, and Henry Schein. And we won't hesitate to take action when companies misrepresent what they're doing with consumers' health information, or otherwise put health data at undue risk. Research like that presented today helps us to identify critical risks to consumers' health information or other sensitive data and better target our enforcement, education, and policy efforts.

And so I want to thank all of the researchers who submitted their work to PrivacyCon, and all of the researchers who are presenting their work here today. What you do is of vital importance, and we look forward to hearing what you have to say.

And a big thank you to everyone who made today's event possible. I want to thank Jamie Hine and Elisa Jillson for leading the planning of this PrivacyCon, and also the many other FTC colleagues from the Division of Privacy and Identity Protection, the Bureau of Economics, the Division of Business and Consumer Education, the Office of Public Affairs, and the Office of the Executive Director, who have worked together to make today's event possible.

Finally, thank you to everyone who's attending virtually. We appreciate the opportunity to

engage with the public on this important and cutting-edge research, and I hope that you enjoy the FTC's fifth PrivacyCon.

So our first panel begins at 9:20, and I'll turn it over to Ellen Connelly and Elisa Jillson for that panel. Thank you.

SESSION 1: HEALTH APPS

MS. CONNELLY: Good morning, everyone. Welcome to PrivacyCon 2020. I am Ellen Connelly, and my co-moderator today is Elisa Jillson. We are both attorneys in the Division of Privacy and Identity Protection at the FTC. We want to welcome you to our first panel of the day, which is entitled Health Apps.

We have five panelists here to present some very interesting research. First will be Quinn Grundy. Quinn is an assistant professor at the University of Toronto and will present her research on the data-sharing practices of medicines apps.

Next, we have Kenneth Mandl of Boston Children's Hospital and Harvard Medical School. He will present his paper on the privacy implications of moving health data, such as electronic health record information, to entities that are not covered by HIPAA.

Then we have Dena Mendelsohn of Elektra Labs, who will tell us about her work evaluating the privacy risks of connected sensor technologies in medicine.

We will conclude the presentation portion of our panel with John Torous and Sarah Lagan from Beth Israel Deaconess Medical Center and Harvard Medical

School, describing their effort to develop a practical framework to aid consumers in their evaluation of health apps.

We have more detailed bios for all of our panelists available on the PrivacyCon website at ftc.gov.

At the conclusion of the presentations, we'll have a question-and-answer period, during which we'll be able to have further discussion about the research presented. We'll be taking questions from the audience during the Q&A portion of the event, so please send your questions to privacycon@ftc.gov and we will try to include them.

With that, I will turn it over to Quinn to start us off. Quinn?

MS. GRUNDY: Thanks, Ellen, and thank you, again, for the opportunity to speak today.

I am very excited to share with you some work my colleagues and I did looking at the data-sharing practices of apps that have to do with medicines. And I'm hoping to spark some discussion this morning about how we think about data sharing within the context of the wider mobile ecosystem.

Could I have the disclosures slide, please?

I first just wanted to acknowledge that this

project was funded by the Sydney Policy Lab at the University of Sydney and that we have no conflicts of interest.

Next slide, please.

"Drug App Comes Free, Ads Included." So this was a headline that ran in The New York Times back in 2011. This app, which is really popular among health professionals, provides information about prescribing, drug information, and clinical conditions. This article reported, however, that Epocrates was generating the bulk of its revenue from pharmaceutical companies that purchased targeted, tailored advertising that was delivered to users on the basis of their personal characteristics and browsing history.

So we know that apps routinely and legally share consumer data with third parties and that this is done in exchange for services that aim to enhance the user's experience, such as integration with social media or to monetize the app. But what about health apps? I think we all have a sense that health information is particularly sensitive, particularly personal, and also that it is valuable. We know that little transparency exists around data sharing and, also, that threats to privacy are heightened when data

are aggregated across multiple sources.

But consumers are in a really difficult position and really have very little way of knowing whether their apps or websites that they use share this data and with whom. So we wanted to add to this ongoing discussion by specifically examining the data-sharing practices of a sample of apps that we thought were likely to share sensitive, specific health information that might be of high value to commercial stakeholders. So these are apps that provide information about medications, whether consumers taking medications or health professionals administering and prescribing.

We wanted to know exactly what data these apps collected and where they sent it and then to extrapolate from this data sharing to understand where that data might travel beyond third parties within the wider mobile ecosystem.

Next slide, please.

Our methods. So just quickly, and there is certainly more information provided in the paper we posted, but we looked at 100 paid and free apps -- sorry, we looked at the top 100 paid and free apps in the United States, UK, Australian and Canadian Google Play stores, and screened it for any apps pertaining

to medicine, so looking at the most popular apps.

We chose 24 of these apps that had some degree of interactivity. We designed a fake user profile, and in a lab setting, we interacted with these apps to simulate use. My colleague, Andrea Continella, developed a tool, Agrigento, that performed a traffic analysis to eavesdrop on the data sharing that these apps performed between themselves and the network. We analyzed the types of data shared and the IP addresses where it was sent.

We were able to identify the entities that had these IP addresses, and then looked at their websites and these companies' privacy policies to understand what they might do with user data. And, frequently, we found that they reported further sharing through integrations or other commercial partners. And so we were then able to identify what we called fourth parties and to simulate a worst case scenario of all the possible data sharing within this wider mobile ecosystem.

Next slide, please.

So in a sample of just 24 apps, a tiny fraction of the health app market, we found that the majority did share user data outside the app with the network and that some apps reported additional sharing

within their privacy policies. We had pre-specified types of user data that might be shared, including names, time zones, medications, or email.

Next slide, please.

We found that, most commonly, apps were sharing technical data, which might seem very benign on the face of it, so things like the device name, the operating system. But we did find that just over a third of these apps shared unique identifiers, such as Android IDs or email addresses. And a quarter shared the user's medication list, which is something that people could use to infer information about other sensitive things, like health conditions.

Next slide, please.

We conducted a network analysis of the data-sharing relationships between the apps and these third parties. So we identified 55 unique entities that received or processed user data, which included app developers and their parent companies and these third parties. We found that third parties received a median of three different pieces, or unique transmissions of user data, and as many as 24 different types of user data.

In this network, you'll see the orange nodes are the apps, and the size of the node is the volume

of user data sent or received. The blue nodes are third parties that we characterized as infrastructure and represented about a third of the recipients. These were providers such as data storage, cloud providers. And because of their business model, which often involves keeping information secure, we reasoned that risks to privacy were low from this type of sharing.

The gray nodes, however, are entities that were involved in the collection, collation, analysis, and then commercialization of user data, and this involved advertisers, social media, or analytics companies. And because of their business models and the way they described handling user data, we reasoned that there might be privacy risks associated with this type of data sharing.

Next slide, please.

So first parties. We're calling first parties developers and parent companies that were receiving user data in our traffic analysis. We found that they received both the greatest volume and the greatest variety. And that might be expected, as this data was likely used to enhance the service that developer provided to users.

However, we also found, analyzing their

websites and privacy policies, that developers were using this data for their own marketing purposes for products and services, but also the ability to tailor sponsored content, to sell advertising space, beyond banner ads, for example, and even to sell de-identified and aggregated data or analyses to third parties, like pharmaceutical companies or health insurance.

So for example, one app said they commercialized what they called the patient insights, from how medicines are used in the real world to healthcare stakeholders, like pharmaceutical companies. And so the sense that because developers were collecting information, that that might be safe and secure and private, may not, in fact, be entirely true.

Next slide, please.

When we looked deeper into the third parties receiving user data, there were 21 entities that we characterized as analytics. We found, when we analyzed their privacy policies, that these entities typically reserved the right to collect de-identified and aggregated data from app users for their own commercial purposes and to share these data among their commercial partners, or to transfer data as a

business asset in the event of a sale.

What was interesting was that for third parties, their privacy policies defined a relationship with the app developer, not the app user. And so if app users were concerned about the collection or sharing of their data, even if it was de-identified or aggregated, they were referred back to the developer in the event of a privacy complaint and couldn't take it up with the third party directly.

Next slide, please.

So fourth parties. We found that the third-party entities reported this ability to share end-user data with 216 different fourth parties, so entities beyond what directly received user data. And we found that these entities could potentially create highly detailed profiles of users, even if they could not identify them by name. So while certain data sources are clearly sensitive and personal, or identifying, like your date of birth or a drug list, others may seem irrelevant from a privacy perspective.

However, when combined, all these little pieces of information from a variety of different sources can create a fairly detailed picture of a user or to associate them with certain groups. So we conducted a network analysis to understand, again, how

data might be aggregated within larger companies and their commercial partners, and we simulated this hypothetical data sharing.

Next slide, please.

So this very busy picture is our fourth party network, and it's the worst case scenario, where, if all the data were shared by all these apps within the network, 44 percent of these fourth party entities may have access to medical information, and all but four of them also had access to potentially identifying personal data. We found that multinational technology companies, digital analytics, and advertising firms occupied highly central and prominent positions within this data-sharing network, with a significant ability to aggregate and potentially re-identify users. And of interest, only 1 percent of these entities could be considered as health-related or part of the health sector.

Next slide, please.

So in discussion, I think what these results suggest is that collection and commercialization of health app users' data is a legitimate business practice and that sharing of user data is both routine and far from transparent. Our analysis suggests that privacy regulation must emphasize the accountabilities

of both those that collect and control user data -- right now, a great deal of onus is placed on developers -- but also that process it, these third and fourth parties that sit behind the scenes.

I think we increasingly understand that the sharing of app user data ultimately has real world consequences. And I think the panelists in later talks today will be sharing some of these things, like bias in algorithms. These consequences include highly targeted advertising or the commercialization of data into algorithms that ultimately make decisions about people's insurance premiums, employability, or financial services.

We're seeing increased scrutiny of collection and sharing of sensitive, personal, or health data, but I think understanding how data are aggregated suggests that in combination, a much wider array of data types might actually be considered health data and used to make inferences about people and groups. So for example, even the existence of a health app or a mental health app on one's phone could be used to make inferences and decisions about a person.

Our current regulation focuses on securing individual informed consent through improving privacy

policies or labels for apps and protecting harms to individuals, for example, by ensuring that data are de-identified. However, when we think about the mobile ecosystem, the aggregation and sharing of data within this wider space, I think we also need to consider the disproportionate harms that can occur to certain groups when inferences are made on the basis of characteristic.

Next slide, please.

So in conclusion, I wanted to share our dashboard, healthprivacy.info, where the full data from this study are available, and it includes additional information about the security analysis we also performed and the apps that we sampled.

I'd like to thank Ellen and Elisa, again, for this opportunity, and to acknowledge my collaborators on this project, and in particular Andrea Continella, for developing the tool we used in the traffic analysis. And I'd like to thank, again, the Sydney Policy Lab and the Australian Communications Consumer Action Network, who we worked with.

Thank you so much.

MS. CONNELLY: Thank you so much, Quinn, for that really interesting presentation.

We're going to move on now to our next presenter, and next we'll hear from Ken. Ken, you're up.

MR. MANDL: Terrific. I'd like to thank the FTC organizers of PrivacyCon for putting together this spectacular program, and I'm honored to be able to participate.

Let me set the context for my talk. At the beginning of the Obama Administration -- and I assume my slides are going up -- at the beginning of the Obama Administration, Congress passed the HITECH Act, and the Federal Government invested $48 billion to promote the adoption of electronic health records. Because I had worked with electronic health records as a physician and a researcher, I knew that these older 1980s and 1990s software stacks would not advance the goals of a learning health system, where the data collected are put to work to improve health, control costs, drive discovery, underpin public health, and empower patients to manage their care and participate in research.

So I wrote in The New England Journal of Medicine a piece proposing that if we're going to invest this $48 billion of federal dollars -- which, by the way, was complemented by probably between a

half a trillion and a trillion dollars of private and public investment in installing these electronic medical record systems and purchasing them -- if we're going to do that, why don't we think about a public interface that essentially turns the electronic health record into a smartphone-like platform that can run apps that can be added or deleted the same way they could on the iPhone?

And when we wrote this, the iPhone was one year old, and we were just starting to see the power of an application programming interface that allowed third-party apps to connect to a platform. The type of business advances, the types of innovation, the competition that you see in an app store, the truly spectacular examples of apps that were emerging, could we have this for medicine, too, even though we were investing in older technology as the sort of backbone of our health IT infrastructure?

So we were funded for $15 million by the Office of the National Coordinator. And what we proposed was an application programming interface that would enable EHRs to run these apps. This was a high risk play, because each EHR was different, had no standard for the storage of data, and was not designed to ever let data out of its walls. In fact, quite the

opposite.

Patients had some access to their electronic health record through portals. Many of you may have used them. But those data are essentially behind glass. You can look at them, but you can't get a computable copy. You can't feed them into a computable process, like an app or an algorithm.

Now, HIPAA, passed in 1996, guaranteed that consumers could get access to a copy of their data in an electronic format if it was feasible. And from 1996 until, essentially, a year or two ago, it was determined by healthcare and healthcare IT vendors that, in fact, it was not feasible. Now, whether that's true, I think, is a subject of debate. But the good news is that now, 10 years after the $48 billion investment began, we have actually new regulation that comes from the Office of the National Coordinator of Health Information Technology, an HHS agency that oversaw the $48 billion investment and that funded us and that now has passed regulation based on the 21st Century Cures Act.

I don't do very much lobbying, but I managed to get this one sentence into the 21st Century Cures Act, requiring an API that provides access to all data elements of a patient's electronic health record, and

that those elements can be accessed without special effort. This underpins the potential for an extremely robust apps economy.

A second API was also developed in our group and managed to make it in under the wire into the regulation, which allows us to get data on populations out of electronic health records as well. The first API is called SMART on FHIR.

Next slide, please.

And these two APIs together allow us to potentially think about healthcare innovation in a parallel way to how Tim Berners-Lee thought about the Web. I think the slides might be a bit ahead. There should be a slide of Tim Berners-Lee showing now on the World Wide Web.

In a sense, what we're trying to do for healthcare is similar to what he tried to do. He wanted to share pre-prints of his articles, and he invented a way to show those articles in HTML. He invented a Web server so that you could serve up those documents. He invented HTTP so that you could link to them, and he invented a Web browser so you could display them. All of these documents -- what Tim Berners-Lee created parsimoniously, and then instantiated through the World Wide Web Consortium,

enabled a tremendous economy to be built on top of these parsimonious rules and specifications.

The APIs regulated by the Office of the National Coordinator, stemming from the 21st Century Cures Act, actually have the potential to create innovation within the healthcare domain.

The next slide should have a picture of the Apple Health app with the heart on it. And the first major company to take advantage of these APIs, even before these final regulations, based on some earlier regulations, was Apple. And Apple had a spectacular success. They used our API, called SMART on FHIR, to connect the health app to hundreds of health systems so that patients at all those health systems could download data from the health system onto their phone and expose it to other apps.

And there it is. There's the API and the health app being announced on the Apple stage. To the right of the health app, you see this little Blue Button 2.0. This is less well known, but it's actually a very important effort, made by CMS, to enable all consumers to have access to their claims data through the same SMART on FHIR API. And as I mentioned, though not the subject of the talk today, a second API, called Bulk FHIR Access, is going to give

us data on whole populations.

The next slide has a picture of the USCDI. The data that we're talking about is regulated as the United States Core Dataset for Interoperability and defines which health system data will be available through these APIs. This data set will expand over time, but now includes things like medications, diagnoses, laboratories.

The next slide shows the data protected by HIPAA on the left and the SMART API in the middle, where the patient can request the data, for example, to be downloaded into their Apple Health app. And then the magic that happens here is that the patient gets a copy of their data.

The regulatory piece, which has not been fully addressed, is that the data goes from HIPAA-covered, in the health system, to FTC-covered afterwards. And what happens as the data are passing across the API is critical for protection. The FDA has the most enforcement power over privacy in the US, but it does not prescribe what those privacy requirements are.

The next slide shows some aspects of privacy policies that are in the rule, that they be written in plain language, that they be made publicly accessible

all the time, that they include statements of whether and how the data is accessed, used, or sold, that they share this with users before accessing the data, and that they require express consent. So it establishes some elements of what needs to go in a privacy policy, and that is a good start.

The next slide, Analysis of Current Approaches, shows us that, yes, there are a few community-based efforts to address this. There is a model privacy notice. There are questionnaires that some of the electronic health record companies have actually developed to ask app developers what their intentions are. There are external codes of conduct. An early one comes out of something called the CARIN Alliance, and it gives us an attestation that is enforceable later, by the FTC, as to what that company will do with data collected by the app.

The next slide shows that there was opposition to this rule on the basis of multiple special interests. I strongly supported the rule publicly, but I have to agree with one of the points that was made in the opposition to the rule. And the rule was passed over this opposition, and I'm going to talk about some approaches that we're taking to address the point.

The point is that when data traverses that API, it loses, potentially, a lot of protection. And the opportunity here is to enable the FTC to handle the proper stewardship of those data. I addressed some of these points about the privacy of data once it has traversed the API and lost the HIPAA protections, in The New England Journal, around what do we need to do to be data citizens in the 21st century?

We have to be very cognizant that there will be, as an exception to the rule, I'm sure, but nonetheless, predatory app companies. We may have multiple forces, partially driven by privacy concerns, where we don't get the market economy of apps competing with each other and adding value to the health system. If we're not careful about the security, we'll have abuses and breaches, which will lose confidence. And also, we must be very careful about widening the digital divide when we deal with these technologies and when we deal with people's attitudes towards privacy, which may, actually, vary across this digital divide.

So I want to talk briefly about how there is a stop-gap technical fix that is enabled while we think further about how to strengthen the FTC's role. And what that is -- and what we should do now is go to

the slide that shows the SMART app privacy manifest, which is a couple of slides down.

And the opportunity here is the following: The API provisions were accompanied by very strong regulations against information blocking, so that a health system cannot prevent a patient from choosing an app that they wish to connect to their electronic health record. An electronic health record vendor cannot prevent a patient from connecting an app. Overall, that's very good, because it gives patients agency, and it gives app developers and innovators the opportunity to have a large market.

The problem is that it could be perceived of as information blocking, just to tell patients and warn them about bad apps because bad apps may be in the eye of the beholder. And so the Office of the National Coordinator, in the regulation, actually addressed this with a potentially innovative solution. And that is that in the OAuth process that enables the authorization and authentication of the user and the app to the electronic health record, there is an opportunity to present the manifest of privacy policies. And, in fact, some of the electronic medical record companies have begun to do this.

And so there is, specifically regulated, an

approach that this will not be information blocking if basic information is provided. What kind of information could we provide; the location of the privacy policy, the data storage policy, the data usage policy, the data-sharing policy; who made the app developers send data to and for what purpose; what relevant data; the apps method for approaching patients before sharing their data with other parties, as we heard about from Quinn; and we can also put in trust entities badges if the apps have actually attested to certain practices.

However, what we may also want to be sure that we do is to also -- and we can go to the last slide, which is this timeline -- is make sure that this decade of work that has gone into liberating information from electronic health records to empower consumers and provide them with computable copies of data actually results in a safe ecosystem. Part of this is defining what the privacy policies are and making sure, perhaps even from a regulatory perspective, that those elements are there.

Research is needed on how patients understand those privacy policies, and I believe the FTC could have a strengthened role in enforcement of those policies, as well, to make sure that when there

are breaches of what is promised, that there is a strong enforcement reaction. And it's very critical to protect consumers from harms related to health data. And if we can make consumers feel safe in this environment, I think the opportunity is almost unlimited.

Thank you very much.

MS. CONNELLY: Thank you so much, Ken.

Dena, you're up next.

MS. MENDELSOHN: Hi. My name is Dena Mendelsohn. I'm the Director of Health Policy and Data Governance at Elektra Labs. We offer services to better evaluate and dispense connected health-monitoring technology, many of which feed into the health apps that you're hearing about today.

Prior to joining Elektra earlier this year, I served as senior policy counsel at Consumer Reports, where one of my most recent projects was reviewing the data practices and security of a handful of reproductive health apps.

Today, I will discuss a paper published by my colleague, titled "Modernizing and Designing Evaluation Frameworks for Connected Sensor Technologies in Medicine."

Next slide, please.

In today's presentation, I'm shifting gears slightly from the preceding speakers and will pan out to consider the ecosystem that feeds into and works with health apps. I will give you a broad overview of why clinicians are increasingly using biometric monitoring technologies and what type of due diligence we recommend before adopting this remote monitoring technology. I will conclude with what we recommend to simplify the decision process. Sneak preview, it's a label, somewhat akin to a nutrition label that we're all familiar with.

Next slide, please.

But, first, let's talk about why. Why collect digital measurements in real time at home? Well, the simple answer is that in research and care, remote sensing offers a more holistic view of a person's lived experience, especially when we're looking at chronic conditions that impact a person's daily life. Do we want to just know how they're doing through a few status points throughout their day? Well, not really when there's a better alternative, where we know how they're doing continuously throughout the day and over a longer period of time.

So while it would be simple to just step away from health apps, for those who are concerned

about their data rights, that really takes away some very powerful tools for them, and so it's not what I think any of us would recommend.

Next slide, please.

We believed in the value of the remote health-monitoring technology before COVID-19 took over, but the value of these technologies is even more clear during this difficult time. Uptake of remote monitoring technology, like connected sensors, are likely to rapidly increase during this pandemic, especially following guidance from the FDA and CMS that encourage widespread use. I think we've all seen a lot of articles about this in the lay press. Yet, public discussions of the risk of these technologies has been limited.

Next slide, please.

We should be at the Due Diligence is Necessary slide. And this is where, in our paper, we provide a deep dive into the due diligence that is critical when selecting connected sensor technology, whether it feeds into a health app or not.

Next slide, please.

What you're seeing here is a broad overview of our five-point holistic framework for balancing the benefits and risks of adopting connected health

technology. Again, many of this technology feeds into the health apps that we're talking about in this panel. The first three dimensions evaluate the data and subsequent results generated by connected biometric monitoring products.

The fourth dimension, utility and usability, evaluates the ease of implementation and adoption of the product. And the last dimension, economic feasibility, has the reader consider the cost and the value of adoption. As explained in the paper, evaluations should be multidimensional and a single score should be avoided.

Next slide, please.

So on this slide, we're looking at step one of the evaluation framework. And this is less about health apps and more about ensuring that the technology that's being used will generate information about a user that is suitable, both in terms of what measurements are made, the accuracy, and the appropriateness in the situation where it will be implemented.

Next slide, please.

As discussed in the paper, suitable technology must be verified and validated. Simply put, the technology must be accurate, both in the

measurements it makes, as well as any algorithms that it applies to the collected data, and that the technology were for a specific use case in mind. After all, not all technology is appropriate in all contexts.

Next slide, please.

The second part of the evaluation framework in this paper considers security.

Next slide.

On the Cybersecurity Considerations slide, the paper recommends including whether the company has a coordinated vulnerability disclosure policy and what's in it; does the organization publish its security support lifetime and issue secure, prompt, and agile software updates once security issues are discovered; and, finally, does the organization track and share a Software Bill of Materials.

Next slide, please.

A third component, and probably of a special interest to viewers today, is to look under the hood of data rights and governance. Given that you're streaming PrivacyCon, you probably know why data rights are an important safety tool for users of technology. When it comes to technology involved in health and healthcare, individuals' right to data

governance is pretty uncertain.

Next slide, please.

As it is, in our healthcare system, we have strong protections for patient bio specimens, like blood or genomic data, but protections are murkier for digital specimens. The same can be said of data created by health apps. Make no mistake, wearables, health apps, and in-home sensors offer great promise for affordable, accessible, equitable, high-quality care. But in the modern era, data rights have become a safety issue that extends beyond the body. The digital health data that folks generate may threaten both their health and their financial welfare, which you're hearing a lot about today.

Next slide, please.

We've seen enough headlines to know that there's a problem with how data is collected, used, and shared.

Next slide, please. Shaky data rights in the United States means that when clinicians recommend some health technologies to their patients, or a friend recommends it to another friend, they could be unwittingly putting the individual at risk. That's why the third part of the evaluation framework asks these foundational questions about the data practices

of technology under consideration. As explained in the paper, there could be gradations in manufacturer data practices.

In our evaluation framework, the minimum threshold is that the manufacturer has a EULA or terms of service and privacy policies that are publicly accessible online. But, really, we know that that's just a baseline. It's also important that documents are comprehensible or understandable by a broad audience. And at the end of the day, being fully transparent about practices is not the final solution. Transparency is not the solution, but, rather, manufacturer and app developers need to commit to privacy-protective practices. As we explained in the paper, the highest quality data practices means that the EULA and terms of service do not contain exculpatory language.

There should also be an opt-in or opt-out of third-party transfer or use of data, where appropriate. And, ideally, these rights should remain unchanged, even in the case of a change in ownership of the connected technology or the sensor manufacturer.

Next slide, please.

Finally, parts 4 and 5 of our framework

consider whether a product has features that users need and whether it's designed in a way that folks will actually want to use it. And, finally, no evaluation will be completed without the consideration of the cost and value of the technology.

Next slide, please.

We should be looking at the Nutrition Label slide. Now that I've considered the holistic evaluation framework, I'll remind you that excellence in one dimension does not necessarily imply excellence in another. Indeed, significant deficiencies in any one dimension may lead to problems when using connected sensor technologies in research or in practice. Thus, we propose a framework that simplifies the evaluation process of connected sensor technologies for the intended use, but it does not give an individual score that would make a decision for the reader.

As remote health-monitoring technologies become increasingly commonplace, more and more people need to decide the risk/benefit type of evaluation that we explained in the paper. But this analysis will need to be more straightforward. As the paper concludes, they propose that a connected sensor technology label could be a useful piece of

infrastructure for an evaluation framework, which would make it easier for decision-makers to understand critical aspects of technology in a streamlined and accessible format.

It's extremely likely that remote health-monitoring technologies, paired with health apps and some connected in other ways, will become a very common thread in how individuals manage their own health, how healthcare is provided, and in the context of biomedical research.

I would encourage viewers to read the paper that I discussed today to get a deeper understanding of the features of connected sensor technologies and their benefits and risks and how they should be evaluated ahead of deployment. If viewers from the healthcare sector are interested in learning more about digital medicine to enhance public health, I would encourage them to check out the Digital Medicine Society, or DiMe, which is a professional society for digital medicine.

I also want to acknowledge the authors of this paper, my colleagues Andy Coravos, as well as Megan Doerr, Jennifer Goldsack, Christine Manta, Mark Shervey, Beau Woods, and Bill Wood. I also want to thank the FTC for inviting me to speak today and for

its efforts in moving PrivacyCon online this year.

Thank you.

MS. CONNELLY: Thank you so much, Dena, for that really interesting presentation.

And now we'll move to our final presenters for this part of the panel. Our final presenters are John and Sarah.

So John and Sarah, I'll turn it over to you.

DR. TOROUS: Oh, thank you for having us, and as going forth, I think you'll hear some themes that are repeating and some parts that are new.

But we'll start with the first slide. We'll see if it gets pulled up, Actionable App Evaluation. And let's see, is it up? I think it's not up yet.

MS. CONNELLY: We're experiencing a little bit of a time delay with certain browsers on the slides. So if you could maybe just start off and, hopefully, they'll catch up pretty quickly.

DR. TOROUS: So as I said, we'll talk about actionable health app evaluations. And, first, we want to thank our donor, the Argosy Foundation, which made this work possible. We couldn't really have done any of this without their support. And I think what we're talking about today - and I think Sarah and I are coming from an interesting position, where we're

doing clinical research, but we're also delivering clinical care. So we're looking at how these apps work in real world settings and how policies really impact care decisions and patients today on the ground.

And we know from experience there's many good smartphone health apps and wearables that can improve care. As we've heard about from other speakers, there's also some pretty concerning dangerous ones that can directly harm care, threaten care, or harm the whole field. And we know, again, that a lot of these healthcare apps wearables are pretty clever in that they call themselves "health and wellness devices." They don't really go under the medical category, so they work hard to kind of avoid different types of regulation.

So looking at the slides of privacy concerns, again, we know that many of these things live outside of HIPAA and other kind of privacy laws. And we know that when a lot of patients come to see us, they actually expect, when they go on to the commercial marketplaces and download an app or a wearable, that if it's related to health and they see things about health, they intuitively expect that it's going to offer health protection. So do many of our

physicians, therapists, psychologists, social work colleagues, as well, and nurse practitioner nurses, and, again, that kind of set a line between how is it regulated, where is the data going.

And, again, on the Privacy Concerns slide, you can see the same thing that Quinn Grundy presented, in that you don't always know where your data is going. And on the second Privacy Concerns slide, you can see our team did a paper last year, where we actually did something called a "man in the middle" attack, and looked at where was data from popular mental health apps, popular apps for depression and smoking, if you downloaded them, where was your data going?

And the trick was we actually did read those long, complex privacy policies, and what we've found is even if the privacy policy promised you and pinky swore that your data was really going to stay secure and safe and it wasn't going to go anywhere, it kind of still went somewhere. Often, it went to Facebook Analytics, among other sources. So even if the app developers did have a privacy policy, sometimes it wasn't actually followed as well, which was pretty concerning. And that slide, you can see The Washington Post covered the article saying, "Smoking

and Depression Apps are Selling Your Data," which was a little bit concerning.

And, certainly, these privacy concerns we've heard are still with us today. This is just a headline from February 2020, so not that long ago, about a popular therapy app that's disclosing different aspects of users' data. I think in mental health, we're in a unique position, that a lot of digital health actually focuses on mental health because we can both collect data from sensors and apps that informs care. And in mental health, we can also offer people treatments via videos and technology. So a lot of this is actually happening in the mental health space, and privacy concerns have actually shown up a lot in the mental health space, as well as other spaces as well.

So you can see on this slide that says, "Exaggerated Claims of Effectiveness," in a different study with a group led by the Black Dog Institute in Australia, we actually read the app stores to say what are these apps claiming. If I'm a patient, I'm a clinician, I'm a physician, I'm an NP and I'm looking at these apps, if you read the app store claims, that they really kind of -- over half of them make claims that could be seen as medical, implying effectiveness.

We actually went back and tried to tie it down to what is actually claimed in the literature, what is actually proven. And really, it's less than 2 percent. So there's a huge dichotomy between what a consumer is seeing and what is actually supported. And I think there's different consequences, we've heard different speakers, to this misinformation.

On the Perils of Misinformation slide, one really concerning aspect we saw was that a lot of mental health apps just aren't updated. The developers aren't keeping them current. And some of these apps are offering incorrect suicide hotlines. And I think the quote speaks for itself, "Nonexistent or inaccurate suicide crisis helpline phone numbers were provided by mental health apps, downloaded more than 2 million times." So again, I don't think anyone's trying to give incorrect or false information, but, again, sometimes these things are just not really able to live up to the goals and standards that they would want to.

I think a lot of times the way that people find apps, be it, again, colleagues that would work in the hospital, patients, the people we talk to, is they look at, well, what's the top out there in terms of the search, or which one has five stars, or which one

has over 100,000 downloads. And that's not always the best approach to do it. If you type in schizophrenia, this app that's really a pawn game shows up. It's stigmatizing. It's incorrect. It actually doesn't work on a lot of phones, and that's probably a good thing. But, again, just because it shows up highly in a commercial marketplace really isn't going to tell you a lot about the app.

Prior research, on the left, clusters, really shows you that even apps that have high star ratings, it doesn't really tell you much about their clinical utility or validity. And this was more than mental health. This looked at apps and diabetes, heart disease, as well as depression. And that kind of hockey stick graph, that sharp decline on the right, where the slides with stars and download metrics are misleading, shows you that, really, the average person who downloads one of these apps, they don't actually use - about 95 percent of people aren't going to be using it after two weeks. You see engagement really drops off.

So even if the app is highly downloaded, the real question is, can people actually stick to it? You can't really learn that from metrics.

So I'm going to have Sarah take over on this

slide that says, "Deriving a Practical App Evaluation Framework."

MS. LAGAN: So in light of these concerns, we're now going to briefly discuss our efforts with the American Psychiatric Association to develop a framework specifically for the assessment of mental health apps, but applicable to health apps broadly as well.

So on the next slide, you'll see how there are numerous app evaluation schemes. So there's a clear need for an evaluation system beyond app store metrics, as we saw with these many concerns. And to respond to this need, there have been numerous app evaluation frameworks that have emerged, including the NHS in England, Denmark's MindApps system, and over 45, as of 2018, with far more emerging in the two years since then.

So if we go to the next slide, the Potential for Harm with Lists and Static Ratings, many of these frameworks rely on lists or static ratings, which may fail to account for nuance in diverse app needs. Just as there's no A-plus medication or talk therapy, people react to and use apps differently. Even the same app may be used in different ways, depending on individual variation and preference and needs.

Further, the app market is constantly changing and very dynamic, and it's hard to know if these lists respond to the most current version of the app.

So if we go to the next slide, what we did was we looked at 45 different frameworks, back in 2018, and we sorted the 604 unique questions from those frameworks into categories. So as you can see on this graph on the right, short-term usability questions were highly overrepresented compared to questions regarding privacy. The privacy questions are the ones in pink. So you can see on the right-hand graph how usability questions were just far more predominant, even despite the privacy concerns that you've heard raised throughout the presentation today.

If you go to the next slide, we use these questions to inform the framework we created with the APA a few years back. And as you can see this pyramid graph here, there are five levels, Accessibility, Privacy, Clinical Foundation, Engagement, and Therapeutic Goal. Corresponding to each of these categories is an ethical principle. So our framework is really grounded in the ethics that guide care.

And in the years since it has emerged, if you go to the next slide, we'll see how it stacks up

really well on privacy questions specifically. So this recent scoping review of different evaluation systems for apps featured the APA model. And you can see, highlighted on the left-hand side, how the APA model they found to be extremely thorough in addressing the various components of privacy. So what data is being collected; to whom is it shared. And on the right-hand side, you'll see how the app has been pretty widely cited and mentioned in the literature since 2016.

So then we can go to the next slide, a Framework to be Customized and Adapted. It's been referenced in numerous different papers, highlighting its adaptability beyond just mental health apps and towards health apps more prominently. So our next question this year was, how could we use this framework and make it even more actionable for consumers, clinicians, patients, and any user of apps?

DR. TOROUS: So what we wanted to work on was saying, well, we've built these principles, we've kind of guided people on what to look for, but that can put some more onus on the patient, on the clinician in the visit. And we talked about, so how can we make it easier for people to understand this. And one of the problems was, again, a lot of app-

rating systems will say, is it easy to use? But, again, what does ease of use mean? Who is it for? And if we say an app is easy to use, really that's putting a value judgment on different people trying to say what it is.

So we broke down things like ease of use into things like engagement style. Does it have peer support? Is it AI driven? Does it have videos? Does it have gamification? The idea isn't to judge it, but we wanted to make our criteria with different elements that could be objectively reproduced as kind of yes, no, or numbers. So we kind of smushed that APA pyramid into over 100 questions, which are more objective, to help people understand what an app could or couldn't do, what it offers. And, again, the goal isn't to offer judgment; it's just to say what features or what elements does it have or not have?

So the idea is to build a system powered by the community. This was a theoretical model that we published last year. And, right now, we're kind of at the A. We're looking at clinicians and patients using it, giving us feedback. And the goal is to get more towards B, where we do get app developers involved as well.

But you can actually see our project live

today at Apps.DigitalPsych.org, and you can actually use it. I think what I'm going to show you guys is this is what the broad database looks like. Again, we try to be fully transparent, so this is a screenshot of the website. You can see what it looks like today. But the idea is you can imagine this screenshot app should be very, very wide because it's going over 105 questions of different apps. And people can sort them and people can say, hey, what are all the apps, again, if we look at Spanish, that have really great privacy features, and we'll show those. Or someone may say, I don't actually care about privacy and I want to find all the apps that have video, and I don't care.

And the idea is we want to make people aware. We want to make sure people make informed choices, but we don't want to force what people's choices are. As other speakers have said, we want to make sure people can pick what that nutrition label is. They have to be aware of all the information, but we're not here to say, this is the best one for you. Someone may say, look, it's very important my app have text messaging and that's the most important thing and other features don't matter. So people can easily search our database and learn about what features are in an app.

There was recently an article in STAT News last week kind of showing how we can help people make apps. But the real question is, is there something that can actually help change clinical decision-making, change impact? And one thing we can do is because we have this database, we can query it. So one question we can ask of the apps we've looked at, do apps support downloads, do they offer more privacy features? As it showed a code in that line, but the answer was no. In contrast to last year's PrivacyCon with the apps we've looked at so far, we said, do apps that cost more offer more privacy features? If you pay more, do you get better privacy? And, again, from the subset of apps we've looked at, the answer was no.

We can also use this to help patients make smarter decisions. We can do patients with training. This was an app we don't endorse or not endorse any app. But before, we asked a group of patients, would you be interested in downloading this app? And, basically, it was 50/50. And after we had patients use the tool and ask questions, you can see that their decision-making changed. People said no.

We can also do this with clinicians. Again, just an example. Blue was before and then orange was

after. You can see we took a lot of clinicians who were in that three middle range. Someone said, hey, I'm not as interested in this app now. So it's possible to quickly let people search for apps, learn, and change how they're making decisions. So we're expanding on those.

And I'd say that, certainly, I think clinicians and patients both are pretty excited to learn about this stuff, they just don't always consider it because they think that these protections are inherent.

So we'll close by, again, thanking our donor who made this work possible and the FTC for inviting us.

MS. CONNELLY: Thank you so much, John and Sarah, and thank you to all the other panelists as well.

We'd like to move on now to our Q&A portion of the panel and, hopefully, engage in some good discussion, expanding upon some of the ideas that you've mentioned and maybe touching on some new ideas.

So I'll start us off, and I'd like to start with a question or two that are probably at the top of everyone's mind these days, and these are questions related to the pandemic. So as you've probably seen,

there have been a multitude of recent news articles regarding, for instance, a pandemic-induced mental health crisis in the US and a significant increase in consumer demand for things like therapy apps during the pandemic.

Are there practical steps that a consumer can take now to protect her privacy while also obtaining useful health-related supportive services? So, John and Sarah, you've, in particular, touched on some of these issues in your presentation, so I'd like to maybe start the discussion by giving you a chance to expand upon this particular part of your work, and then I'll move on to the other panelists.

DR. TOROUS: It's a very good question. In the pandemic, as people are looking for more mental health resources, how can they find useful ones and not end up, perhaps, trading all their information in?

We've seen this clinically in the patients that we support. People do want extra help. And I think basically, what we always do with people first is we [indiscernible] check for a privacy policy. You'd be surprised how many apps don't even offer the level one that Dena talked about even a privacy policy.

But usually what we actually do with

patients is we look at how much information the app may be wanting, if it wants GPS for different levels. And then what we do is say, what is the risk/benefit? Usually, by talking with patients, people, again, are usually shocked that the app is collecting this much data, but then sometimes, oftentimes, they say it's not worth the benefit, but it is. But I think as long as people are informed and aware, that's a very good first step and people kind of realize the risk/benefit and going through that.

Usually, as people bring apps to us, we're adding them to our database and then going over it with patients, and sometimes we use our database. If an app doesn't come up with a good match, patients will say, well, what if I was willing to compromise on this or if I wanted more privacy? So usually, we have a discussion around it and it turns out to be, I would almost say, therapeutic and informative for all parties.

MS. CONNELLY: Thank you. I'd like to see if anyone else has anything to add. Maybe Quinn, do you have anything to add? Or Dena?

MS. MENDELSOHN: Yeah, I'll just add in the first step is making sure that individuals understand that, in many cases, HIPAA doesn't apply. So as

speakers said a few times today, there seems to be some misunderstanding or assumption that when we're talking about health, that all health is protected the same, and it's just simply not.

And then going from there, just reminding consumers that health apps is a very large market. So there are choices. It's not that you always have to give up your data. You need to be careful about picking which one you're going to go with and just be intentional about your selection, rather than simply downloading the most popular app or the one that one person may have recommended.

MS. CONNELLY: Thank you, Dena.

Quinn or Ken, do you have anything to add?

DS. GRUNDY: Yeah, I might offer a slightly different perspective. I think the pandemic has laid bare, in many areas of our lives, preexisting problems and really exacerbated them. And so I think this is a great example where there's actually maybe greater awareness around privacy and security of data than ever before. And I think what that will hopefully lead to is some collective demand that there be better protections.

And I can't really think of another consumer sector or industry or product where the same amount of

responsibility is placed on consumers for ensuring that products are safe to use. And I think as we learn more and more about the consequences of lack of privacy or privacy breaches, that hopefully, we will see some better regulation.

And an example would be there's no regulation, for example, placed on the app stores or app distributors to ensure that the products they market are safe for use, and we don't see that in other sectors.

So while I think there are some practical steps and consumers are in a position where they have to make choices for themselves, I don't think that, ultimately, it should be a consumer's responsibility to make sure that products are safe and private.

MS. CONNELLY: Ken, did you have anything to add?

DR. MANDL: I'll just add that, yeah, it's definitely the Wild West. I think one thing a consumer can do is to look for endorsements by professional organizations that they trust. Hopefully, those professional organizations are educated on the issues we're talking about today, enough to know what to endorse. It won't always be the case.

And the other caveat, unfortunately -- because I'm sure many of these apps are very useful -- is that privacy policies and terms of use can change, including for the data that you've already contributed. And so I think we really do need stronger protections going forward so that consumers can take advantage of this emerging apps economy.

One advantage in these API-based apps, where we have the transition that I talked about from a HIPAA-covered entity to the FTC regulation, is there, we really know what the data going in are and we have the opportunity to regulate those data as they go into FTC jurisdiction. With a mental health app, where it's really health-related but not coming from the health system, I think the oversight of those is even more complex. As complex as it is to regulate the health API-based apps, regulating apps that provide a health benefit is, I think, even more complex, but comprehensive legislation is probably what we need.

MS. CONNELLY: Okay. Dena, I see a hand raised, and I saw that John and Sarah did a lot of head nodding, so I'll give you another chance after Dena.

MS. MENDELSOHN: All right, thank you.

Yeah, I just wanted to thank Quinn and Ken

for bringing that up. In the immediate short term, we are not getting any privacy laws passed in the next short term, couple months, and so individuals do need to be very savvy in the marketplace. But like everyone else is saying, it does seem quite inappropriate to shift the burden to consumers to do a lot of homework, and it really makes an assumption that consumers are in a position to always protect themselves, when really that is not the case.

Another concern that I also have is that when we tell people to rely fully on privacy policies, we're basically putting developers and manufacturers in the position of creating their own laws and then following them. And then we're expecting the FTC to be able to enforce on every individual law, which also does not seem reasonable at this point.

So looking forward, what we definitely need is for lawmakers to promulgate comprehensive data protection for individuals.

MS. CONNELLY: Thank you.

John and Sarah?

DR. TOROUS: We'll agree, even from the study we presented, where we showed that the apps aren't really even following their own privacy policies. But I wonder if, as laws and legislation

eventually take effect, there needs to almost be a focus on educating people to be aware of it, too. I think there may not be the demand for it because I think all of us tuned in and listening are aware of these issues.

But I think a lot of times the shock, when you show someone what data an app is taking, again, a clinician, a patient, it doesn't matter who, people actually don't expect that this much is happening or this type of data movement is happening. And again, I think it's because they say, well, when I'm in a clinic visit, I expect kind of privacy. This app is kind of talking about clinical things.

So I think raising even just awareness among people and educating them is probably a good first step. It's not comprehensive, but there aren't that many systematic efforts to do this. Or even clinicians don't have great resources to turn to learn about these issues. I think, again, it would be almost nice if we could force everyone to watch what is happening today. It would probably make a good first step in this.

MS. CONNELLY: Thank you so much.

I'd like to now change gears a bit and I'm going to throw it to Elisa, who's going to ask a

question about the Cures Act.

Elisa, I think you're on mute.

MS. JILLSON: Hi, can you all hear me now?

MS. CONNELLY: Yes.

MS. JILLSON: Yes. Okay, great.

So as Ken mentioned, following passage of the 21st Century Cures Act, the Department of Health and Human Services issued new rules intended to support patients' access to their electronic health information. Some observers believe that these new rules will significantly increase consumers' adoption of health apps, use of health apps, that are not covered by the HIPAA detailed privacy and security safeguards.

What are the implications of your research for the projected shift in how consumers use health apps? From a privacy perspective, how ready is the health app universe for this shift? And I guess my last question -- I know many of you have touched on policy implications and where more regulation or different regulation may be needed -- but coming back to the research, where is more research needed so that we are in a position to prod the app universe into the right direction?

DR. MANDL: This is a fantastic question.

From a utilitarian perspective, I have good news. The uptake of the apps economy innovation marketplace has been relatively low so far. That's for a couple of reasons. One is that the regulation is new and doesn't take full force until 2022. The other is that it's complicated to create these apps and to educate consumers that they even exist and to get them to use them. So from a technologist point of view, that's a big headache. From a privacy point of view, it gives us the advantage in that not that many people are being exposed to this risk yet.

The other aspect of the good news is that, by far, the most common consumer app that connects to this API is the Apple Health app. And to date, Apple, for its health app, has taken an extremely rigorous, privacy-first perspective. Apple does not mine the data. There is tremendous value in those data that are in those patients' health apps, and Apple leaves it encrypted, available only on the patient's device, and backed up, also encrypted, to the patient's or consumer's iCloud account. So it doesn't look across them. It leaves it with the consumer. And it has a special process, much more rigorous, than its process for general apps, for apps that will access the health data that has been downloaded to the patient's phone.

So the good news, again, is that uptake is slow, and where there is uptake, right now we have a lot of safety. But the issues and the caveats that we have seen throughout these talks are what we are facing not too long from now. And in addition to data that is going to be equally concerning, certainly the data that patients and consumers enter into mental health apps, is no less concerning than anything coming across that API.

Nonetheless, the data coming across those APIs will include, actually, clinical notes and summaries, eventually, hopefully potentially, images, things that are very revealing of many aspects of the patient. And I think we need to reinforce what happens as the data traverse those APIs with real standards for privacy policies and real means to enforce them, and tremendous education and research into how patients actually understand those policies and whether they can follow them and what the real risks are.

The other aspect, I think, is comprehensive privacy legislation so that, on the other end of this, these data that are health-related and health-relevant, are, in fact, in some way that the consumers are protected from the use of these data. And that's

going to take some real creativity, to come up with legislation that both promotes innovation and also protects patients.

MS. JILLSON: Thanks, Ken.

Do others have anything to add? Are there other areas where additional research is needed to make this app universe ready for us?

DR. TOROUS: I'll just briefly add, I think we still need to understand what both consumers and patients value in the data, what they are -- understanding kind of how people understand what their data is worth, what they're willing to trade, compromise. We're not telling people never share your data, but I think we still haven't, again, educated people on what it is, what they have, why it's valuable, when it matters, more than less.

I think, as Dr. Mandl says, the stakes kind of got higher. It's on us to make sure at least everyone is aware. We don't have to put the burden on them, but certainly they need to know what they have.

MS. CONNELLY: That mute button. Okay, I think we'll move on to another topic. And I'd like to make some linkages between at least one slide that John and Sarah had up at this conference and some research that was presented at PrivacyCon 2019. So

some observers of the app market have argued that you get what you pay for. Free apps sell your data to turn a profit. The paid apps are a bit more privacy-protective. Research that was presented at PrivacyCon 2019 challenged that idea, that paid apps are necessarily more privacy-protective than their free counterparts.

And so as I mentioned, John and Sarah, you had a slide on this that suggested some similar results from your analysis. I'd like to get some thoughts from all of the panelists about how does the free versus paid distinction play out in the health app context? And also your thoughts on whether additional research is needed here, and if so, what kind of research.

I'd like to start with Quinn for this, and then maybe move on to John, Sarah, and the others.

DS. GRUNDY: Sure. So I think, yeah, the work that John and Sarah and others have done obviously debunks the assumption that if you've paid for an app, your data will necessarily be private. I think one area that our research highlighted that maybe needs some more attention is the relationship between developers and third parties. In particular, there are a number of third-party services that are

used to monetize apps or to enhance the features of an app, whether that's user analytics or error testing or social media integration that are offered to developers in a freemium model.

So developers can access these services without cost and, often, that's in exchange for access to de-identified or aggregate user data. Often, developers who pay for higher tiers of service, sometimes there are different data-sharing agreements. The problem is that consumers have no way of knowing what kind of agreement developers have with third parties, what kind of data-sharing protections are in place, and the relationship between the user and the third party is far from transparent, and they actually, in many cases, have no relationship at all.

And so I think greater scrutiny and transparency with these behind-the-scenes relationships needs to occur so that consumers can understand what is ultimately happening with their data, whether not it has their name attached.

MS. CONNELLY: Thank you, Quinn.

John, Sarah?

DR. TOROUS: I think what Dr. Grundy said is exactly correct. I think the business model of apps is a different topic for a different day. But a lot

of these apps are moving towards subscription models, so it actually also becomes complex. So they'll have a free version that's kind of a limited trial or limited features, and then you kind of can pay to continue using it. So business models are evolving. And there aren't actually that many truly free apps, and the ones that are free are usually kind of just like information resources that don't really do much, not in a good or bad way.

But it's also interesting kind of as the business model of these apps evolve, how does the privacy around them? And when you pay for a subscription, what do you get or not get? I think that's a topic we have to learn a lot more about, as well as if the employer is paying for the benefit. There's a huge move, at least in mental health, to try to say the employer will pay for this. What does the employer have access to or not? So many open questions.

MS. CONNELLY: Thank you.

Dena, I'll give you the last word on this topic.

MS. MENDELSOHN: Okay. Well, I'll keep it brief, but I just wanted to push back on the notion that a paid app should have better privacy protections

than unpaid ones. This could create a major issue, where lower-income individuals are put in a position of picking between a free app that may not be as privacy-protective versus having to pay in order to get access to, perhaps, an essential service, like a mental health app.

And so this is yet another reason why we need comprehensive data rights set in law so that we have a baseline that everybody, regardless of income or ability to pay, can expect from their health apps.

MS. JILLSON: Thank you all for those thoughts. We have just a few moments left, so I'd like to ask if you all have any wrap-up thoughts. We had an audience question about what legislation is needed in this area. I think that's probably a question that would take more than one minute of wrap-up. But if you could briefly, in your closing remarks, address where you think research should be headed and, if you'd like to, where you think regulation or legislation should be headed as well.

And we can start -- Ken, why don't you start us off?

DR. MANDL: Well, I think I would focus some of the research on this transition across the API from a HIPAA-covered entity under consumer direction to a

third-party app. There we have a controlled environment and a regulatable environment. And getting that piece right will help consumers enormously in protecting their privacy and their integrity in the face of using apps and also in helping to prevent misuses of their data.

The research needs to be done in what patients expect at that moment, what they can understand, how much external protection they need, and where regulation versus sort of community standards becomes the most effective focus. But I'll emphasize that because the FTC could potentially be overseeing the regulation of a very large amount of health data for the first time, data that HHS is used to regulating, and the FTC is not yet used to regulating. I think we have an opportunity to really think this through together, as a community and as a nation, on how to make the FTC most effective in taking on this new role.

MS. JILLSON: Quinn, maybe we can go to you next.

DS. GRUNDY: I think at the moment our existing legislation regulation and the marketplace puts the most responsibility on the groups with the least power to do something about this, so consumers,

and to an extent, app developers. And I think the focus of regulation or legislation needs to shift to some of these really big players with much more power, including app stores and distributors, data aggregators and digital advertisers, who currently are very much behind the scenes and engaged in a lot of these sometimes dangerous and harmful practices but aren't really the topic of discussion at the moment.

MS. JILLSON: Dena?

MS. MENDELSOHN: I think at the end of the day, it's on our lawmakers to enact legislation that sets a data rights framework that could serve as a baseline for health apps and other connected technology. And that way, health app developers can focus on creating the best technology that can win in the marketplace and consumers could trust that the technology that they've chosen to further their health and their lives will not be used against them.

MS. JILLSON: Thanks.

And John and Sarah?

DR. TOROUS: It's hard to follow all of that up. So I think we would say perhaps we do need to start using and investing these frameworks in real world settings and actually, again, educating people, giving them resources they can use today.

On a more flippant note, if anyone has a name for the database that we've built, we'd love your help in naming it. Calling it the App Database is a little bit boring. So please send us any names you have. We're open to it.

MS. CONNELLY: Okay. And with that, we are over time. So I want to thank - Elisa, and I really want to thank all of our panelists for this really interesting discussion and great presentations. We appreciate it. We'll have a short break, and our next panel, which is Bias in AI Algorithms, will start at 10:50. Thank you all so much.

MS. MENDELSOHN: Thank you.

SESSION 2: BIAS IN AI ALGORITHMS

MR. ROSSEN: Good morning, everyone. My name is Ben Rossen, and I'm an attorney in the Division of Privacy and Identity Protection at the Federal Trade Commission. And it's my pleasure to welcome you all to our second panel of the day.

Today's PrivacyCon is primarily focused on the privacy of health information and mobile apps, but this panel has a little bit of a broader focus on what is a very important issue and, surprisingly, is one that we haven't covered in a previous PrivacyCon. Namely, that is algorithmic bias and the risks of data discrimination.

So we are extremely lucky today to have two really terrific panelists. First up, we're going to have Muhammad Ali. He is a PhD candidate at Northeastern University, and he's going to be presenting his paper entitled "Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes."

Next, we are very lucky to have Professor Ziad Obermeyer from UC Berkeley's School of Public Health, and he's going to be presenting his influential paper about bias in managed healthcare algorithms, entitled "Dissecting Racial Bias in an

Algorithm Used to Manage the Health of Populations." And you could find their full bios on the event website.

We're going to have two 12- to 15-minute presentations, after which there will be an opportunity for some Q&A.

And with no further ado, I'm going to turn it over to our first panelist. So, Ali, I'll let you take it from here.

MR. ALI: Thank you. Are my slides online right now? Okay, I hope they are.

Well, thank you so much for the introduction and thanks to everyone who is watching this. So today, I wanted to talk a little bit about discrimination in online advertising. And if you've been following the news, you've probably read an article or two about it. But a lot of the focus in the past was focused on the targeting side of things, how these online platforms are built in a way where they provide this breadth of options to advertisers, in essence, enabling them to exclude certain users from seeing their ads.

But I'm not going to talk about that. What I wanted to focus on was the delivery side of things, where once an ad starts running, the algorithm is

making decisions on who to show the ad to. So that will be the focus of this talk.

But before I talk about my results, I wanted to give a brief climate on what the Facebook advertising system looks like. That's what we focus on in this study.

Next slide, please.

So here you can see sort of -- if you have a Facebook account, you can go to the Create Ad option in the top right, and in a couple of clicks, you'll end up on this section. You see you can target by location. There's a bunch of demographic variables here, age, gender, language. And at the bottom there, you can see that there's detailed targeting. These are interests that Facebook is constantly inferring about its users, whether you're interested in coffee or comics, and then they present all of these attributes to advertisers to target. And that has been the focus of a lot of the prior work.

Next slide.

For example, these are some of the examples. On the top here, you can see, back in 2016, ProPublica showed that they could target people looking for housing and exclude people by their ethnic affinity, as Facebook was indexing at the time. And, later, it

showed that even if Facebook goes ahead and blocks these features from being excluded, as they later did, a malicious advertiser can go ahead and find other proxies that correlate with race, and then go ahead and exclude that. So there's a lot that a malicious advertiser can do here, but that's not the focus here.

Next.

So we sort of look at the advertising system in these two tables. There's the advertiser, who is controlling the targeting part, where they design the target audience, what the ad looks like, how much money they want to pay. But then once the ad is created, it goes to review. The advertising platform is making decisions on which user they want to show these ads to. And they're running an auction. They're doing some estimates of relevance. We want to understand whether the differences -- any sort of discrimination can arise in this space. So can there be delivery skews on this second phase?

Next.

And we do that simply by actually buying ads from Facebook, because there's no clear way -- there's no data set where you have information about targeting and then the eventual information about delivery. What we had to do, we had to create our own ads, sign

up as an advertiser on Facebook, and then ask them how those ads are doing. Facebook is happy to report breakdowns by age, gender, location, multiple other things. So we used the APIs to collect all this information on the ads that we ran ourselves. We thought this was the best way to do this.

Next slide.

And one of the first set of ads we found were these two extremely stereotypical ads that we expected would skew a certain way. So one is advertising bodybuilding and the other one is advertising a makeup kit, pointing to Elle or bodybuilding.com, both websites that we don't own. And we targeted these two ads to the exact same set of random phone numbers in the US to see, given that the targeting is the same, how does the delivery affect?

Next slide.

And we see that there's these large differences, where one ad has, eventually, 85 percent of male audience and the other just has 5 percent. So it's clear that just the targeting, just the delivery phase, can cause these large differences, regardless of the targeting.

Next slide.

So that's the first question that I asked on

how these differences can arise in the delivery phase, yes. But we want to understand it better. Like how do these differences even get there? Like what elements of the ad is Facebook looking at? Are these differences because users are clicking on these ads more? Does this decide a priority? I'm going to try to go through all of these one by one and see.

Next.

So this is what a standard ad on Facebook would look like when you are advertising a link. You can see there's so many things you change here. There's the text on top. There's the image. There's the URL. These are just the user-facing attributes. And behind the scenes, there's other attributes as well, such as the daily budget, what audience you're selecting. And we wanted to tweak each of these to see what causes the most difference.

Next slide.

And we realized that even before we changed any of the interfacing attributes, as I mentioned, just changing the budget itself causes differences in how many women see the ad. So we ran this ad for Indeed, the job search website site, from one of our pages. And we noticed that the more money we were paying, the higher fraction of women in the eventual

audience we were reaching, arguably because women are more competitive on Facebook or because they're more expensive for some reason. But these are differences that the advertiser would not be able to realize what's happening because -- so we sort of stick to a $20 budget for all of our experiments, so these baseline effects disappear.

Next slide.

And then we started to tweak the attributes of the ads themselves. So when we started running this, my expectation was they're running some sort of natural language processing and they look at the text that I put in the ad and that's how they decided who the ad is relevant to. Turns out I was wrong. We ran an ad with just a baseline, a white image, with a text on the white image. And we see that there's no differences between the bodybuilding and the cosmetic site. Adding the headline causes some differences, but not the sort that I mentioned earlier.

Next slide.

But adding the image immediately causes these large differences. We see that as soon as we add the image of the guy pumping iron and the bodybuilding, the initial skews that we saw just immediately replicate. So it seems like in these ads

we were running, the image is the strongest factor to the classification algorithm and its relevant estimate.

Next slide.

And one of those other things -- which this was also in our initial hypothesis -- it might be because people are clicking on these ads more or because people are interacting. But it turns out, we polled the API over the 24 hours multiple times, but it turns out that some sort of relevance estimate was made as soon as the ad started running and the platform sticks to a decision throughout the course of the ad. So there is clearly some initial decision being made.

Next slide.

And this was one of the harder things to measure, but we wanted to really be sure how much of this difference was because of any humans in the loop versus algorithms. By humans in the loop, I also mean users who might be giving telemetry data to Facebook, basically scrolling over my bodybuilding ads differently than cosmetics, or any sort of modulators just that might be in the loop.

So we wanted to create ads that would make no sense to people but would make sense to an image

computer vision algorithm. How we do that is we take images and we try to make them transparent. This is an example of that. You can see that this image looks slightly transparent. It's because -- you can see on the right there are RGB values for multiple pixels here. So each pixel has an RGB value, and then the alpha channel, which controls the transparency. And this is slightly transparent because I've turned down the alpha channel all the way to somewhere in the middle.

Next slide.

And if I was to turn the alpha channel all the way down close to zero, it would look basically a blank white square to a person. But the computer vision algorithm can take these RGB values and work with them. And it's funny because when I sent these slides to the organizers, they were confused. They said something is missing in these slides. But it's sort of built like a reverse CAPTCHA, where it doesn't make sense to a person, where it makes sense to a computer.

Next.

And we use this technique to basically take images where we knew, working with the algorithm, they were skewed towards men and images that we knew were

skewed towards women. And beyond both visible and invisible images, so that any sort of user interaction has gone away. It's just the image algorithm.

Next slide.

And you can see here, for example, the two blue-colored dots on the top. You can see the hollow ones are the ones where the male images were made invisible. Between the visible and invisible, there's barely any statistical significant difference. So the gender estimate, the gender skew remains the same, regardless of what it is. Because the user is seeing just a plain white square. It's not any sort of data that was being incorporated there. It's just that the image algorithm sees a certain image, it classifies it, and it sticks to its judgment.

Next.

So we went through all of these sort of to gain a better sense of how the algorithm is working. So we understand that it's mostly the image that's causing all of these differences. A lot of these differences are made as soon as the ad starts running, and humans are not as involved as we thought. And we say "at least" because we're not sure, because these ads aren't run for weeks or months. So we don't know what would happen if we got hundreds of clicks on

them. But at least in the few days that we ran these ads, we see that a lot of these decisions are algorithmic.

Next slide.

But one of the other things we really wanted to measure was whether Facebook is capable of producing any sort of racial skews, and Facebook wouldn't report us breakdowns as it does with the gender, where we can ask the APA for information. So to get at racial information, what we do is we take voter records from North Carolina. So we build this methodology where we divide the state of North Carolina into regions, where we only take information of black voters from the voter records and upload that to Facebook to create an audience, and regions where we only take information about white users.

So from the voter records, we can get information like first name, last name, zip code, and a lot of other things, and we can target these people. So when Facebook reports the location back to us, we know that we only uploaded black users in this area, so we can infer their race. And to test whether this works or not, we run yet another set of stereotypical ads.

Next slide, where we essentially take the

top 30 country albums, top 30 hip hop albums, all pointing to RollingStone.com, the same website, just different articles with images. And we see very, very strong skews, where the country music ad goes to 80 percent white users in the audience and the hip hop ad is only 12 percent white users and the rest of the audience is black. So this sort of gives us confidence that this reverse inference methodology that we come up with for measuring race works, and we can use this to measure these effects in more important categories.

Next slide.

And by what I mean by more important categories are protected categories, employment, where it's illegal to discriminate. So a lot of the examples that I showed so far, they might be benign. Judging whether someone likes sneakers or not doesn't seem too problematic, but doing the same thing excluding someone from an employment opportunity would create some sort of liability.

So what we do is we create these job ads on multiple ads. For example, this is a job in the lumber industry, the cleaning industry. All of these ads point to Indeed.com, actually job searches. So if someone clicks on it, they actually go to an actual

job search. And we target the exact same set of people for all of these ads.

Next slide.

And we see the same differences exist, even for these jobs ads that we saw earlier. For example, on the left, you can see the gender distribution. You can see that the number of job ads are close to 90 percent male, while the janitor ones are skewed towards women. And on the right, you can see the racial split, and you can see that the lumber jobs skew towards white people and the janitor actually skews slightly towards black users. Without the advertiser ever asking anyone to do so, this is the exact same set of people that both of these ads are targeting.

Next slide.

And we see not just in these two categories. We've done it for a variety of jobs, supermarket workers, secretaries, nurses. And you can see that across gender and race, there's so many differences that occur on the delivery side of things, even when the advertiser might not have intended to discriminate in any way.

Next slide.

So in essence to sort of summarize, what we

do is we provide these new methodologies to be able to measure Facebook's advertising system. And we show that regardless of how an advertiser decides to target, a lot of these differences can arise in the delivery phase. And not just in benign categories; it can also bring into protected categories, like employment.

So what are the real world implications for all of this? And I'd like to mention, last year, the Housing and Urban Development Department, they decided to sue Facebook because Facebook was enabling discrimination in housing opportunities. So our paper, we believe, sort of provides a way to investigate whether these differences -- how much of these differences arise from the delivery part versus how much of these differences are responsible by the algorithm itself, who's deciding who to show the ads to. So it's a methodology towards that.

We also think our paper sort of provides a unique nuance on the Communications Decency Act, Section 230. So this provides a lot of immunity to online publishers from all the content that they're hosting. So it's the responsibility of the people posting and not the publisher's. But what we show is that if so many of these decisions on which user

eventually ends up seeing something are contingent on the delivery algorithms, on the AI that's running in these systems, then it's not so clear then who's entirely responsible.

And, finally, I'd like to emphasize that we're still at the phase where we need more transparency into these systems. Whenever something goes wrong, online advertisers cannot continue to blame the advertisers for being discriminatory, when we clearly show that so many of these differences don't even depend on the advertising. A lot of these decisions are because these algorithms are optimizing so heavily for relevance that they might end up skewing these ads.

Next slide.

Yeah, that's all I have for today. I would like to profusely thank my collaborators, Piotr and Alan at Northeastern, Aleksandra at USC, and Aaron and Miranda at Upturn. And thank you, again, for listening.

MR. ROSSEN: Ali, thanks so much.

Next up, we have Professor Ziad Obermeyer. He's going to be presenting his paper, "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." Ziad, I'll turn it over to you.

DR. OBERMEYER: Thanks, Ben, and thank you so much, Ali.

I think that the work that Ali just presented was such an ingenious example of the kinds of ways that researchers have tried to essentially study algorithms in the wild. So if you think about all of the things that that research team had to do to kind of understand what exactly Facebook was doing, and in some ways, probably even better than Facebook understands what they're doing themselves, you know, it's this careful process of pinging the system, seeing what happens, reconstructing results. And all of this stuff is done, essentially, from the outside. Because in a lot of these settings, when we want to study algorithms that are operating at scale in our society, we can't get inside.

And we can't get inside for some reasons that are not so great, like the algorithm developers don't really want us to get inside. But also some reasons that are legitimate, that there are trade secrets and things that we legitimately don't want to make public.

And so I wanted to talk through one example from our work where we had an enormous luxury relative to most studies of algorithms, which is that because

we were working in collaboration with a health system that had actually purchased one of these algorithms, we could see everything about it. We could see all the variables going into it. We knew exactly what the algorithm was doing. And maybe, most importantly, for the purposes of making the case that there was racial bias, we could actually follow up what happened to patients and document the impact on health outcomes.

And so I think that this one example, or at least I hope, can teach us some general lessons about, essentially, how to be good users of algorithms. And that's on the consumer side, but also on the regulatory side as we try to make sure that bias doesn't get into these algorithms, and if it does, how to hold organizations accountable.

So our example that I'm going to tell you a little bit of background on up front is about our system's effort to help complex patients. So in general, our health system does a not-so-great job of helping people with complex health needs. They often end up in the emergency department or in the hospital, if they're on many medications that often conflict.

And so over the past few years, the health system has gotten very interested in trying to intervene early on these patients. And the idea is if

you imagine a person with heart failure, a person with diabetes, there's a window of opportunity to help that person early, when problems are still able to be nipped in the bud. And so what the health system has invested in very heavily is what's called high-risk care management programs to do exactly that.

So the idea is that these patients are treated like VIPs. So patients with chronic conditions are given a special phone number to call. There's a special team of nurses who can make home visits. They can arrange for a next-day primary care appointment. So it's really they want a low threshold for these patients to call in, reach for help so that this team of trained experts can nip all these problems in the bud.

And the goal is twofold. The goal is, of course, to help patients so that their health problems don't go from small problems to big problems, and the second goal is to save the health system the money that's associated with those problems turning into big problems, people ending up in the hospital. So as you can imagine, that SWAT team of specially trained nurses and extra primary care slots and home visits, all of that is fairly expensive. And so you can't do this for everyone. You have to choose your patients

carefully. And that's where algorithms come into this story.

So it's fundamentally about resource allocation. We have this scarce resource of extra help programs and we want to target those resources to the people who need it most. And if you think about most health systems, they're managing tens, if not hundreds of thousands, of patients. That's not a great job for humans to do. And so a lot of health systems have started investing in algorithms to at least start that screening process for them.

And so if you take the industry estimates seriously, the scale of this is just enormous. So the industry itself estimates that around 150 to 200 million people are screened by this family of algorithms every year. The particular software that we're using is one of the largest in that market, and so that's what we're studying.

And the way these algorithms are generally used is almost as a first step. So there's a primary care population. And the algorithm just runs in the background and generates a score for everyone in that population, and then the health system does something with that score. So in the particular decision that we're studying, the top few people were just

fast-tracked into this high-risk care management program, and about the top half, except that top 2 percent, 3 percent, those people were shown to their primary care doctor and the primary care doctors were asked, should this person be in this high-risk care management program?

So a lot of variety in the institutional practices, but, ultimately, the algorithm does a screening step, and then that screening is used to decide lots of things about the patient, but in this case, should that patient be enrolled in one of these programs?

So on the next slide, there's a graph. And I'm just going to talk through it slowly because I'll show you a few graphs that look like this and I just want to make sure they're all clear. So on this graph, on the X axis on the bottom, is the algorithm. So this is what the algorithm thinks about people, and it's arranged from very low risk on the left at zero to very high risk on the right. And those top few percent, to the right of that vertical dotted line, those are the people they get fast-tracked or autoidentified for this program.

On the Y axis is a measure of health. So this is basically at a given level of what the

algorithm thinks about you, how healthy do you end up being in the next year. Concretely, it's a count of how many chronic conditions you have that flare up over that year.

The two lines show two groups of patients. The top line, the purple line, is black patients, and the bottom line, in gold, is white patients. And as you can see at every point in this distribution, black patients, at the same score as white patients, have worse health, on average. And so I think that violates what you could think of as our working definition of bias.

So the algorithm is being used to guide a decision. And so two people who have the same algorithm score are treated the same by the algorithm and, thus, by the health system who uses the algorithm. So those patients should go on to have similar health needs, irrespective of the color of their skin.

And what we find is that if you just look at that high-risk group, where people are fast-tracked into the program, the algorithm, operating on its own, judges that high-risk group to be a group of patients that's only 18 percent black. When we did a very basic analysis to say what would this look like if the

algorithm had no bias based on need, that number would rise to almost half, to 47 percent black. So this is not a trivial amount of bias. And, again, the definition of bias that we're working with is at the same algorithm score, people should have the same needs, and that turns out not to be the case.

So on the next slide, what we're trying to illustrate is where we think this bias got in. As I mentioned, we knew exactly what this algorithm was doing, what it was predicting, how, what variables. And it turns out that if you step back -- this is a very complex question. Who has health needs? So in most data sets, we don't have a variable called "health needs." And so what we do instead is we pick a proxy variable that's measured in the data sets that we have access to.

And what the algorithm developers did in this case -- which is a very common choice; this is not just about this particular developer; this is a very common strategy -- is we used costs as proxy for health needs. Now, that's not unreasonable because, in general, when you're sick you go get care and you generate healthcare costs. The problem is that even though, on average, that relationship is true, that you generate costs when you need healthcare, that

relationship is very different for black patients and for white patients. So when you need healthcare, you're less likely to get it when you're black, and that leads to lower costs.

So in this graph, we're showing you on the X axis, instead of the algorithm, a measure of health. So increasing health needs further to the right. And what you see is that white patients always have more costs on average, no matter where you are in this health distribution. And in our sample, black patients cost a substantial amount less every year at the same level of health.

So on the next slide is our hypothesis of tying this all together. Using proxy measures is inevitable, but some proxy measures are biased, and we think this is a very common mechanism by which bias gets into algorithms. In our example, it was using cost as a proxy for health and not realizing that costs were just lower at a given level of health for black patients. But you can imagine many other situations like this.

We often use arrests or convictions as proxy for criminality, but that is not an unbiased measure of criminality. We use income to measure creditworthiness, and that's going to introduce all of

the biases we already know about and differences in income by ability. So all of these things, because they're subtle questions about correlations with underlying truth with race, they can be subtle, and that's why this wasn't caught.

It wasn't caught by the people who developed the algorithm, even though they were very well-intentioned. It wasn't caught by any of the clients that purchased the algorithm, even though these were people who have a deep commitment to fixing disparities and improving population health. And it wasn't caught by the humans who were either using the algorithm or being affected by it. And so that, what you can think of as a market failure, is the reason that I think there's an important role here for regulation. And so the question is how?

And so I'll just leave some of this to the discussion, but I'll just say that where anything starts, making sure that the algorithm that you develop or buy isn't biased, regulating and holding organizations accountable. All of this starts by having a very clean definition of what we mean by bias.

So in our case, it was two patients with the same risk score should have the same health needs

because this risk score is being allocated, is being used to allocate a health resource, and it shouldn't matter what color their skin is. That definition is the beginning of lots of methods that you can use to test for bias, to query algorithms that an organization is thinking about buying, and for regulators, to offer guidelines to industry.

And, critically, none of these things require compromising trade secrets. All of these things can be done from the outside. We don't need to understand or interpret the algorithm. All of these things can be done with our basic level of data access that we have.

So to wrap up, I'll just tell you that after seeing this work, we actually reached out to the company that developed this algorithm and we worked with them -- they were incredibly responsive and positive -- to replicate our results and their data and to patch their software by predicting a measure that was closer to health and not so close to cost. And when we did that, we saw really large reductions in bias. And I wanted to mention that because we've expanded this effort out to work with a number of different health systems, insurers, algorithm developers. And our email address is on that last

slide, if you want to reach out.

Thanks so much.

MR. ROSSEN: Great. Thank you so much. So we have an opportunity for some Q&A. And I know if the folks who are watching on the Livestream have questions, you can submit those by Twitter or otherwise, and we have somebody who's going to pass those along to the moderators.

Ziad, I'll start with you to just get that conversation started since you hinted at this a little bit already in your talk. What is the takeaway for developers and healthcare systems and regulators in terms of applying the lessons from your work as a practical matter? And are some of these applications already out there in the field given the work that you're doing with health providers and developers in light of your paper?

DR. OBERMEYER: Yeah. Thanks for asking. I'll say, first off, that it's such a treat, as an academic, that anyone in the real world is interested in your work, and so it's been a real privilege for us to work with people who are actually doing things in the world to try to understand and solve these problems.

I think it all really starts with coming up

with a working empirical definition of what bias looks like. And I think that a lot of the ways that we tend to do this in practice so far are we look at, is there a race-based adjustment? That doesn't guarantee that there's bias. The absence of a race-based adjustment does not guarantee that there's no bias. So I think really delving into the substance of what the algorithm is doing, what it's being used to do, and then coming up with a context-dependent definition of bias there that we can test empirically is the first step.

And so when we're working with these organizations, the first thing we do is we go really in-depth to understand, okay, here's what the algorithm is being used for. Here's the real thing that we're trying to get at. Here's what the algorithm actually does. And is there a difference there? So setting up a very clean definition of what bias is is the basis for software developers to audit their own products before they go into the field.

If you are purchasing an algorithm, you can set up queries to actually answer those questions. If you are a regulator, you can set up a definition for a given application, and then you can hold people accountable to it. So I think that's really the core

of what we did, and I think the work that Ali presented as well. It's really trying to translate the somewhat abstract notion of what bias means into an empirical data-driven definition in a particular data set.

And that's hard because there is no automated process that you can do for that. You actually need to really understand how the algorithm is being used and what disparate treatment or disparate impact would look like in this particular situation, and then set up a set of empirical tests following that.

MR. ROSSEN: Thank you. That's really interesting. And to sort of follow up from that, given that there is no off-the-shelf way of doing this, with these types of algorithms that are purchased from third-party developers, which is still, I think, the most common way that a lot of companies are getting their AI tools, is there a market failure there, in terms of who has the incentives or the obligations to really examine these types of algorithms and both the resources to look at, is it somewhere where regulation needs to set in or are there steps that your sort of ordinary companies are able to take to evaluate these risks?

DR. OBERMEYER: Yeah. I think, you know, empirically, at least in the case that we've studied, and I think in many others, there was a market failure because there was this problem that wasn't caught by anyone. I think the first part of fixing that is actually to put a name on it and to make it transparent that this is a problem. In all of my conversations with industry, I don't think there's a single software developer who wants to put out a biased algorithm.

And so a lot of them are already taking steps to do that internally, but I think because all of us are just learning about what bias looks like in different contexts and what it means, I don't think that there's a consensus definition on how you even do that if you you're the one that's developing the algorithm or if you're the one that's purchasing it. And so I do think that this is an area where regulatory guidance would be incredibly valuable. Because now that there's a lot of attention, there's a spotlight on these issues, nobody wants to be the company that is putting out an algorithm that someone later audits and finds to be biased.

So I think having regulators just set out a definition of what this looks like would be incredibly

valuable, because as in most things in medicine, prevention is much better than treatment. It's a lot easier and it saves a lot of pain on lots of different sides. And so I think having consensus around what that looks like would be really, really important.

MR. ROSSEN: I realized that I misspoke earlier when I said folks could submit questions on Twitter. The right way to do it is actually through email, privacycon@ftc.gov. So if you have questions, feel free to send an email that way and we'll pass them along here.

Ali, I'm going to turn to you as well. I thought your paper was really fascinating. One of the things that it reminded me of is a story from a few years ago about when Amazon tried to build a recruiting tool. There's only so many stories that are out there about algorithmic biases, but many of the reasons you mentioned, about not having that kind of window into how these things operate.

But they had caught this algorithm that was going to be used as a recruiting tool because they identified, before it was rolled out, that it was systematically discriminating against women, despite the intentions of the developers and despite every effort they made to try to fix that problem. And the

findings in your paper reminded me of this, because some of these issues, certainly with Facebook, have been identified in the past. As you mentioned, Housing and Urban Development brought a lawsuit, and folks have been looking at this issue with a pretty keen focus on platforms like Facebook.

Is there some reason to think that, regardless of their efforts, skewed ad delivery is just an inherent part of using these types of tools? And if so, what should the platforms be doing to respond, or is there a need for regulators to step in?

MR. ALI: Yeah. I mean, yeah, that's a loaded question. But I think the case of -- I like that you mentioned the Amazon case. The Facebook case is slightly different because their advertising tool is a one-size-fits-all thing. It's the same tool that's used for political ads, controlling democracy; the same tools used for selling sneakers, the same tool that's used for -- so what works in one context doesn't really work so perfectly in another. But for the Amazon case, it was very easy for them to test it in that very controlled case and see that.

But I think now Facebook has also started to make other tools for the housing and employment ads, where they're trying to actively address this because

so many people have brought forth these concerns. But these differences arise essentially because, as Ziad pointed out, there's always proxies for making up bad metrics. Because they're trying to optimize so heavily for relevance, what works in one context ends up hurting people in the other context. So I think the only way to go forward is to be cognizant that these algorithms actually have an effect on people and then measuring them.

You can only try to iterate on the measurement and trying to fix these things and realizing that the way these algorithms are designed -- because you're so heavily optimizing for some sort of machine-learning metric of loss or trying to accurately optimize some exact thing, it just ends up picking up more layers and hurting people in the process.

So, yeah, I think the only way to do that is to iterate on trying to fix it, and I think Facebook is only now starting to get into understanding that, okay, these things actually have harms. So they're now in that phase where the developers are actually trying to measure and counter these.

DR. OBERMEYER: I'll add one just interesting thing about the Amazon example that you

brought up, Ben, which echoes Ali's point, is that, in some ways, algorithms can actually serve as a very valuable role of exposing bias in humans. So what was the algorithm in that case doing? Well, it was predicting some variant of, is this person going to be invited back to be interviewed by us?

Now, as Ali mentioned, that's a proxy for the quality of the applicant. But when the algorithm spit out these predictions that were predominantly white and male, that actually was like holding up a mirror to the recruitment process, that was the bias. That was the source of bias to begin with.

So in a funny way, algorithms can actually work to expose these biases in the human processes that are used to train them, and I think that that's a kind of underrated contribution of algorithm. Everyone gets mad at the algorithm, but it's not the algorithm. It's us. It's just reflecting back what we're doing.

MR. ROSSEN: That actually leads me to a question that we received from the audience, which is for you, Ziad, which was about, what were the alternative proxies that you ended up looking at in your work, as opposed to costs? How did you choose them? And is that process of choosing unbiased

proxies something that is replicable?

DR. OBERMEYER: Yeah. It's a great question, and I think that it does go back to understanding exactly what we want the algorithm to be doing. So we want the algorithm to identify people in whom we can intervene early and make a difference. So from that point of view, it's actually not obvious that you want to be predicting total costs. Total cost brings together a bunch of things that you can think of as, like, good costs, like people taking insulin, which costs money, and bad costs, which are things like people getting their toe cut off because they didn't take their insulin, which also costs money.

So when you put those together into a total cost metric, you're conflating a bunch of things that are not the same. And so what we did is we came up with a metric of avoidable costs, so things like, you know, getting your toe cut off because you didn't take your insulin and not the insulin itself. We also have lots of different measures of health that are applicable to different populations, and some are not.

So it took a lot more work, kind of like just substance knowledge-intensive work to come up

with these. But I do think that in most of the data sets we use, there's a rich set of alternatives. Some are more work than others, but I think the message from our work is that that extra effort can be hugely valuable because it can make the difference between a biased algorithm and one that actually works against the structural biases in our society.

DR. ALI: I'd like to go back and talk about that because it's very interesting what Ziad said about how these algorithms sort of hold a mirror to us and tell us how we're being biased. I really like that argument, but I hate when computer scientists use that argument to just evade all sort of responsibility.

I think a very common thing that computer scientists do is that, oh, the algorithm isn't biased, it's the data that's biased. But I think it's that very point where -- as someone who's trained as a computer scientist, who's been in way too many machine-learning classes, it's important to understand that just because the data is biased doesn't mean you let the thing go through. It's that very opportunity where the algorithm's holding a mirror to you to understand that you're now automating this harm that was accumulated over years. And that's where you need

to start auditing these systems.

As Ziad said, you need to have clean definitions of bias and work with those until you reduce that harm.

DR. OBERMEYER: I think that's a great point. I think there is a tendency to throw up our hands and say, well, we can't have algorithms because the data are biased and the data are biased because our society is biased. And all of that is true. But with a lot of work to take into account structural biases and historical inequalities, we can actually make the difference between good algorithms and bad algorithms.

MR. ROSSEN: I know we're running out of time, and I think that's a great place to end the conversation. I know our next panel picks up immediately after this one. I want to give a big thank you to both of our panelists. Really fascinating work. The papers are available on the FTC.gov website. And thank you so much for having us.

DR. OBERMEYER: Thank you.

SESSION 3: THE INTERNET OF THINGS

MS. ROUGE: Hi. So this is Phoebe Rouge, and today, for our third panel, we're going to be talking about privacy and the internet of things. We have three presenters here.

The first one we're going to have, Daniel, who did his research at Northeastern University, and he's going to be talking about his research to look at the network traffic from various internet of things devices.

DR. DUBOIS: Thank you for the introduction. Yeah, so now I will talk about information exposure from consumer IoT devices. And I will also thank my collaborator, Northeastern University, Jingjing Ren and David Choffnes, and from Imperial College London, Anna Maria Mandalari, Roman Kolcun, and Hamed Haddadi.

Next slide.

Usually, I start the presentation by asking the audience if they have any IoT device. Typically, the majority says no, but, actually, most of them then realize that they bought a TV in the last 10 years. That TV is likely a smart TV and that can be connected to the internet and that's a full IoT device.

So what motivates this work is that IoT devices have access to private information. They have

the sensors. For example, smart speakers can listen to you, like a smart camera, smart doorbells, and watch you because they have a camera. And smart TV knows what you do, for example, what TV programs you watch. So all this information is actually shared with their own companies, those IoT devices, and their main purpose is actually to be internet connected. So there is a potential of privacy exposure from that.

And we have seen that that actually happens. The press has actually wrote many articles where there are devices, for example, sharing audio with Amazon workers and other persons like that. And this problem is very important because there are around 10 billion IoT devices deployed currently. And we want really to understand what they are doing.

Next slide, please.

In this work, we focused on the devices that are typically deployed in a smart home. We call them smart home devices, like appliances, smart lights, and other devices like that. So what we are interested in is to understand what the devices are doing. What does this mean?

We want to understand what is the destination of the traffic of them. Are these devices talking to their parent companies or are they talking

to some other third party? And, also, is the traffic staying in the country where it is generated or is it crossing the geographical boundaries? That's important because having state regulations may be different in another country, and sometimes even within one country, it may be different if the traffic stays inside or travels.

And, also, we want to understand is the traffic protected by encryption or not? What information is being sent? This is important because if a company is sending private information, it's likely that the user maybe is not aware of that. So we want to understand where things like that are happening and, also, if any information has been sent unexpectedly. For example, if you have a smart speaker, most people know that they have a microphone, but that microphone should not be transferring the voice all the time, but only when it's used. So we want to understand if that's true or not.

Next slide.

Answering those questions is not easy. It's actually a hard problem to measure privacy from IoT devices. And the reason is that the devices are typically black boxes that are much harder to analyze than mobile apps, for example. And the reason is that

manufacturers don't provide specifications, and probably for intellectual property reasons, all their information of how they work is not disclosed.

To overcome this problem, we want to use some technology. For example, we want to employ destination analysis and information inference so that we know at least what they are doing without having to look inside the device. In addition to the techniques, we need a proper tool to do that, because at the time of this work, there were no tools available to analyze those devices. So to solve this problem, we created some software that is able to collect the traffic and analyze that from IoT devices.

And we deployed this software in two IoT labs, one in the United States and one in the United Kingdom. The one in the United States can be seen on the left, the picture on the slide. And that lab is actually furnished as a studio apartment, where all the devices are put and arranged in a way that are easy to use for their intended purpose. We actually recruited like 36 students to use that lab. Of course, they signed an informed consent form. But then we could use their data to see how the device are using for their intended purpose.

Next slide, please.

So the devices that we considered are home IoT devices, in particular, smart cameras, smart hubs, home automation devices, like [indiscernible], smart thermostats, smart TVs, smart speakers, and many types of appliances from smart freezers to smart vacuum cleaners, for a total of 81 devices. And we were able to run 34,000 controlled experiments with these devices that were partially automated with our software.

We also monitor how those devices behave when not being used. For each device, we monitored 112 hours of inactivity to see what they are doing. And, finally, we also looked at what the devices do when they are actually used by a study participant, and we monitored them for six months.

Next slide, please.

So once we set up our environment and our analysis framework, now we want to answer the question that I said before. And the first one is, where is the IoT network traffic going? We have a lot of plots from a lot of studies seeing where it's going. But what is important to know about this is that the traffic is actually going to some entities that are not the main manufacturer or the parent company. Most of the traffic is going to other companies. Most of

them are cloud services and CDN providers. This is not necessarily a problem, but still, the traffic is still going under control of another entity.

What is probably more interesting is that some of this traffic -- like we have seen situations where the traffic is going really to a completely wrong company. For example, imagine that you have some smart TVs, or at least most of the ones that we analyzed, and it's contacting Netflix. Doesn't look strange, but imagine that you never installed Netflix, you never open it, and you never logged in, and that TV is still contacting them. So that is a bit of a problem that we found from the devices under test.

Also, we have seen that the majority of the devices send traffic to another country. Fifty-six percent of US devices contact other countries and 84 percent of UK devices contact other countries. Strangely, the UK devices contact a lot of US destinations. So this looks strange, but probably not too much if you think that those devices are typically developed by some smaller companies that maybe don't have the means to create an infrastructure in every region. But still, there are different regulations that apply in each of these regions, and we don't know how is this compliant like with the US and also

European regulations.

Next slide, please.

In addition to the destination, we were also interested to the traffic itself. Is this traffic encrypted or not? At the beginning of the presentation, I said that most of the traffic is encrypted, so it's really hard to understand how the devices are behaving. But we analyzed more in detail, and we have seen that a lot of traffic is encrypted, a lot of traffic is unknown. That means that we don't know if it's unencrypted or encrypted, but it's still encoded in a way that cannot be read. If you are optimistic, we can see that it is encrypted as well, but some investigation has to be done.

But still, there is some traffic that is red in the figure that is unencrypted, especially from cameras, that are some of the cheapest devices that you can buy. We looked at this unencrypted traffic and we've seen some negative trends.

Next slide, please.

So that negative trend is that a lot of devices across many categories are actually leaking unique identifiers, like MAC addresses and device IDs. Also, other content is being sent unencrypted, like some actions from simple devices, like turn on and

turn off, firmware updates activity and also when the device was set up for the first time, which behaves differently from when it's used later.

Next slide, please.

In addition to unencrypted traffic, we wanted to see if the encrypted traffic is also carrying some information. And the answer is yes. How did we do this? Well, simply, we look at our experiments. We tried to see how the traffic looks like when a camera is used to produce a video, and then we infer some patterns from this traffic and use these patterns to recognize when the video was sent in our traffic.

And by applying this methodology, we have seen that more than 90 percent of the devices that we tested that are able to produce a video or voice actually leak this information from encrypted traffic by using our technique. So one problem of this is that this technique can also be applied by any other entity. For example, an internet service provider has access to all the traffic that is produced in a household where the IoT devices are deployed. So they can infer activities and they can see what is done and what is not by those devices, which is a violation of privacy.

Next slide, please.

The last question we wanted to answer is if the devices behave unexpectedly or not. We have seen some cases where the devices behave unexpectedly. One of them is from popular doorbells, from actually different manufacturers. This doorbell was actually sending a recording of the video when a person was moving in front of them. This feature was not documented at the time and was not even possible to disable. So just owning and using those devices means that the device is self-recording when users don't expect that to happen.

We are also seeing cases of smart TVs, not just contacting Netflix, but also other companies, that are not related to the apps that have been used during our experiments, such as Google and Facebook.

And last, but not least, we have seen some very popular smart speakers being activated when you actually don't use them. Typically, they have a record. For example, Alexa can activate as my speaker, but unless they activate it if you say something that is different. For example, you could say, I like something, and some smart speakers activate. So this might just be a limitation of the device or maybe the manufacturers really want to know

what you like. So when you say, I like something, the device activates and sends a recording.

We're seeing other cases of unexpected behavior. For example, like a motion sensor reporting motion when there was no motion or devices spontaneously restarting or reconnecting. Those are all problems because when the device reconnects, they send all the information again. So they get more chances for violating their user privacy.

Next slide, please.

So all our findings of this study have attracted the attention of the press. So they wrote some articles that actually became very famous and attracted the attention also of the manufacturer. I will say later how we engaged with them to improve their devices.

Next slide, please.

So in summary, all the devices that we analyzed had some sort of problems, and the most important is that 57 percent of the devices and 56 percent, have non-manufacturer destinations or they send traffic to destinations abroad. This is something that is unexpected. And, also, the vast majority of the devices, 89 percent in case of the US, are vulnerable to activity, for instance, meaning that

a profile can actually be created for the users of the devices and how they use them by whoever has access to the network traffic, like the ISP.

This work had some impact. As I said before, the press covered some of our findings and the manufacturer contacted us to get more information about why the devices are contacting Netflix, for example. We provided them all our information, along with our experiment, so that they could double-check. We never got anything back, like yes, we've fixed this or we don't. But at least they are aware of the problem and will see that some of the latest versions of the devices actually have improved a lot, compared to when we performed this study.

Also, all the software we produced is publicly available on the website that you see. It can be used to create, for example, new testing labs, and we are aware that there is one in Italy that has been built. And all the software we collected from all the devices can be used to perform further studies by the companies to understand how the devices behave. And all this data is also available on this same website and has already been downloaded more than 100 times.

So this concludes my presentation, and feel

free to ask questions during the panel session. Thank you.

MS. ROUGE: Yes. Thank you very much, Daniel, for your presentation. That's so very interesting and a little -- so there is a lot of information out there, clearly, from the previous talks today, and there's a lot for consumers to understand.

So Pardis is now going to talk about her work with Carnegie Mellon and trying to package that information in something like a label so that consumers might word these things better.

DR. EMAMI-NAEINI: Thank you so much, Phoebe.

Hi, everybody, and thank you for joining my talk. I'm Pardis Emami-Naeini, and, today, I'm going to talk about our project to specify the contents of an IoT privacy and security label. This is a joint project with my colleagues, Yuvraj Agarwal, Lorrie Cranor, and Hanan Hibshi at Carnegie Mellon University. This work has been recently published at IEEE's Symposium on Security and Privacy, or S&P 2020.

Next.

IoT devices are everywhere. Some of the most common ones, which you might also have at home,

are voice assistants, smart doorbells, smart security cameras, smart thermostats, smart toothbrushes, and smart light bulbs.

Next.

And some less common ones are smart salt shakers, smart forks, smart umbrellas, and the most controversial of all, the smart toilets. And the list goes on and on.

Next.

People are increasingly purchasing smart devices. However, despite the surge in purchasing them, consumers are concerned about the privacy and security of the smart devices they purchase.

Next.

And people should be really concerned about these devices. After all, there's been news on how easily security cameras are getting hacked. But sometimes risk could have been mitigated if users of these devices were more informed. For example, after Ring security cameras got hacked, the company emailed their millions of users to use multifactor authentication. So maybe these devices could have not been easily hacked if users knew about better and more secure authentication mechanisms.

Next.

You may have also heard about Google putting its consumers at risk by forgetting to mention that its Nest secure hub had a microphone, or in other words, failing to inform consumers about the device sensors.

Next.

Another example shows how current key manufacturers are not transparent about their privacy and security practices as then some smart TVs are selling our data to third parties without disclosing it. Then it got revealed that Amazon is sharing unencrypted recordings of users' voices with its employees. Therefore, in many data collection scenarios, consumers are not informed about who their data is being shared with or sold to.

Next.

So what we need here is to find an effective way to show this information to consumers. And this is what we explored in this paper. We designed a privacy and security label for smart devices, somewhat similar to nutrition labels for foods. Our design label covers various privacy and security attributes related to the smart device. And as you can see, we include some of the important information about the IoT devices that IoT companies are not disclosing to

consumers, such as access control, sensor type, data sharing, and data selling.

Next.

Several pieces of legislation have been proposed, both inside the US and in countries outside of the US, including the UK, Singapore, and Finland, that would require IoT labels.

So I'm going to mention a few factors that should be included in these labels, but they don't contain too many details about what the labels should look like. And as you can see from the headlines, these proposals are primarily focused on security attributes without much attention to privacy practices. So our question here was, what should be included on an IoT privacy and security label?

Next.

To capture a holistic view, we invited a diverse sample of experts from industry, academia, government, and NGOs. To elicit expert opinion on the privacy and security factors, we followed a three-round Delphi process. In the Delphi method, the objective is to reach a consensus among a panel of experts without those experts directly influencing each other's opinions. This consensus is usually reached by conducting multiple rounds of interviews

and surveys.

In Delphi method, we have this concept of controlled feedback, which means that the aggregate output of the previous stage will serve as the input to the next stage. We have this feedback loop to allow experts to adapt their responses and eventually converge.

Next.

The first stage of the Delphi process is usually an interview study. We conducted semi-structured interviews with experts and asked them to specify the most important privacy and security attributes to include on the label. These interviews resulted in 47 attributes that at least one expert wanted to see on the label.

We then conducted the first follow-up survey. Each expert was randomly assigned to review one-third of the attribute and then specify their importance, as well as the reasons supporting their decisions. From this stage, we found the most common reasons for including or excluding a factor. And then we presented these aggregate reasons to experts on the second follow-up survey. And this is where we have the controlled feedback process.

On the second survey, each expert was

randomly assigned to review one-third of the attributes and, once again, we asked them to specify whether they would like to include or exclude the factor now after looking at all the reasons from the previous stage. To analyze the interview responses, as well as the opening answers from these surveys, we conducted thematic analysis, which is a recommended qualitative analysis approach, but information is high in subjectivity. We followed a six-step procedure recommended by Braun & Clarke to create the code book, find the themes, and merge them.

Next.

Experts acknowledge the value of the label in informing consumers' purchase behavior. An expert said, "What's good about a label is that it empowers the consumer to make a more active decision about cybersecurity rather than just being completely helpless as to what the security of her device might be. The average consumer doesn't have a privacy, security, or a legal department to review this stuff before they buy it. Enterprises do, but consumers do not, so someone's got to be looking out for consumers and giving the consumers this information."

Next.

In addition to informing consumers' purchase

behavior, some experts reported that the label could be a forcing function for manufacturers to be more accountable and transparent about their privacy and security practices. Moreover, experts mentioned that if the labels get adopted, it could initiate a competition in the market for manufacturers to enhance their practices. And I should mention, "There is value in forcing the company to write a list down, even if the consumer doesn't understand it. If you said, 'list your open ports,' there would be an incentive to make them few."

Next.

As I previously mentioned, experts wanted us to include 47 attributes on the label, which is clearly too many to show on a typical product package. Therefore, we designed a layered label with two layers. The primary layer is the concise format of the label, which could be printed and attached to the package of the product. And then there is a QR code and a URL at the bottom that directs consumers to the secondary layer, which has more detailed information and is in an online-only format. Online formats means that it can be updated as the firmware changes, which is critical as devices get updated often.

Another important reason to have this online

layer is to have a way to accommodate companies updating their privacy and security practices.

Next.

Some of the attributes included on the primary layer, their security update lifetime, type of collected data, availability of automatic security updates, and availability of default passwords.

Next.

Secondary layer has all the information from the primary layer and a lot more. Some of the attributes presented on the secondary layer were retention time, data inference, data storage, and whether there is any special data handling practices for children's data.

Next.

To assess our label's risk communication and information comprehension, we recruited 15 IoT consumers and conducted a one-hour semi-structured interview with each participant. In these interviews, we first asked participants to take a look at the package of a smart device with our label on it and define the attributes, as well as their values. We also asked them to specify the information that conveys risk to them.

We then asked participants to imagine doing

comparison shopping for a smart device from two different companies. We asked participants to compare the labels and specify which company had implemented better privacy and security practices and why.

Next.

By following a user-centric design process, we [indiscernible] improved the design of our labels, and this is the version of our label from last September.

Next.

In addition to the label, we prepared a specification document for users and IoT manufacturers. The content of our specification is based on the previous studies we conducted with experts and consumers and several IoT privacy and security references. In the specification, we provided the taxonomy of the label, consumer explanation for each attribute, list of the items to include as additional information for each attribute, and a list of best practices drawn from various references.

Next, please.

The real world impact. We would like to have our labels adopted. And to ease the process of generating labels, we developed a tool that allows

users to complete a form for different sections of the label and see the label being generated in real time. In the most current version of the tool, users can download the label in the format of JSON, XML, and HTML. Users can also work on the label offline, and then upload the saved JSON file to resume working on it.

Next, please.

To recap, consumers are concerned about the privacy and security of smart devices they purchase. And these devices are not transparent about their privacy and security practices. A label could be useful to provide that much needed transparency and inform consumers' purchase behavior.

Although a few proposals advocated for having an IoT privacy and security label, they are not clear about what the label should look like. I showed you some of these legislations in previous slides. To specify the content of the label, we conducted interviews and surveys in a diverse sample of privacy and security experts and identified 47 pieces of information our experts wanted us to include on the label. To fit this information, we designed a layered label. And what you see on this slide is the most recent version of our label.

To make the content of the label accessible to consumers, we put the most critical information on the primary layer and additional information on the secondary layer. To ease the process of label adoption and generation, we prepared a specification document, as well as a tool, to generate the label.

And, now, we're currently looking for manufacturers and retailers to participate in a pilot deployment of the label for their products. So if you want to show your commitment to security and privacy, this might be a great start.

Please visit iotsecurityprivacy.org to know more about this project and design your first IoT privacy and security label. Thank you.

MS. ROUGE: All right. Thank you very much, Pardis, for that presentation.

So, next, we're going to have Danny, who's an assistant professor at New York University's Tandon School of Engineering, present some work on IoT Inspector, which is a tool that collects crowdsourced information on actual IoT, what the IoT devices are transmitting in real time out in the wild.

DR. HUANG: Thank you, Phoebe.

So hello, everybody. I am Danny Huang. I am an assistant professor, starting fall of 2020, at New

York University.

So as the previous two panelists have talked about, we are constantly surrounded by smart IoT devices, like cameras, Alexas, smart TVs, whatnot. These devices could be constantly watching us or listening to us. But, today, I'm going to talk about a way for us to watch these devices instead.

So as you see in the next slide, here's a video of me watching Roku TV. On the top corner, on the top half of the screen is the Roku TV, running the CBS app. I'm just opening the CBS app and watching the live news streaming, without doing anything.

At the bottom is a screenshot of the network activities of the CBS app on Roku TV. I'll talk about how I obtained this screenshot a little bit later. But here's the big takeaway. On the Y axis, vertical axis, is the number of bits sent and received per second. On the X axis is the time, sped up at 10 times the speed. And each colored bar corresponds to some third-party advertising and tracking services that the Roku TV is talking to at the moment.

So remember, here I'm just passively streaming the CBS News, without doing anything on my Roku TV, and the TV is talking to three or four different third-party advertising tracking companies.

And one of the biggest ones is actually showing in pink. That is actually the Adobe Marketing Cloud. It's a little creepy, right? I'm not doing anything, watching TV, and my TV is watching me and talking to a bunch of advertising and tracking companies.

So in general -- next slide, please -- there are lots of concerns about IoT security and privacy, not just smart TVs, but Alexa, smart light bulbs, cameras. And as the previous panelists have aptly summarized, we don't know what's going on. It's a black box. We don't know what data is being sent. We don't know to whom the data is being sent to, and we don't know even from which IoT devices this data is coming from.

In general, there are two main problems, one for consumers, one for researchers. For consumers, these smart devices are like black boxes. We have no idea what they're going on behind the scenes. And there aren't very many good tools. If you want to start a Wireshark, good luck. It takes some time to set up a Wireshark to analyze network traffic. So that's the first problem for consumers.

The other problem is for researchers. Many research projects on IoT security privacy are limited to lab settings. Like security researchers

would buy a bunch of devices, like maybe dozens of devices in the lab, and connect them to the network, analyze the traffic over Wireshark, and analyze the traffic. The problem is that there are more than dozens of devices. There are literally thousands of smart devices in the world, and how to scale the analysis to thousands of different kinds of devices in the world remains an unknown problem.

So to solve these problems faced by consumers and researchers, our vision -- next slide, please -- is to develop a simple tool for consumers. Our vision is simple. We want to build a piece of software that provides volunteers with usable insights on IoT security privacy with one click. Here, there are two sets of colors, one-click and software. We wanted to make a tool that is simple to use. No hardware needed. No access point needed to be set up, but something they can download with one click. That's the first vision.

The second vision is usable insight, in green. We want to incentivize users to use our product. It's not just a research project. We want it such that users would want to actively download the software to find out more about their smart home devices, whether my camera is talking to some third

parties. That's useful insight.

So to provide this vision, we developed a tool called IoT Inspector, which you can download right now at this particular website on the screen. It is Windows-only for now, but we're coming up with Mac in the next version soon.

Next slide, please.

And here's what IoT Inspector does. At a very high level, it is a tool and it provides a data set. We launched the tool in April 2019. We've gathered more than 55,000 anonymous users at this point, and we're still gaining users and collecting data. Our users are anonymous, but some users have come out and told us that they're using the IoT Inspector. Examples include reporters from NPR, from Washington Post, New York Times. Some of these reporters are trying to analyze smart devices themselves, but lack the technical expertise, and they use our tool, such as in the case of NPR.

There are other users coming from, say, for instance, Consumer Reports, who told us that they're using this software, and the New York City Cyber Command emailed us and told us that they're using this software to analyze smart devices as well. So it's a tool that's currently being used by thousands of

users. You can try to download it, too. Just Google for Princeton IoT Inspector.

So in addition to it being a tool, we provide a usable data set for security researchers. In particular, since we launched the software in April 2019, we've collected network traffic data from more than 55,000 internet-connected devices. And we've attracted attention from more than 10 research teams requesting data, including academic and non-academic researchers. Academic researchers include NC State, CMU, University of Illinois, and University of Chicago, looking at different aspects ranging from, say, for instance, the privacy of smart devices, like, say, for instance, what companies devices are talking to, to security aspects.

For instance, I'm working with a group at UChicago trying to build a smart firewall to protect users from anomalous IoT devices. So essentially, we are doing a service for the community, not just for consumers, but also for researchers as well.

So how does IoT Inspector work? In particular, how do you use IoT Inspector? In the next slide, I'll show you how to download and run IoT Inspector in a test environment. Here, I'm showing a screenshot of Mac. Again, the Mac version will be

coming out soon.

So essentially, you download this executable from a website, and then instead of double-clicking in the finder window, you right-click and select "open." The whole reason is that this offer is currently not being approved in the Mac app store yet, and I'll tell you why it's not approved at the Mac store.

You right-click, and then a dialog box pops out asking you if you want to open it. You click "open." And then, finally, a browser window will show a list of IoT devices on their network. So it's a little bit small here. I'm going to walk you through this particular screenshot of the browser window.

So this is the browser window with IoT Inspector's main screen. In particular, it shows a list of devices on my network, like a Wemo smart plug, a D-Link camera, an Amcrest camera, et cetera. For each of these devices, you can inspect in real time what party it is talking to and how many bytes it is sending and receiving and whether the company being talked to is an advertising tracking company.

So again, I'm going to play the same video as I played earlier in the opening slides. Here's a video of me watching Roku TV. Top half screen is Roku TV screenshot and the bottom half is a live

screenshot, real-time screenshot of IoT Inspector as I'm streaming the CBS News app on Roku TV. IoT Inspector. Here the video is sped up 10 times, but you can basically see the CBS app talking to, basically, four different advertising tracking companies, the biggest one being the pink one, the Adobe Marketing Cloud.

So beyond just showing the live view of smart devices -- next slide, please -- we can see, basically, devices that communicate with advertising services. So, in this slide, we see a list of devices and the remote parties that are advertising tracking companies. The right-hand column shows devices. Under my account, I have a Samsung Smart TV. I have a Google Home. I have an Alexa in my home.

And then I can see on the left-hand corner, the left-hand side, the remote parties, they're identified as advertising services. Say, for instance, Samsung was talking to Samsung ACR, DoubleClick. So basically, the Samsung Smart TV is talking to Google advertising services.

So the question is how does IoT Inspector work to gather this insight? So the next slide has the answer. At the core, IoT Inspector analyzes IoT network traffic through ARP spoofing. And let me

explain to you at a very high level how ARP spoofing works.

So imagine we have a smart camera in the house and it is talking to the internet through our wireless router. Normally, without IoT Inspector, you would have to have your home network router to capture the traffic. But for us, we make the process simple. IoT Inspector, as is shown in the next slide, captures traffic through ARP spoofing. Here, we have an example of IoT Inspector running on a MacBook.

And IoT Inspector was sent to the camera to say, hi, camera, I'm the router. IoT Inspector will also tell the router, hello, router, I'm the camera. In doing so, the computer that runs IoT Inspector convinces the router that it is not a computer, it is a camera. At the same time, IoT Inspector convinces the camera that it is not a camera -- that itself is not a MacBook but a router. So this allows the traffic between a camera and a router to be intercepted by the MacBook that runs IoT Inspector.

IoT Inspector can, at this point, see the traffic going through between the camera and the router without actually rearranging cables or setting up a different wireless network. Basically, a one-click solution.

In the next slide, I'm going to show you examples of findings from real devices from real users by IoT Inspector. Basically, there are two areas of findings. One is security, one is privacy. And this is just the tip of the iceberg. I'm just going to explain a few examples. So security-wise, we found the lack of encryption on many smart devices, including devices made by big manufacturers, like Google and Amazon. Smart TVs on Amazon, some of the apps don't really use encryption. They use, in some cases, just plain HTTP. And in some cases, they use encryption, but they use -- surprise, surprise -- SSL 3.0, which is basically outdated encryption.

We have seen many devices with open unused ports, like cameras that have ports open on port 22. Like SSH, they are never used. But having unused open ports opens up opportunities for exploits by attackers. So these are some examples of security insights.

In terms of privacy, we found evidence of advertising and tracking on many smart TVs, including Roku and Amazon. It also found cross-device traffic. In particular, your IoT devices did not only talk to the cloud, they talk to each other. So basically, one device can potentially gather private information from

another device without your knowledge.

So again, more details about these examples are in our paper. Just visit our website at iotinspector.org.

So in summary -- this is the last slide -- we built a tool that provides transparency. It is a usable tool for consumers, being used by more than 5,000 users, and we've collected a larger scale data set for researchers to conduct IoT-related research.

But beyond transparency, we want to create action. In our next version, we will build IoT Inspectors such that it will alert users of any actual problems and will protect users from these problems. It is our hope that soon IoT devices will no longer remain as a black box, but we will be able to provide transparency and will provide actions to protect consumers.

With that, thank you, and I'm happy to take any questions.

MS. ROUGE: Thank you so much, Danny.

So yes, if you have any audience questions, you can send them to the privacycon@ftc.gov address and we'll try to get to them.

To start out our discussion, though, I will first ask a question of Danny. So one of the

questions you might have, as far as using your IoT Inspector, you know, given it's doing this ARP spoofing - let's say I don't know what that is. So does it potentially introduce any -- would it cause any problems on my network? Would it affect any of the devices in any way?

DR. HUANG: Yes. If you want to download on the website, we have a big warning, where it's saying that it's going to slow down the network. Basically, instead of having traffic directly going through your router, as in traffic directly coming from your smart devices to the router, it takes an additional hop to your computer, and that's going to slow down your network. That's going to slow down your smart devices.

So if you're, say, watching Netflix on a smart TV, you may experience degradation of traffic. You may still have HD content. You may see some blurriness. This is from the real experience of me running IoT Inspector myself in my house.

MS. ROUGE: So this is a question for Danny and Daniel. Since you've been looking at this traffic that's coming from these smart devices, have there been any really surprising results?

I know, Daniel, you talked about, for

example, the smart speaker activating when you weren't expecting it. But have either of you seen anything really notable that was very egregious or anything like that in your results?

DR. DUBOIS: Yeah, so we have seen something that was not bad. We had to investigate that. I don't know if I can name the device of companies now, but there was like one doorbell, but not the most famous one, but one that can be bought on Amazon, that was encrypting the traffic without verifying the certificates. So that means that if you encrypt in that way, encryption is completely pointless. You can actually do man-in-the-middle attacks on the device and get [indiscernible] the device and password. So if it happens with a smart camera, that's a problem.

And, in general, we analyze a lot of categories of devices. And smart cameras, for this point of view, are one of the ones that behave in the worst way. Because, also, the companies are very small. They typically just buy hardware that is made by another company. They also use software that is by another. They customize it a bit. So it's often not updated software, full of bugs, and it's concerning. Besides the unencrypted -- like the fully encrypted Traffic where [indiscernible] like some cameras are

sending traffic to other like residential addresses.

We tried to understand why this was happening. We don't know, but they mentioned that your camera is contacting a bunch of addresses. At first, we thought it was hacked. Maybe it was part of the bottleneck, but we didn't find really any evidence of that. We might think that the device was probably uploading some of this data to other devices in a way to reduce the use of their computational systems. But we're still -- that is a question of development. We see things that are strange, but it's really hard to see what they do because we don't have control over this stuff. We just see the traffic that is strange.

And the only thing we can say is that when you consider IoT devices, think about what you do for a mobile app. You can install an app from a small company. You don't know what this app is doing. It may send traffic. And what you can do is if you really need to use the app, just use it, but knowing that you are exposing yourself in some ways. Or you can basically delete the app. That means you unplug the IoT device from the internet. Sometimes they can still be used without being the internet.

Also, smart cameras, sometimes they allow you to use them on a network that is isolated from the

internet and they still work. So those are only the possible ways that come to my mind where consumers can protect themselves from this situation.

DR. HUANG: And some of the surprising things that we found is actually from smart TVs. One example is that, say, for instance, the Roku -- I'm sorry, the Amazon smart TV screen, for instance, has a built-in feature that basically says you can actually opt out of interest-based advertising. If you think that turning this off, turning off the interest-based advertising would reduce tracking, you're wrong.

So in one experiment, we found that we turned off interest-based advertising on both Roku and Amazon. We still see these devices potentially sensitive information to some third-party advertising tracking services. So yeah, it's one example. Tip of the iceberg for some of the privacy issues we found in smart TVs.

MS. ROUGE: Okay. So I guess sort of following on, Daniel, you mentioned some things that you might want to do if you get one of these smart devices to address some of the concerns.

I guess I'll start with Pardis. If you are a consumer that wants to buy a smart device, and I'm

watching this PrivacyCon and I'm like, wow, there's a lot of things to be concerned about, what's the first thing that you would look for? Like if I bought a smart device, what's the first thing I should do if I unpack it? Is there any setting I should change? Is there anything I should look at to make sure it does or doesn't do, anything along those lines?

DR. EMAMI-NAEINI: That's a very good question. So I think, basically, privacy really depends on your own preferences, definitely. So you may be concerned about some type of data and you may not be concerned about other types of data. But apart from that, I think what is really important for consumers to know about is to know what types of controls they can have, if they want to change them or not.

So basically, when you purchase a smart device, I think the first thing that you should do is to understand the settings of the device, the privacy and security settings of the device, to basically know how you can change data sharing, how you can opt out from data sharing, data selling, for example. Do you have this option?

And another, I think, important thing is to understand the basics of privacy and security

information of the smart device. For example, whether the device is the default password or whether you would get security updates. So there are some critical information, privacy and security information, some basics that you should really know about. And then other than that, the types of controls that you can have. So I think that's the first things that I would recommend consumers to do.

MS. ROUGE: Is there a particular setting that if I bought a smart device I should make sure it has or that I would immediately change when I bought it home?

DR. EMAMI-NAEINI: Yeah. So one thing that I'm concerned about, for example, is data being shared with third parties or my data being sold to third parties. And something that I would look for is, can I opt out from data sharing? And so this is the first thing that I would look for. But as I said, privacy is very subjective, so it really depends on your own preferences.

MS. ROUGE: Got it. That makes sense.

I guess, Danny, I would ask you the same question. You're looking at all of this data coming out of the smart devices. Is there something specific that you would look for as a control or something you

would want to change when you brought it home?

DR. HUANG: The first thing I want to do when I buy a new device is to run it over IoT Inspector and see what's it doing, basically.

And just echoing what Pardis said earlier, maybe different people have different privacy preferences. For me, I don't have a lot of tolerance for weird behaviors, but for others, maybe they would be okay with it. So I think having a tool like IoT Inspector allows users to gain transparency into exactly what's going on with the whole network and make a decision themselves, whether to return the product or continue using the product.

MS. ROUGE: Daniel, I'll just ask you the same question.

DR. DUBOIS: So usually, the problem is that, depending on -- like a normal consumer, is not able to configure to the privacy settings in the correct way because usually they are complicated. Sometimes, like in my experience installing like 81 IoT devices, I had trouble to configure some of them. So even if you have a PhD, it might not be enough to do that properly.

So what I do, and I cannot suggest other people do that unless they have the technical

capabilities, is to try to isolate the devices from the public network as much as possible.

There are some open source tools, like Home Assistant, that are difficult to use for most people, but maybe in the future, there will be easier versions of that. And those tools can actually isolate the IoT device from the internet and they can control what the devices are doing and what they are not. And those tools are open sources so they can be analyzed. The code is open for everyone. And if your IoT device is behind a tool like that, it's much safer for use than if they use like a black box solutions, that you don't know exactly who they are talking to, what they are doing, what they are saying, and everything is like a question mark.

MS. ROUGE: All right. I'll start with Pardis again on this question. So as people become more aware -- you know, we see lots of headlines. We have this whole event, we have your research and the others getting out there. There's a lot of marketing talk, as far as how much IoT is going to proliferate, and we definitely see a lot of devices being sold.

Do you think, either in the course of your research, as you were asking questions or when you explain your research to others, do you see any

changes in people's feelings about IoT, as far as these are devices, okay, these clearly require a lot of care and feeding? Do you see people changing their minds and thinking differently about how IoT devices should be used in their home?

DR. EMAMI-NAEINI: Great question. Yeah, so in the interviews that we've conducted over the years, we've found that participants are concerned about the privacy and security of smart devices. And they know -- for example, smart speakers are very famous. So they know that, for example, they are doing some weird stuff because they've seen that on news, for example. And so they're very concerned.

But at the same time, when you ask them whether they'd purchase the device or not, they would still purchase it. And this is not really about whether they're concerned or not. I think it's mostly about whether there are alternatives in the market, and if consumers know that these alternatives are better, in terms of privacy and security.

So I think there are basically two issues, that you don't really know which devices are better and you don't even know how to define better privacy and security, because at the time of purchase, you have no information about the privacy and security of

these devices. So I think if you can solve these two issues, in the market if you can have better products, and if you can convey this to consumers that these are really better products, then I think consumers would be better able to apply their concerns. Now they're concerned, but they don't do anything about their concerns.

MS. ROUGE: Thank you. So one question I wanted to make sure -- to circle back -- we got from the audience. Danny, you had mentioned that your app is not approved for the Mac app store. And I'm wondering, could you just quickly explain why that might be?

DR. HUANG: ARP spoofing. It is an attack, basically, but we are turning this attack for good. That's a short answer.

MS. ROUGE: That makes sense. So yeah, we're right up at time, but I guess I just wanted to give each of you a chance to sort of -- if there's kind of one thing that you would want consumers to come away with from this presentation and from PrivacyCon, what's one concept that you'd like them to come away with?

And I guess we can start with Daniel.

DR. DUBOIS: Yes. So one thing that is

important to know is that IoT is not going away. It's becoming more common in our lives. So we cannot think that we'll reduce that type of exposure by just not buying this stuff. You can already see that. Try to buy a TV that is not smart. You will not be able to find one. And this might become common with many other objects. Of course, you don't have to connect them to the internet.

In my house, I have a device, a cooking device, that doesn't have any interface on it. It needs at least Bluetooth to work, because it requires a phone. And even if it needs Bluetooth, then the companion app of the device connects to the internet in some way. So we have to learn how to use these devices properly, and we need to keep doing research on the privacy concerns on them, because regulators will notice when these things are happening.

And as it happened already from the apps, the privacy regulations will be updated and the devices will be safer to use, hopefully, and there will be more transparency.

MS. ROUGE: Great.

Pardis?

DR. EMAMI-NAEINI: So this is not directly related to my presentation, but it's related to the

interviews that we've conducted about this study. So I want consumers to know that smart devices are not pieces of furniture, that you would just have them in your home and that's it and then you don't need to think about them. Because I've seen a lot of these anecdotes, that people think that -- they're getting used to these smart devices and they don't really care about them. They don't really do anything to change their settings or even think about them.

But that is not the case. These devices are powerful and they will get more powerful in the future. So they have these sensing capabilities. And you should treat them as things or people who can listen to you, or even can see you. And if you treat them like that, you will change your behavior in front of them.

MS. ROUGE: That makes sense.

Yeah, and Danny?

DR. HUANG: Just echoing Pardis' point, these devices are getting more powerful and they're getting more prolific, so what do we do? There's no current signs of them improving, in terms of security and privacy, so what we do?

Two suggestions as you walk away from these presentations. One, set up a separate network, just

for smart devices. Many home routers allow you to set up a guest network. Just connect your smart devices to a guest network. So increasingly, you are working from home, you probably don't want your regular computers to be talking to and from the smart devices, if they are ever hacked. So one, set up a separate network.

Two, for devices like smart TVs, they have capability of tracking you and following you around. So, say, for instance, you want to start looking at some shoes on your website and start seeing these shoes in a smart TV, so what do you do? Use a separate account, a separate email address for your smart TV account. For me, I use a -- create a completely new Gmail account, just for my smart TV, so that I don't have advertisements that follow me around.

MS. ROUGE: Those are good practical suggestions.

All right. Well, we went a little over into our lunchtime, but thank you very much for your presentations and the discussion. This was really interesting. And we'll be back after lunch with presentations about specific devices, like cameras and such. So we will see you back here

then.

(Lunch recess.)

SESSION 4: SPECIFIC TECHNOLOGIES:
CAMERAS/SMART SPEAKERS/APPS

DR. BANKS: Okay, welcome back. So in case you are a bit drowsy from lunch, we're going to shake things up, do things a little bit differently.

Our researchers today have lots of empirical data to share, so what we'll do this time is that we'll have Q&A after each presentation. So immediately following the presentation, we'll have a brief Q&A with the presenter. So for the audience, that means you'll send your questions at the end of each presentation. So please send your questions to privacycon@ftc.gov as you think of them.

With that said, let me introduce myself. My name's Lerone Banks. I'm a computer scientist in the FTC's Division of Privacy and Identity Protection. Welcome to Panel 4, where we'll talk about some empirical data related to specific technologies.

With that said, let's get started with Madelyn Sanfilippo from Princeton. She's presenting privacyrisks in disaster-response apps. Welcome, Madelyn.

DR. SANFILIPPO: Thank you. Hopefully, we're currently on the title slide. The projects that I will discuss on behalf of my wonderful collaborators

today focus on privacy issues around apps during emergency circumstances. The paper we submitted focused specifically on hurricanes and natural disasters, but I'll also discuss some implications along the way for other crisis contexts, including the current public health emergency, as we're exploring these things in follow-up research.

Moving on to the next slide, this is pertinent given that many factors shaping social norms and emergencies, such as individuals' inclination to share more personal information under disaster situations, as documented in many previous research studies, extends to other crises. During a hurricane or a fire, people think it's appropriate to share their location with first responders, for example, just as during a pandemic, many people are willing to share information for the purposes of contact tracing with public-health officials, although not necessarily with other actors.

As we see on the next slide, there are a variety of digital platforms to structure information flows during emergency circumstances, with government agencies both providing their own platforms and channels, as well as recommending others, in addition to the prevalent use of tools, like Facebook Safety

Check.

In this study, we focused on those apps that were recommended to users during hurricane season, as can be seen on the next slide. These apps can be divided into five distinct categories: Those apps developed by government agencies, such as FEMA; those apps developed by trusted organizations that partner with the public sector to provide relief, such as the Red Cross. There are also apps that are general weather apps recommended during these times.

Additionally, there are hurricane-specific apps from private sector developers that can be divided into two additional categories: Those that are transparent about their development, in contrast with those that appear, either by name or branding, to belong to government agencies, despite private development. This latter category is problematic from a consumer protection and deception standpoint. And many frequently change their names, though not necessarily their code or behaviors. As Apple or Google take them down from the marketplace, they simply reenter the market with superficial changes.

Moving on to the next slide, we framed our analysis of these apps in terms of privacy as contextual integrity. This is to say we conceive of

privacy as the appropriate flow of personal information in a context, in contrast with the privacy harm. That can be understood as inappropriate information flows. In this sense, information flows themselves can be deconstructed in terms of information subjects, senders, recipients, and types, as well as transmission principles, in order to compare them and understand where violations of users' expectations might occur.

We use this framework to make comparisons throughout our overall research framework, as seen on the next slide. We compare the context of privacy policies as endogenous governance and regulations as exogenous governance of information flows with actual information flows and practice, which we identified from a combination of static analysis of permissions; dynamic app analysis of flow traces, including the recipients and decryption of traffic to identify information types; as well as user experiences, described anecdotally in reviews and simulated through our own controlled experiments within virtual mobile machines.

As we can see on the next slide, the governance of disaster information flows is extremely complex due to polycentric arrangements of institutions, with different agencies having

significant say, in addition to federal regulation and distinctions between Personally Identifiable Information as PII and Sensitive Personally Identifiable Information as SPII. The key points for our purposes today are highlighted on the next slide.

Specifically, I would like to note both the ambiguities of routine uses, which is likely a source of discontinuity between government and partner organizations, whose routine uses vary significantly; and the nuance of trusted partners, including other government agencies at various levels; utility companies; hospitals; and relief organizations, from the Red Cross to religious groups, and things like Team Rubicon. These partners are subject to restrictions, which are actually similar to those on federal agencies under the Privacy Act, including limiting redissemination and to need-to-know circumstances.

Some of you may remember that this came to be an issue around FEMA inappropriately sharing too much information about hurricane and wildfire victims with contractors in 2019.

We can move beyond governance to look at flows on the next slide. This is a visualization of information flows from the apps within our set that

are sharing location information, for a variety of different reasons, with many third parties. Some of these flows violate not only regulations, such as the Red Cross sharing location of victims with Flickr and social media companies via an installed third-party library, but many are not disclosed in their privacy policies.

I will differentiate between these types of violations in a minute. But, first, I'd like to briefly revisit some additional concerns raised by users on the next slide.

In addition to requirements from governance, user expectations should also theoretically be met under conditions of contextual integrity. Some users noted that user permissions or options to control personal information did not work on the very apps being promoted as the best to use during a hurricane. Concerns about persistent tracking were particularly significant in these complaints.

Further, others noted, in relationship to the Red Cross apps, which are depicted on the next slide, that some of the persistent tracking information was too accessible to anyone who requested it, tracking individuals in real time and indefinitely, though both of those problems have now

been corrected to an extent. The concerns about the ability to track victims via Red Cross by former intimate partners in cases of domestic violence were not, however, addressed via the updates. We tested these and other complaints, as well as explored what options users actually had to control their privacy, as depicted on the next slide.

Specifically, we found that many of the apps considered provided insufficient or misleading options, with the most obvious problem being that, despite user preferences not to share location with a particular app, if they shared location with another app, it might be shared in order to personalize their disaster apps as well, in addition to personalization of other outside apps. Five hurricane apps shared location with one another, as highlighted in red. So if an individual were to share with any one of those, it was happening in all five of them. The nature of those relationships directionally is further described in our paper.

In comparing all facets of our privacy analysis, we classified the privacy compliance of all of the apps in our study as categorically represented on the next slide. Here, we differentiate between apps that are wholly compliant with policy, those that

comply either with their own privacy policy or regulation, but not both, and those that are compliant with neither.

Looking more specifically on the next slide at those that are compliant, as highlighted in green, there were three apps that behaved appropriately, transmitting no personal information to any third parties, complying with all expected regulation and behaving in practice as was disclosed in the user agreement.

On the next slide, as highlighted in yellow, we see apps that did not act in ways consistent with information flows described in their own privacy policies, but that did not actually violate any laws or requirements under contractual obligations with the government. These apps simply violate user expectations.

On the next slide, as highlighted in orange, we see apps that comply with their privacy policy but that are otherwise problematic. Some of these, such as Dark Sky and Global Storms, inappropriately share data with trusted partners, though they themselves are not trusted partners. The others violate user expectations and are problematic from a deception standpoint, rather than privacy violations, as they

appear to be NOAA apps, when, in fact, the National Oceanic and Atmospheric Association does not provide a consumer app.

On the next slide and highlighted in red, we see our classification of Red Cross apps as problematic due to information flows that are not disclosed in privacy policies and that violate contractual obligations with FEMA on user data. So note that when I explained the user experience violations in a previous slide, this is different than actually violating the policies. We have brought these to the attention of the Red Cross, and we would attribute, based on the information that we have, that these issues are primarily due to the use of third-party libraries and a lack of communication between technology and policy offices within the organization, rather than some particularly malicious intent.

Overall, the implications of this study are multifaceted, as depicted on the next slide. The study helped us to identify what aspects of context shape the unusual and under-addressed social norms that apply about information sharing. Particularly, emergencies' end and duration ought to be considered as an aspect of this context. Further, people do not

normatively object to the information flows during disasters, but rather to the aggregation and reuse of this data outside of disaster circumstances or to specific actors as recipients in this context.

Finally, users ought to be able to reasonably expect what flows are going to happen in practice because privacy is contextual. However, certain events, including hurricanes and the current COVID pandemic, change expectation.

Moving on to the next slide, we build on this study to explore privacy issues around contract tracing apps in our subsequent research, and thus far, see very similar patterns. There are, again, the same categories of apps, including government apps and privately developed government-imposter apps. There are also general health apps that have been repurposed for the current context. Further, there are also efforts by major platforms, which are sort of in parallel to the Facebook practices during hurricanes.

Yet, a really major difference is the increased objections to the potential for misuse prior to data collection or use of the app by broader sections of the public. Overall, what we see is that governance ought to apply to all parameters of information flows, not just to a small subset of

actors and information types. This would provide clarity around regulatory expectations and practices so as to better inform users, and would likely meet user expectations under emergency circumstances better.

The final slide, an overarching issue that connects this work to an emerging concern in privacy research is that of context collapse in an emergency and other circumstances as introduced by digital technologies. Many of the concerns around contact tracing, as well as around Red Cross's Safe and Well Program, stem from overlap of actors, information, and circumstances that people feel are inappropriate, particularly in terms of the long-term consequences.

I'm happy to address any questions that may have been sent to FTC right now, though, also, feel free to reach out to me or my collaborators.

Thank you.

DR. BANKS: Thank you very much for that informative presentation, Madelyn. This is very timely research.

And you actually mentioned the idea behind one of the sort of first questions I think that we have, which is about contact-tracing apps. And so I understood you to say that, in some -- well, I guess,

the first question is, how much analysis have you done with contact-tracing apps? And what has been the outcome?

DR. SANFILIPPO: So this is, I think, a really logical direction to take the previous study, the paper that we shared is sort of a pilot study, exploring how we could bring all of these different types of data analysis together in order to understand relative levels of compliance and violation of user expectations. What we have done thus far with contact-tracing apps is to collect, obviously, all of the apps and begin testing some of the user concerns that have been articulated.

However, the dynamic app analysis following traffic has not necessarily happened in every case because some of the apps we're considering haven't actually been deployed, and so we're sort of doing preliminary analysis on some of these based on user concerns. But we're able to categorize them according to the same parameters about who has developed them, how transparently they've been developed, and what types of exogenous regulations or policies might apply to them.

DR. BANKS: I see. And so based on some of that preliminary analysis, you are seeing at least

some early signs that suggest some of the issues identified in some of the other apps that you've looked at in your paper are also starting to crop up in these contact-tracing apps. Is that right?

DR. SANFILIPPO: Yeah, yeah.

DR. BANKS: Yeah. Given that, do you have some recommendations for the approach that regulators should take towards analyzing these apps? And, particularly, given the fact that, unlike maybe other disasters, like hurricanes, where there is somewhat of a definitive beginning and end, the pandemic, unfortunately, does not necessarily have that clear-cut delineation.

So the question is really, do you have some suggestions for how regulators should approach it? And are those suggestions different based on the nature of a pandemic, which, I think, is different from other disasters?

DR. SANFILIPPO: I think that's a really good point. However, I think that duration and sort of a time element of this particular emergency context could still be addressed. So it would be my recommendation not that we think about when the pandemic context is done definitively as a whole, but, rather, when individual harms that could be associated

with exposure to someone with COVID actually terminate.

There is an end to a period in which someone may have been infected through this. And so it's not necessarily a matter of maintaining all of that data set from beginning to end of pandemic, but, rather, maintaining it only as long as is necessary in order to trace particular harms and to protect public health.

Further, I think making guarantees that this data won't be used for other purposes would be much more consistent with individuals' concerns. For example, the level of trust between a public health department and trust in particular commercial platforms is not necessarily equivalent. And so I imagine that much more compelling arguments and impetus to use some of these things could be made if the actor responsible for this data and making assurances that it will not be used for other purposes or after a period of time would be much more trustworthy from the perspective of users, at least in terms of the complaints that we're investigating, or concerns we're investigating at this point.

DR. BANKS: I understand. Let me ask you one more question from the consumer side.

So I assume that you've probably analyzed more privacy policies than a typical consumer has ever actually read, right? Do you have some advice on how consumers can read them more effectively to address those concerns, some tips that you may have, given your comprehensive analysis, that you can advise consumers on how to find the most relevant information or interpret legal jargon?

DR. SANFILIPPO: Yes. It's not easy to read privacy policies at all, and that is something I spend quite a lot of time doing. In particular, one of my coauthors, Yan Shvartzshnaider, has done extensive research on how we might better visualize or represent this information in a way that's more easily comprehensible to users.

And so on the one hand, I would recommend that people try to communicate particular information flows that they are structuring in their policies in a standard format, as opposed to in the middle of large, complex paragraphs. That information maybe still needs to be there from a legal perspective. But from a consumer standpoint, looking at a table that says this information is being collected this way and will be used for this purpose is a lot more understandable.

From a user perspective, right now, I think

flagging particular issues that you might be concerned about or third-party advertisers, for example, and looking specifically for those things amongst the text is, perhaps, one of the most useful ways you can sort of skim these policies without necessarily reading through all of the legal jargon yourself. You can sort of flag particular concepts, or third parties, or uses that you're uncomfortable with, and read to see if they are covered within a policy.

DR. BANKS: That makes total sense. Hopefully, the consumers that are listening today will take some of that advice. Thank you very much for your great work.

DR. SANFILIPPO: Thank you.

DR. BANKS: Next, we'll have Christin Wilson, who will present the team from Clemson's work on getting malicious skills into Amazon's Alexa Skill Store.

Welcome, Christin.

MR. WILSON: Thank you. So good afternoon, everyone. Before I begin, I would like to thank FTC for providing me this opportunity. I would also like to thank my research team at Clemson University, especially Dr. Long Cheng, Dr. Hongxin Hu, Song, Jeffrey and Daniel.

So we are excited to present our paper, "Dangerous Skills Got Certified: Measuring the Trustworthiness of the Amazon Alexa Platform." So a brief introduction, the user base of Amazon Alexa has been rising rapidly over the last couple of years, and this actually encourages third-party developers to build new skills. So here, "skill" refers to a voice app, so that's what the Amazon Alexa platform calls it.

So a skill has to be certified by the team before it's published to the end-users. And a weak rating system will result in malicious skills entering the store. So these can be privacy-invasive, this can disseminate inappropriate information to users, et cetera. So we are especially concerned about children and the skills meant for them.

So on the next slide, we have our three research questions. Number one, we want to evaluate whether the certification system is efficient and trustworthy. Number two, do policy-violating skills exist in the skill score currently? Third, how do Google Assistant's certification systems compare?

Next slide.

So before we move further, let's just discuss how can third-party skills collect data. So

there are two methods. The first method is to configure permissions in the skill. So when a developer develops a skill, he can just configure some permissions. So what happens is when a user enables the skill, a prompt will be sent to his phone -- the Amazon Alexa app on his phone, and you have to provide permission. And this data is actually taken from the developer account, so you're not providing it. It's just taken from the account.

The second method is to collect the information through voice. So this is directly done during an interaction. So Alexa will just ask you, what is your name, and you can speak it back to them. So in this case, no prior consent or permission is taken while this skill is enabled.

So now, we come to the next slide.

This is the first research question. Evaluate the certification system. So we developed skills that violate 7 privacy and 14 content policy guidelines. So this is actually provided by the Amazon team to the developers in the developer documentation. We do have some ethical disclosures. We have obtained approval from our university's IRB. We do not use or share any of the information collected, if any. And we remove the skill as soon as

we see that it is certified. And for high-risk violations, we do try to provide a disclosure.

So the next slide discusses the first subsection. It's violation of children-specific policies. So these policies mainly focus on the collection of data from children and the content provided to children. So like in the image, you can see that you do not want a skill asking a child for personal information or encouraging him to drink or smoke, or do something illegal without telling his or her parents.

So in the next slide, you can see that we were able to get 119 skills certified in this category. So there are two types of skills, one that could collect data and one that would provide some inappropriate content. For the ones that could collect data, it had the following features. So it could collect personal information -- and remember that the users are actually children and it's completely restricted by Amazon to collect personal information from children.

The second one was it could save the collected information in the developer's database. So we could save it in our DynamoDB. The third one is they didn't provide a privacy policy. The fourth one

is no prior consent is taken from parent or guardian before the collection of data. So when the skill is enabled, nothing is told to the parents, and when the child uses it, personal information will be taken from them.

And then no access is provided to view, delete, or modify the collected information from our data set. And, also, the developer account details that we used were fake, so they can't actually contact us to ask about the collected information. The other skills that were published had content not suitable for children or encouraged them to use services outside of Alexa.

So the next slide contains the next two subsections. These are a violation of general content guidelines and privacy requirements. So these policies are mainly for, like, the general audience. And we were able to get 115 skills certified. So these either had the data-collection violations, just like we discussed with the children skills, or it had some data that is not supposed to be told, like there were, like, health-related information, promotions, disturbing content, advertisements, promotion of alcohol, drugs, illegal activities, et cetera.

So in the next slide you can actually see an

example of a skill that we submitted. You can see that the skill "moral stories" was live on the -- it asked for the user's name. It got the full name and generated a story with their name. But you can see that no privacy policy was provided. It's a kids-category skill, and you can see that the name was actually saved in the DynamoDB database. So we made this just for illustration purposes, and we ourselves provided that name and it's not actual user data.

And next slide.

So the experiment results, we've been able to certify 234 skills in total. So it's not 234 unique skills, but it's like 234 different certifications, I would say. So it was conducted over about a year. And we did have to resubmit some of these skills. So initially, some of these skills were actually rejected, so what did we do? We just had to use a simple counter to delay the session in which our privacy-policy violating response is delivered.

So if we set the counter as four, the first four responses from a fact app will be perfectly fine, and the team would actually certify the skill based on that. And after that, since the counter was fourth starting from the fifth, the policy-violating response will be very good.

So in the next slide, we are actually discussing our observations. So the first observation is the inconsistency in checking because we got different responses for the same exact skill each time we submitted it. The second one is limited voice checking. So they're not actually looking at the code or anything. They're just talking and just seeing if the conversation matches.

The third one is overtrust placed on developers. So this is actually evident in the image shown. We marked that the skill does not contain advertising, but the skill actually contains -- so the certification team, instead of actually checking for it, they just trust the developer and certified the skill. The fourth one is humans are involved in certification, and it's not an automatic process yet.

The fifth one is negligence during certification. So this is because, for some skill sets, especially the ones that can ask for the name from users -- this was actually a story skill and this was asking for the name in the first session itself. But we got some rejections for that. The certification team actually had a problem with the content of the story, like the story had some violence in it or something. But they never complained about

us asking for the personal information without providing a privacy policy, and that, too, from kids. So this shows a negligence from their part, I would say.

Next slide.

So our second research question was to look for existing policy-violating skills in the store. We only tested 825 skills. There are about 100,000 skills, which we can't actually check. So we just get skills that either had a negative review or had a privacy policy provided.

So by looking for skills that had a privacy policy provided, what we wanted to do was, like, Alexa only requires skills that collect personal information to include a privacy policy. So this was our assumption that they might be collecting personal information. Made us look through them. And we identified 52 skills with possible privacy violations. Again, we use the word "possible" because we can't really ensure whether some policy violations actually existed, because we can't access the code. There was also 51 broken skills that didn't work. So there is no constant check being done to see if the skills are working perfect.

So the next slide, we have a few examples

about some privacy-policy problems we saw during the manual testing. So the image on the right is actually a skill developed by Amazon, and it's actually a weather app, and it's available by default on all Alexa devices. So it mentioned in the description that it collects the user's device location, like any other weather app would do, but it does not provide a privacy policy in the usually allotted space.

So since this is an Amazon-developed skill, it's okay because you can actually find one in the bottom of the page. But there are other skills that are not developed by Amazon and mention about collection of data in their description, but don't really provide a privacy policy.

The other image is an example of a badly written privacy policy. We don't really know what the developer actually meant by that line. There are also examples of privacy policy URLs leading to the Google search webpage, other developers' privacy policy, et cetera. So these were links provided before certification. So during certification, the team could actually see the URL, and they just might not have gone through it or just neglected it.

Next slide.

We also did a preliminary comparative

measurement on the Google Assistant platform. So we got 15 out of the 85 kids actions certified. And for general actions, we got 101 out of 185 certified. So actions is the Google equivalent of skills. This data just suggests that Google's vetting is better, but, again, this is a preliminary study, so we can't really state that.

We did see some inconsistency in feedback here, too. And the post-certification vulnerability exists here as well. So this vulnerability means that once a skill is certified, you can make changes, and then it will be deployed to the live audience without requiring a recertification. So yeah, this still exists in both Amazon and Google, and I think this has been discussed in some other papers as well. We did manual testing on the 76 kids sections as well -- or they call it "actions for families" -- and we found one problematic action. So this goes to say that there are policy-violating skill actions in the Google directory as well.

On the next slide, we have a responsible disclosure. We have reported our findings to both Amazon and Google. The Amazon security team is still working with us on investigation and resolving this issue, but it's still going on. It's not resolved

yet.

The Google team, on the other hand -- the counter-abuse systems actually issued us an award as part of the Vulnerability Reward Program for our work.

And coming on to our final slide, you can see that we have provided a website link. More details and video demos are actually provided in the website. You guys can take a look at that. If you have any questions, I can take them now.

DR. BANKS: Great. Thank you very much, Christin.

So first, for the audience, if you do have any questions for Christin, please do email us at privacycon@ftc.gov right now if you have some questions. He presented a lot of information and, hopefully, you have lots of good questions.

So while we're waiting for a few audience questions to come in, I'll ask you a question that kind of starts at the end of your presentation. You mentioned that you reported these results to Amazon and Google. First, congratulations on the award from Google. That's an accomplishment in and of itself.

So my question is about Amazon's response. So how receptive were they initially, I guess, to your findings? And what was their feedback and,

particularly, your claim about the ease with which it is to circumvent their process? And how did they address the issues that you raised?

MR. WILSON: So I would say they were very eager in our results. They got into a call with us as soon as they got the email. We did have a lengthy meeting discussing about what work we actually did, how we see it was. Even they asked, like, did you guys try some other method? And we were like, no, this was collectively very easy to do this, so we didn't have to go for harder techniques and stuff.

But, yeah, they're were really eager to know about our work. They are still investigating. They have tried to get as much information from us, so they're actually looking at the certification log of every skill that we actually published. So I think they are doing a good job, but it's still being done, so we don't know the final result yet.

DR. BANKS: I see. Well, at least it sounds like it's a pretty collaborative process --

MR. WILSON: Yes.

DR. BANKS: -- which is not always the case in these types of instances. So I think it's good to see that.

MR. WILSON: Mm-hmm.

DR. BANKS: Let's see. I think there might be a question coming in. But another question is about static analysis and whether or not static analysis would have been effective in identifying your techniques for bypassing their checks. Because, I think, in your paper you mentioned that Amazon didn't -- or the architecture prevents static analysis of certified skills or that they don't do static analysis. So with that, should Amazon do static analysis and would that have been helpful to you?

MR. WILSON: So with the current architecture, what actually happens is the back end is completely invisible to the certification team, I would say. Because this is a black box, we still don't have confirmed results. This was a question that we raised to Amazon, and they haven't responded to that yet. So from what we know, they do not have access to the code, the back-end code of the developer. So they have no means to actually check it right now. So all they do is, actually, they talk to the skill, and they look at the responses that are coming in and then they decide.

So yeah, this was one of the solutions. You should be taking permissions from the developer to actually view the back-end code and maybe just block

it from being made. The developers can change it any time right now, and it will be deployed to the live audience. So this is not something that should be done. It should be blocked is what we are saying.

Many other researchers also said this in other papers, that developers should not be allowed to change the code. But until recently, the developers could only change the back-end code. The front-end code could not be changed. But, recently, they have changed that, too. So now, you can update both the front-end code and back-end code without requiring a recertification.

DR. BANKS: Okay. I kind of want to make that point clear. So it sounds like what you're saying is that there's a pretty significant blind spot that Amazon has for third-party code, in that the certifiers within Amazon cannot actually see the code. And if the developers make modifications to that code, that does not have to get recertified.

Can you really make clear the significance of that blind spot in terms of what the potential vulnerabilities that can arise from that are?

MR. WILSON: So with the blind spot, what the problem is -- like even if you create a chat bot that can go test the skill for 1,000 times, the

malicious thing can happen on the 1,001st session. So you can't actually find it. So unless you have the back-end code, you can't actually find out all the policy violations.

And, again, because of the skill certification vulnerability, even if the certification system is really good in detecting all these problems and rejecting all the skills, a skill can actually pose as a good skill initially, get it certified, change the back-end code, completely change its functionality, and then just -- yeah, it doesn't make sense.

DR. BANKS: I see. So that sounds like maybe an opportunity for regulators to step in and say that if you're going to offer skills or that any organization will offer skills, then they should be able to have access to the code in order to do a comprehensive analysis before it's made public.

MR. WILSON: Yes. Again, like I said, this is still an assumption because that's what we know, that it's a black box. We have contacted Amazon, and they haven't responded to us about this yet. But from what we saw in our experiments, we can say that they are not looking at the code. Otherwise, they would have definitely found some of these mistakes.

We usually used to name their variables as first name, last name. And if they just look at the code once, they can see we are collecting the full name of the user. So if they actually looked at the code, we would say that they would have definitely found this. But, yeah, still, this is an assumption.

DR. BANKS: Okay, I understand. Let's see. We have a question from the audience.

Can you elaborate more on what you found with the skills in the kids category? The question asks whether or not there were specific violations of COPPA. I don't know if you're familiar with the details of COPPA. But even if you're not, can you talk about what you did see, specifically, in terms of what type of information was being collected within kids skills and what you perceive the violations might have been?

MR. WILSON: So I think, regarding the collection of data, it was mostly collecting either the device location or the user's location or the user's name. So I did see a lot of story skills asking for a name to personalize the stories. It just makes it interesting, I would say, to get a story based on you as the character. So that was mostly it.

But, again, we tried providing the full name

and the skill [indiscernible] with it. But we do not really know if the skill was actually collecting the full name because, I think, personal information is going to be the full name, according to COPPA, and not just the first name. So like I said, since we don't have access to the code, we don't really know what they are collecting. Are they keeping both the first name and the last name or are they just taking the first name, all those kind of things.

DR. BANKS: I understand. So it sounds like that's an area for some closer analysis.

MR. WILSON: Yes.

DR. BANKS: So thank you very much, Christin, for your work.

Next, we'll have Aerin Zhang. She's here to present CMU's research into consumer attitudes about video surveillance and facial recognition.

Welcome, Aerin.

MS. ZHANG: Thank you, Lerone. I appreciate this opportunity to present our research at PrivacyCon.

So today, I will present our work on understanding people's privacy attitudes towards video analytics technologies. This work is part of the Personalized Privacy Assistant project. There were 17

million surveillance cameras in the US in 2018, and if that number is not impressive enough, 1 billion cameras are expected to be deployed globally by the year 2021.

The massive amount of video data captured by these cameras motivates video analytics technologies, which use computer software to automatically process and understand videos. Such technologies have been greatly improved due to recent events in deep learning and computer vision, and they are becoming increasingly sophisticated. Such software can be easily applied to real-time IP cameras or store footage from any cameras. Those analyses often happen without subject's awareness or consent.

Important information about the data collection, like how long the footage is retained, whether the information could be shared with other entities or the purpose of analysis, is often not available to data subjects. Privacy regulations, like GDPR, include stricter laws to govern the use of video analytics. The regulations require entities that use video analytics notify data subjects and enable them to opt in or out of some practices at or before the point of collection.

But there are several different types of

video analytics technologies today. Facial recognition is the most prominent type and also has several variations. It can identify an individual by matching an image of a person to a database of known people. There's also anonymous face detection that can be used to estimate demographics of the person.

Another type is facial expression recognition that detects individuals emotions. Other than facial recognition, scene detection is also one type of video analytics. This image shows how the software is analyzing the video feed to count the number of passengers in the subway compartment.

Next slide, please.

The gap between the current disclosure practices and the requirements of the regulation draw our attention to the lack of guidance on how to do a better job at communicating these data practices and what choices to expose to data subjects. In order to facilitate appropriate notice and choice about these different types of data analytics deployments, we first want to understand people's privacy expectations and preferences with regard to these deployments.

We asked the following research question. Do people know about these deployments? And how do people feel about them? Especially, we are interested

in people's surprise levels, whether to expect these practices at certain places or not, their comfort level, their notification preference, meaning whether they would like to be notified or not, and if yes, how often they want to be notified. We're also interested in whether people would allow or deny those practices if given a choice.

With these research questions in mind, we designed an experience sampling study. So the experience sampling method is a longitudinal research methodology which enables us to engage and survey participants in the moment as they go about their normal daily lives. As a result, this method allows us to collect higher quality, more ecologically valid research data than static online surveys. In total, we collected detailed responses for more than 2,300 deployment scenarios from 123 participants.

When recruiting, we tried to avoid convenient sampling of students and tried to reach out to the local community. We ended up with a rather diverse sample. Here's a pie chart showing the different occupations of our participants. This study is approved by Carnegie Mellon's Internal Review Board and by the following agency's Human Subject Protection office.

Next slide, please.

So we first did an extensive survey of news articles about real-world deployments of video analytics technologies. We identified four major categories in a variety of contexts. The first category is for security, which includes automatically detecting petty crime scenes, like pickpocketing, break-ins, or using facial recognition to identify known criminals and bad actors.

The second important type is for commercial uses. It's been used to count the number of people in a facility in order to optimize operation, like staff management. Or it's used for targeted advertising based on demographics, individual profiles, or reactions when people are looking at items. Yes, you can be advertised based on what you look at and your facial expression. It has also been used to rate people's engagement at museums, movie theaters, and comedy clubs.

The third key usage revolves around identification and authentication. Facial recognition can be used to replace work IDs, membership, and loyalty cards. It has been used to track attendance at gyms, schools, workplaces, and even churches.

And the last category of uses is more

advanced. Facial recognition and emotion analysis can be used by health insurance providers or hospitals and doctors to make health-related predictions, and by employers to evaluate employees' performance and monitor their productivity.

In total, we identified 15 unique purposes and a baseline purpose, which involves only generic surveillance with no video analytics.

Here, I'm going to briefly explain the study protocol. So participants download and install the study app on their own Android devices, and the app first asks them to fill in a pre-study survey. After that, participants were instructed to go about their regular daily activities, and the app sent them push notifications prompting them to complete a short survey based on their current locations.

Because the GPS location is not precise enough indoors, participants first confirmed the place they were at by selecting from a dropdown list of nearby places. Then the app displays video analytics deployment scenarios relevant to the place they were visiting, and then they answer four in-feature questions about their surprise level, comfort level, and notification preference with regard to this scenario.

On the days participants receive push notifications through the app, they also got an email in the evening to complete a daily summary. The summary asked participants to revisit notifications they received during the day and to provide additional responses. The process will happen for 10 days, and participants finish the study with a post-study survey.

Now, I'm moving on to some of the results we found in our study. Due to the length of the presentation today, I'm only showing some of the results and more can be found in our paper. So this slide shows a summary of participants' comfort levels organized around 16 different purposes we previously identified. It was clear to us there is no scenario where everybody feels uniformly about. People's responses vary greatly for each purpose. For instance, scenarios related to security appear to surprise participants the least. Close to 72 percent would feel somewhat or very comfortable about these scenarios.

On the other end of the spectrum, we observed considerably less acceptance by event scenarios, like health and productivity predictions, where only 70 percent feel somewhat or very

uncomfortable.

Participants are least comfortable with employees making predictions about their work productivity. So after the 10-day study, 75 out of 123 participants grew more concerned about these practices. Eighty percent of these 75 participants developed stronger awareness of the possible deployment of video analytics technologies as they received notifications on their phone every day. They were not aware that video analytics could be used for so many purposes at such a diverse set of venues and with this level of sophistication.

One participant commented, "Some of the scenarios and growth of the technology you mentioned, I had never considered. Freaked me out."

Twenty-seven percent emphasized the privacy issues of these technologies, like the lack of notice or consent. Twenty-five percent expressed concerns about specific usage of these technologies. One said, "I didn't realize I could be marketed to based on what I'm looking at in a store....I found this whole practice disconcerting." Four percent were worried about implications like how the data is shared, what could be inferred from the data, and the potential abuse.

Next slide, please.

So just to give a complete picture, I'm going to show some opinions of participants who stayed equally concerned or actually grew less concerned over the course of the study, even though they are minorities. So 27 percent of them claim they are already familiar with these technologies. Twenty-three were not bothered by the practices. They said something like, if you're not a criminal, you shouldn't be worried about facial recognition. And 21 percent expressed some level of resignation, describing the technology as ubiquitous and out of their control. Fifteen percent did not believe that the scenarios showed to them were real. And 13 percent who learned the benefits of these technologies become more accepting.

Now we move on to the next slide, showing results on those notification preferences. We asked participants, how would you want to be notified? The choices range from notify me every time to do not notify me. Again, we observed that people show diverse notification preferences.

This graph shows how their preferences changed before and after the study. More than half ended up with different preferences, and the majority

are looking for some type of selective notification solution instead of being notified every time.

Next slide, please.

So interestingly, we observed that people grew more concerned in general, but opted for less frequent notifications as time passes. This change in preferences is attributed to some level of privacy fatigue as people got a better appreciation of the number of times they are likely to be notified. So one participant described their fear for privacy fatigue as they received many notifications.

Next slide, please.

Even with our 10-day study, we already observed privacy fatigue. So remember the regulations which expect people to manually opt in or out of video analytics each time they encounter such functionality, but because of the increasingly widespread deployment of those softwares, this could result in an unrealistically high number of privacy decisions.

So the natural question to ask is, how could we reduce user burden and assist users in making privacy decisions? So I want to first provide some context of how obtaining consent works with video analytics data collections. But there are some recent technical advances that made it possible to obfuscate

people's faces in real time, allowing people to opt out of video analytics.

There are also academic efforts to build a privacy infrastructure and a privacy assistant app for Internet of Things. Such an app running on people's smartphones would alert users of nearby IoT sensors, for example, cameras with video analytics software enabled, and present them with potential choices, like opt in or opt out.

However, with all the efforts, the high user burden remains a problem. So with the data collected from our study, we're able to use clustering techniques to reduce user burden. We first grouped like-minded users to generate privacy profiles and then leverage clustered profiles to make predictions of people's allow or deny decisions. So using this method, we're able to predict 94 percent of the allow/deny decisions with 89 percent accuracy.

Next slide, please.

So it is worth taking a closer look at the clusters of the like-minded subjects identified by our clustering algorithms. This graph shows privacy profiles of six clusters. Each cell represents whether people in this cluster allow or deny data practices for a specific purpose. The color blue

means "allow" and the red means "deny." Darker colors indicate a stronger cluster consensus.

We see cluster 1 and 5 are polar extremes. Cluster 1 are privacy conservatives while cluster 5 are mostly unconcerned. They merely allow data collection for all purposes. The other four profiles are more nuanced, with a mix of red and blue cells. For example, cluster 4 mostly allowed data collections for security purposes and denied others.

So I'm going to summarize what we have learned. Through the study, we observed people's lack of awareness of video analytics and their desire for greater transparency. The current practices of notifying people by putting signage that states "this area under surveillance" is not sufficient and also not compliant with regulations. People want to know when they are subject to video analytics technologies.

We also saw how participants can be overwhelmed by the number of privacy decisions they might encounter as required by regulations, and they are looking for selective notification solutions.

Lastly, we demonstrated the feasibility of reducing user burden through machine learning by predicting the majority of decisions with high accuracy.

So this concludes the presentation.

DR. BANKS: Thank you very much, Aerin.

Again, for the audience, if you have any questions for Aerin about what consumers have to say about video surveillance, please send your questions to privacycon@ftc.gov right now. We have a few minutes left.

All right, so thank you, again, Aerin, for that work, and thank you for spending time talking to consumers. I'd like to spend some more time doing that, too.

Let me ask you one question as we wait for some audience questions to come in. In other contexts and other privacy research, we often hear the term "privacy paradox" thrown around, and other people use it, too. In your interactions with consumers, did you observe any privacy paradoxes or any counterintuitive behavior or responses?

MS. ZHANG: I think the privacy fatigue that we described had something to do with the privacy paradox, but the privacy paradox deals with actual behaviors that we, in the study, did not really measure. So we are basically asking their opinions. So the privacy paradox describes the discrepancies between the actual behaviors and their saying that

they care about privacy.

DR. BANKS: I see. I have one other question for you. Oh, in your study, comfort is strongly correlated with allow and deny decisions from consumers. Is that right?

MS. ZHANG: Yes.

DR. BANKS: And were you able, based on consumer responses, to get an understanding of what things companies could do that were the most effective at increasing their consumer's comfort level? Essentially what I'm asking is, are there things that companies can do to increase consumer comfort and reduce surprise?

MS. ZHANG: I think by listing a lot of the attributes, like the purpose for which this data is collected and for how long the data is retained, by disclosing those informations, companies will receive more acceptance from users because we have seen that people are -- once they know the benefits and they know the whole picture of how facial recognition is used, they become more accepting to some level.

DR. BANKS: I see. And, presumably, that would be outside of the privacy policy that they're less likely to read and maybe somewhere prominent and easy for consumers to understand, hopefully, right?

MS. ZHANG: Yes.

DR. BANKS: So thank you, again, Aerin.

And I'd like to thank all of our researchers today. You're doing great work, and it's informed me a lot today. And I hope our audience got as much out of it as I did. So thank you very much. And I think we have another panel coming in immediately after us. Thank you, again, and thank you to the audience for your attention.

SESSION 5: INTERNATIONAL PRIVACY

MR. WOOD: Okay. Welcome to the fifth panel of this year's PrivacyCon. The topic of this panel is international privacy.

I'm Dan Wood. I'm an economist in the Bureau of Economics at the Federal Trade Commission and the Division of Consumer Protection.

With me are four panelists. The first is Guy Aridor. He's an economics PhD candidate at Columbia University. And the research he's going to be talking about today is about how the European Union's General Data Protection Regulation, or GDPR, how its opt-in requirement affected the mix of consumer data observed by intermediary web services.

The second panelist is Garrett Johnson. He's an Assistant Professor of Marketing at Questrom from School of Business at Boston University, and he's going to be talking about how GDPR affected market concentration among web-technology vendors.

Our third panelist is Jeff Prince. Jeff is a Professor of Business Economics and Public Policy at the Kelley School of Business at Indiana University. He's also the Harold A. Poling Chair of Strategic Management and Co-Director of the Institute for Business Analytics at the Business School. And his

research he'll be presenting is about measuring individual's valuation of online privacy across countries and also across privacy domains.

Christine Utz is our last panelist. And she's a PhD student at the Chair for Systems Security at Ruhr University Bochum in Germany. And her research is about how design choices in GDPR consent notices affect how users interact with those notices.

So without further ado, I'll turn it over to our first presenter, Guy Aridor.

MR. ARIDOR: Yeah. Thanks, Daniel.

So today, I'm going to talk about the effect of data privacy regulation on the data industry. This is joint work with Yeon-Koo Che at Columbia and Tobias Salz at MIT.

So the fundamental tension at the heart of data privacy regulations is that, on the one hand, consumers increasingly want control over the data that firms collect on them. It's just been amplified in recent years as a result of a number of high-profile data breaches, as well as an increase in the scale and scope of data that firms collect on consumers, rendering consumers to be unable to understand what kinds of data is collected on them.

On the other hand, firms are becoming

increasingly reliant on this consumer-generated data. There's a worry that such privacy regulation might impact their function.

There's two main uses of data in the digital economy. The first is that this data is the fuel behind a lot of the machine-learning technologies which are becoming more and more deployed in the digital economy. And, second, they're crucial for the targeted advertising, which is how many websites derive their revenues.

So next slide.

What did we do in this paper? So we looked at the European Union's General Data Protection Regulation. In particular, we focused on the consent aspect of the legislation, which gives consumers additional control over the data that firms collect on them. And what we try to answer in this paper is try to empirically say something about the tension I previously discussed.

First, we try to ask, do consumers make use of the privacy means provided by GDPR? And then how does this impact the overall pool of data that firms observe and how does this materially impact the firm's ability to predict consumer behavior and accrue advertising revenues?

Next slide.

So the empirical setting for this paper is the data is provided to us from a third-party intermediary in the online travel industry, which spans the majority of this industry across the globe. This intermediary's sort of an ideal setting to study the consequences of the GDPR on data-reliant firms for several reasons. The first is that the data that this firm collects is directly at the heart of the consent portion of GDPR, such that properly implemented consent should allow consumers to opt out of data collection from this intermediary.

Second, the primary business of this intermediary is to collect user search and purchase histories and to predict whether or not consumers are going to purchase a flight or hotel, and then conditional on this prediction, show some advertising, which is how most of their revenues come about.

And so in particular, we observed the following. We observed a high degree of consumer search histories, we observed advertising revenues, and we observed the output of the proprietary machine-learning algorithm, which are all the necessary outcomes to talk about the original tension.

Next slide.

So what do we do in this paper in terms of our empirical strategies? So we used a relatively standard tool from economics, known as difference-in-differences, which allows us to get at the causal impact of the policy. And, in particular, our treatment group here are the travel agencies in major European countries and the control group here are travel agencies in non-EU countries. And our analysis revolves around the GDPR implementation date, which was Friday, May 25, 2018.

And it's important to point out that our specification will allow us to look at the causal impact of the policy overall and not necessarily on particular manifestations of the policy.

And so we look at the period from beginning of April 2018 until the end of July 2018. And I'm going to report two specifications here. One is just going to give the overall causal effect over this time period and the second is going to give a time-varying one.

Okay, so next slide.

The first thing we do is try to use our specification to indirectly measure consumer opt-out. So it's important to understand how GDPR opt-out manifests itself in the data that we see from the intermediary. In particular, when a consumer opts out

of data collection, their data is not showing up in the database at all. And so what we do is we measure opt-out indirectly by estimating the difference between the observed users and the number of users that would have been observed had GDPR not been around.

And so to measure this, we're going to consider the following outcome variables. One is the total number of unique cookies and the second is the total number of recorded searches. And so, again, this will allow us to indirectly have an estimate for how many consumers opt out, but it's also going to tell us about how the overall scale of the data that the firm sees changes.

Next slide.

So this is the time-variant specification. So you can see a sharp drop at week 22, which is exactly the week of GDPR implementation. And you see a steady decline of roughly 10 percent to 12 percent, and this is consistent across the different outcome variables.

Next slide.

Okay. So now what do we turn to? So we've established that there is a 10 percent to 12 percent drop in the total number of users that firms observe.

Now, our next question was, are there any changes in the composition of the users? And to do this, we look at a measure of persistence. So what we do is fix a website, J, and fix a week, T. And then we collect all the cookies that that website observes in time T. And then we ask, what fraction of these are still around one week later, two weeks later, three weeks later? And we ask, what happens after GDPR? Are the resulting consumers more persistently identifiable?

Next slide.

So what we find is, again, you see a sharp increase at the onset of GDPR. And the effect size of this is roughly around an 8 percent increase in persistence. So I'm not going to go into the details here, but in the paper we sort of were curious, what's the mechanism behind this increase in consumer persistence?

So next slide.

So our main conclusion is really that it's important to distinguish between different means of privacy. So before GDPR, consumers could do things such as delete their cookies or use private browsing. And what would happen there is that the consumer's data would end up in the firm's database, but with a new identifier, whereas under GDPR, such data is

eliminated completely. And so a substitution between these may lead to a different data-generating process and longer consumer search histories.

Next slide.

So this figure sort of illustrates exactly what I'm talking about. So in the far left, you can see the identifier column is the identifier that the intermediary is observing for a particular user. And then in the other three columns, you're going to see consumer histories.

So if you focus on the first panel, the full-visibility panel, that gives the true data, that if the firm perfectly observed everything is what they would see. And so what you see is four distinct consumers with distinct search and purchase histories.

In the middle panel is the obfuscation regime, which is the pre-GDPR. And so let's suppose the first three consumers, they don't change their behaviors at all. But suppose the fourth is privacy conscious, so periodically deletes cookies or uses private browsing. So now what happens with this guy is his identifier is now partitioned into two users. So the intermediary thinks that it's saved two people, but it's actually one person. And as you can see,

consumer 4 and consumer 1 now have identical histories, and consumer 5 consumer 2 have identical histories.

Now, what happens under GDPR? So this privacy conscious consumer can now opt out of the data. So you see that the firm only observes three users now, but they arguably have cleaner identifiers.

So there's sort of two takeaways here. One is moving from the second panel to the third panel is going to mechanically increase persistence. And the second is that it might actually help the firm predict consumer behavior because they have clear user histories. And we have an extended discussion of this in the paper.

Next slide.

So now, what we want to do is we want to look at what happens to advertising revenues? And so all we're going to report here is just the results of our specifications without many more details, but it's important to contextualize the advertising setting here. So we're not thinking about behaviorally targeted advertising where advertisers are bidding directly on consumer histories. We're thinking here of keyword search advertising, so similar to Google-sponsored search. So an example is advertisers

are bidding on consumers who search from a flight from New York City to L.A. So any changes to bidder behavior are reflecting the average value of a consumer.

Next slide.

Okay. So these are results using the same specification as before. So first, what we find is the total number of advertisements that get clicked on has a similar effect size drop as we saw before, which is roughly 13 percent.

Next slide.

So we look at revenue. And so revenue is a bit interesting. So I don't report the time-varying graph here, but what we see is there is a sharp decline at the onset of GDPR and a slight increase afterwards. So we find a negative-point estimate, but it's relatively imprecise and statistically insignificant. And the reason why, we think, is because -- if we go to the next slide -- the average bid for a consumer -- and this is for GDPR -- actually increases, which points to the fact that advertisers had a higher average value of consumers after GDPR. And so this partially offsets the loss from opt-out but not completely.

Okay. So that was in advertising. Then,

finally, we're going to turn to prediction. So yeah, so we should be on the consequences for prediction slide.

So we asked, what's the impact on predictability of consumer behavior? So I think this is interesting for two reasons. The first is obviously, you know, directly impacts from revenue in terms of the amount of personalization and product quality that they're able to offer. But I think the second reason why one should be interested in consumer behavior -- or in predictability, is also from a consumer privacy standpoint because it's privacy in the modern age, at least colloquially. It's not just about what kinds of information do people have about me, but sort of what can firms predict about my behavior? And so we think that this exercise is interesting to look at from that light as well.

And so the key idea here for interpreting our results is that the privacy decisions of others affect my predictability, right? So all our data gets pulled together. So if the firm has less data from opt-out, this is going to impact the firm's ability to predict for a consumer that opted in.

But the second thing is, if you think about our results from persistence in the sort of stylized

diagram I showed you before, it might actually be possible that the substitution from cookie obfuscation to opt-out actually leads to cleaner identifiers than exerts an externality to consumers by making them more predictable.

And so in our setting, we can use the same special specification as before precisely because the firm trains their prediction model for each site on whatever data they accrue from that site, which means that changes in data from one website don't affect the firm's ability to predict on another website.

And so what do we find?

Next slide.

So we do a short-run exercise, which we just put it through the same specification as before. And we find that there is slight improvements in prediction, but the big takeaway, I think, for us, the prediction didn't get substantially worse, according to the measure utilized by the intermediary.

Now, we were a bit worried that the short-run effects might not give enough time for the intermediary to adjust its prediction rhythm. And so what we do is we do a back-of-the-envelope logarithmic exercise where we sort of take the changes we saw from our earlier difference-in-difference estimate in the

change in the overall scale and longer consumer search histories and we asked how should those affect prediction. And what we find is a roughly similar result.

Okay, next slide.

Well, what did we do today? So we looked at the impact of GDPR on a number of different outcome variables. I think there's two high-level takeaways that are very closely related. The first is I think we highlight how government-mandated privacy protections do interact with other privacy needs. And this can be important for understanding the value of such regulation.

And second is that we highlight that consumer privacy decisions have externalities on other consumers, which is not something that legislation such as the GDPR really thinks about. And, finally, just in terms of welfare, going back to the tension we talked about before, do consumers benefit? Privacy conscious consumers clearly do. For the others, it depends on the alignment of the preferences of firms and consumers. Do firms suffer? Firms lose a significant number of consumers from opt-out, but remaining consumers are higher value, so it's not wholly negative. And, finally, the ability to predict

is not substantially worse.

Okay, thanks.

MR. WOOD: Okay, great. Thank you, Guy.

Next up is Garrett Johnson.

DR. JOHNSON: Thanks, Dan.

I'm honored to be back this year to present our second GDPR paper with the same set of coauthors. It's with Scott Shriver at Boulder and Sam Goldberg at Northwestern.

Next slide, please.

Our main research question is, can privacy policy hurt competition? Now, there's a theoretical tension between privacy and competition policy, but this claim lacks empirical evidence. One reason for this tension is economies of scale, that larger firms may have more resources to comply with regulation. We propose a novel mechanism, though, which is B2B choice of data vendors. That is, if privacy regulation pushes firms to limit data sharing, firms may prefer to keep their larger vendors because these vendors have better products.

Next slide.

So as in the last talk, we're studying the GDPR, which is a landmark privacy policy and a leading example to the world. And we, too, are going to use

its enforcement deadline of May 25, 2018, as an event study.

Now, the GDPR is very complex, but its many elements contribute to increasing both the logistical cost and legal risk associated with processing personal data. And this is going to have important consequences for the web. We study the technology vendor industry that provides an ecosystem for the web to thrive. Specifically, these vendors help websites to monetize themselves with ads, to load and share content, as well as measure and optimize site traffic.

Now, in order to provide many of these services, vendors often have to share what the GDPR considers to be personal data. And as a result of this, the industry has faced intense regulatory scrutiny with at least three EU countries releasing major reports or statements criticizing the industry. But, so far, the regulators have not issued any fines.

Next slide.

So today, I'm going to briefly discuss our data and then discuss our results in three stages, talk about the GDPR's impact on vendors, its impact on concentration, and then differences by website.

Next slide, please.

So we begin with our data.

Next slide.

When you visit your favorite website, your browser interacts with the first-party domain at that site. So in this example, I visited TheGuardian.com, and my browser's interacting with that domain.

Next slide.

At the same time, your browser is interacting with potentially dozens of third-party domains owned by vendors, selected by the website to provide these services. So here, I've used an extension for Chrome called "Disconnect" that allows me to visualize these vendors. And you can see many familiar logos, including Facebook, Twitter, and Yahoo. Many of these vendors are helping to monetize The Guardian with ads.

Now the GDPR is challenging the study because, normally, we cannot observe how firms use and share personal data. However, in this instance, the function is being outsourced to the browser, so we are able to observe a website's network of vendors.

Next slide, please.

Our data collection precedes as follows. First, we use a VPN service to simulate ourselves as originating from within the EU, specifically from

France. Second, we use a specialized piece of software, called webxray, developed by a researcher at CMU named Tim Libert. And that allows us to record all the third-party domain interactions when we visit a website. And, finally, we repeat this for 28,000 top sites regularly throughout 2018, and these sites in our data sample are the top 2,000 websites in each of the 28 EU countries, as well as the US, Canada, and globally.

Next slide.

So for our results, we begin by looking at the GDPR's effect on vendor use.

Next slide.

So this figure shows the average number of vendors per site over 2018. And immediately prior to the GDPR, sites used 14.4 vendors on average. One week later, this falls to 12.4 vendors, which is its lowest level. And this is a 15 percent reduction in vendor use, which we refer to as the short-run effect of the GDPR. Now, obviously, we would have preferred to collect a longer pre-period, but we know from auxiliary data and related research that the pre-trend here is flat.

Furthermore, websites appear to have waited to the last minute to make changes to their website,

which is why three-quarters of the drop in vendors happens within just a few days of the enforcement deadline.

The reduction in vendor use is short-lived, however, and erodes by the end of 2018. The post-GDPR growth may just arise from a dynamic market that's expanding over time, but I'm going to show you some evidence later that this growth is consistent with sites' beliefs about enforcement falling over time in the absence of fines.

Next slide.

So now, we've seen that vendor use falls post-GDPR. We're now going to turn our attention to concentration.

Next slide.

To fix ideas, we know that vendor use falls post-GDPR, and most vendors are actually worse off post-GDPR, in terms of the number of sites that they're working with. But here, we're instead asking a different question, which is, do the larger vendors get a larger share of the smaller pie after the GDPR?

Next slide.

Now, in order to measure market concentration, we begin by defining market shares, and our market-share definition relies on reach, which is

just the number of websites that use a vendor. So in the sidebar example, you can see that Google Analytics has a reach of two sites and Adobe Analytics has a reach of one site. And then to calculate relative market shares, we just take the vendor's reach divided by the total reach so that in the sidebar example, Google Analytics has two-thirds market share and Adobe analytics has one-third market share.

Now, note, we are not observing any revenue or costs cost that's changing hands between vendors and publishers. We're only observing these vendor links. Our measure of concentration then, which is the Herfindahl-Hirschman Index, or HHI, is just the sum of the squared market shares. And this index is going to be increasing and the level of concentration, so zero is a perfectly competitive market and 10,000 points is a monopoly. Because as a relative definition of HHI, if all vendors fall by the same percentage, then the relative HHI is going to be invariant.

Next slide, please.

Now, we plot relative HHI over time, and we see the evolution of concentration is the mirror image of the average number of vendors. In particular, concentration rises 17 percent post-GDPR in the short

run, and we think the short run is informative, certainly directionally so. When evaluating a policy that has not been enforced, we think that the period where beliefs about enforcement are higher is more relevant. And we'll provide some more evidence about that belief later.

We conclude that the GDPR increases concentration. And the intuition for this is that vendors with large shares have large shares because they provide greater value, whether it be because they have lower costs, deliver greater revenue, or have superior privacy compliance. Whatever the reason, websites prefer to retain the large vendors when the GDPR purchase websites to reduce data sharing, which is why we see this concentration increase.

Next slide.

So in the last slide, I showed you aggregate HHI. But we want to define markets more narrowly, and we do so by using an external categorization based on the type of service that vendors provide. And, now, we can see that the concentration increases in the top four categories that represent over 94 percent of categorized vendors; in fact, the largest category is advertising and, here, we see concentration rise 25.3 percent. And in the next three categories of hosting,

audience measurement and social media, we see concentration is still increasing between 2 and 6 percent.

Next slide.

Now, I want to quickly examine one of our three extensions that illuminate the mechanism for the concentration result. We consider the role of the big two companies -- Google and Facebook -- and there are many associated vendors. As before, with all vendors, we see that HHI rises 17.3 percent. However, when we exclude the vendors associated with the big two, concentration actually falls 6.2 percent. So maybe we need to update the old adage that nobody gets fired for hiring IBM to also include Google and Facebook.

Next slide, please.

Now, I want to quickly illuminate some differences by website that tell us something about the economics of how websites are making decisions under the GDPR.

Next slide.

To begin, I want to break apart the short-run drop in vendors by characteristics of the sites.

Next slide.

For instance, here, I break apart sites by

the share of traffic they get from EU users. We can see that sites with between 90 percent and 100 percent of EU users, on the right-hand side of the figure, drop a little over two vendors on average in the short run, where sites with between zero and 10 percent, the lowest estimate on the left-hand side drop a little over five vendors on average in the short run.

We think that this reflects the incentives in the GDPR that place a 4 percent penalty on global revenue. This means that sites with few EU users have relatively little to gain, in terms of revenue from the EU, but relatively more to lose from a penalty on their global revenue. The GDPR incentive then has the perverse effect that sites with the greatest share of EU users do the least to cut vendors.

Finally, notice the discontinuity for sites with zero percent EU on the far left-hand side illuminated in orange. Many of these sites are actually not subject to the GDPR and, therefore, do not need to make changes.

Next slide.

Finally, we examine the post-GDPR evolution of vendors in 2018.

Next slide.

One of the things we noticed after the GDPR

is that the average number of vendors grew slowly in countries like Denmark and the Netherlands, but grew rapidly in countries like Bulgaria and Poland. Now, the GDPR is meant to harmonize regulation within the EU, but it's still enforced, in part, at the country level. So we found a survey measure from the EU that measures regulatory strictness specific to data protection. We found that regulatory strictness is negatively correlated with the post-GDPR growth in vendor use. This suggests that site beliefs about the probability of GDPR enforcement help to explain the 2018 evolution in vendor use.

Next slide, my last slide.

We started out with a theorized tension between privacy and competition policy and, today, we're able to show you the first empirical evidence of this tension. The GDPR had its intended consequence of decreasing web-technology vendor use and its associated data sharing. But it had two unintended consequences. First, we saw an increase in vendor concentration and, second, we saw that sites with the most EU visitors reduced vendors the least, an apparent side effect of the GDPR's penalty design.

Thank you.

MR. WOOD: Well, thank you, Garrett.

Our third presenter is going to be Jeff Prince.

Jeff, you have to unmute.

DR. PRINCE: Thank you. There. I'm unmuted now. Perfect. Even better.

So thank you again to the FTC organizers for the opportunity to speak and for Dan for moderating this session.

This is joint work with Scott Wallsten at the Technology Policy Institute. And we received financial support from the Inter-American Development Bank for this work. So we're looking at how much is privacy worth around the world and across platforms.

Next slide, please.

So prior speakers have already kind of highlighted this with the GDPR. But this is across many countries around the globe. Governments around the world are grappling with data privacy policy. And as economists, we're always thinking about the tradeoffs of policy. So at a very rough level, we can think about balancing privacy preferences for the citizens with the benefits from use of the data.

And one thing that has been emphasized in many places is that it's particularly difficult to measure the privacy preferences. And that's something

we're trying to get at with this project.

Next slide, please.

So what we do is we use conjoint survey techniques to measure the willingness to accept for online data information. We compare and contrast those willingnesses to accept for a range of data types across different countries, six of them, Argentina, Brazil, Colombia, Germany, Mexico, and the United States. And then we do it within four different platform contexts, so with your bank, with your carrier, with Facebook, and then with your smartphone.

Next slide, please.

And so for the surveys that we put out, they offer choices with different levels of data privacy and with different monthly payments associated with those different levels. And we had specific reasons for doing it this way. This is rooted in real data markets. So one example, if you go to DataCoup.com, you can go and get money right now for your data, so easy money hanging out there for all of us. So you can offer up some of your data and they'll give you monthly payments for access to that.

So the goal of designing our surveys was to try and make them as realistic of an actual choice

that someone would make with regard to their privacy. And the design of these surveys is well suited for measuring tradeoffs, which is what we're interested in.

Next slide, please.

And so here we have an example for Facebook, just to give you a sense of what we're talking about. So here a respondent is presented with four different options. And each option has different information that's being shared and then monthly payments associated with it.

And so to give you a very clear example of a tradeoff, if you look at option two versus option four, the information that's being shared is the same except for option two, you're not going to be sharing your texts. Option four, you will, but with option four, you get paid more. And so it allows people to make the tradeoff with, you know, is that additional money worth it to me to give up that information or not? And so then that choice helps us to pin down how people value different types of privacy.

Next slide, please.

And so we used the firm called Dynata. Back when we used it, they were referred to as Research Now. They administered these surveys online for us.

We had 325 surveys for each type, and a type being a country, platform, and then we also included a randomized prompt, where you either got it or you didn't. And the prompt essentially indicated to people the value of sharing their data.

And so we included this to try and get a sense as to how flexible people's preferences were. So are they malleable to being prompted about the value of data or not? And so I'll speak to those results as well when we get to them.

Then we had basic screenings on age, so they had to be adults. And then they needed to have existing accounts for all but the banks.

Next slide, please.

And so for the analysis, we analyzed the choices that people made, choice data that came back through the online surveys using standard conditional multinomial logit, going back to McFadden's utility-based formulation. And the real basic idea is take the willingness to accept -- to get the willingness to accept for data privacy, just take the utility for data privacy and divide that by the utility of money. And this is going to be measured in monthly payments.

So next slide, please.

So the next few slides present some of our results and tables. This one is a very aggregated result. So what we have here is, averaged across platforms and countries, what is the willingness to accept in dollars per month, using a purchase price parity index to compare across countries for the different types of data privacy.

And a summary of what we find here is that, A, there is a lot of variation, as you can see. The financial information is particularly well guarded, so those are on the higher end. Also, fingerprint information had high WTA, along with text information and contacts.

Next slide, please.

This one here has a lot more information. I know it's a lot to process all those bars, but let me just highlight a couple main points from those. One is, if you notice, the orange bar is Germany. That one is almost always the highest one for each of them. I think for 8 out of 10 it's the highest.

Another, I think, broad point to take away from this is that for the rest of the countries, there's a lot of similarity and there's not a fixed ordering. So it's not as clear that one country generally has higher willingness to accept than

others, with the exception of Germany.

Next slide, please.

And then here's virtually everything we have in terms of averages broken down across all the dimensions, so across platform and country for the different types of online information. Again, I know it's a lot to try and swim through, but let me just highlight a couple high-level observations.

One is, if you look at the bars across the different countries for the different data types, again, you see a lot of similarities in even the absolute values but even more so the relative values. So if you look at the wireless, upper left quadrant there, the red bars are always the highest, the orange bars are almost always the second highest, followed by the brown. And so there's a lot of consistency in the relative preferences for different types of privacy across the different countries.

Next slide, please.

So some key takeaways. Overall, what we find is relative values are quite similar across our six countries. And then another set of results that were harder to present in tables but are also in the paper is that, at a rough level, the within-country variation -- so if you think about the distribution of

preferences for a particular type of online information with regard to privacy within a country, it is quite similar across countries. So how spread the preferences are, the WTA measures are, across citizens within a given country is similar across countries.

And so some key takeaways from those results is that public and private policies may want to offer similar relative protections, at least to the extent that these countries are representative. If you think about tiered protections, where you think about private firms could offer different levels of protection for different prices, those would likely have comparable appeal across countries. And then the distribution of support for public policies is likely to be similar across countries.

Next slide, please.

And then as I highlighted, Germany is different, at least within the set of countries we looked at. They're different overall and with financial information in particular. So a key driver of their outlier status is with regard to financial information. There is a very high willingness to accept in terms of giving up information that has to do with financial specifics. And then, also, with

regard to the distribution of preferences within Germany, they appear to be the most homogeneous. So the spread of WTAs for different types of privacy is notably smaller for Germany than for the other countries we looked at.

Next slide, please.

And then a few other results that I think are worth highlighting that we found. When we break it down across sex, women versus men, the willingness to accept for women was notably higher than that of men for different types of online privacy, often by about an order of two times. If you go across age, for the older cohort versus the younger, the willingness to accept was substantially higher, often by two, three, or even four times as much. Income, though, does not predict the willingness to accept very well.

And then last but not least, with regards to that leading statement I mentioned earlier, the preferences did not seem to be impacted by that, which suggests that they're not easily swayed by prompts that one might put out with regard to the value of data and giving up privacy, or its potential value. People's preferences seem to be unimpacted by that. And with that, I will conclude. Thank you very much.

MR. WOOD: Okay. Thank you, Jeff.

Our last speaker is going to be Christine Utz. Christine.

MS. UTZ: Thanks, Dan, and everyone at the FTC, for having me back here at PrivacyCon.

What I'm going to present today is a direct follow-up to our GDPR paper from last year's PrivacyCon. This is joint work with my colleagues Martin, Sascha, Florian, and Thorsten and was previously published at ACM CCS 2019.

So I'm sure you've all -- next slide, please.

I'm sure if you've seen all of these before. These are consent notices, colloquially known as cookie banners. And these are little popup boxes that show up on lots of websites these days. And they inform you of the website's data collection practices and ask you for your consent.

The legal foundation of these notices is the privacy directive of 2011 from the European Union. But especially after the GDPR enforcement date, they've seen a large surge in prevalence across websites. And as I presented in last year's PrivacyCon, we saw an increase of about 16 percent between January 2018 and after the GDPR enforcement

date.

Recently, there have been some vendors of third-party consent libraries that have started to also implement the new CCPA Do Not Sell requirement in these consent notices. So maybe we'll be seeing more and more of these also with the CCPA as a legal foundation.

Consent notices can be arbitrarily complex. So you can have just the basic one with just an allow button, like the dark one. Or you can have a more fine-grained selection, where you can select different categories of cookies, like in the one on the top right corner.

Next slide, please.

So we saw all of these notices become more and more complex. And we can come up with a couple of research questions. So how often do people interact with these notices? Do different changes in the parameters of the user interface of these notices influence what decisions users make? And why do they choose to interact and not interact with these notices? And what do people expect to happen when they allow or deny cookies?

Next slide.

So we decided just to find answers to all of

these questions in the field. So we had the opportunity to team up with a German e-commerce website. And that site has about 20K unique visitors per month. Most of them just google something, and then find an article on the website, read it, and then leave the site. It runs on WordPress and uses common third-party services, like Google Analytics or embedded YouTube videos or a design framework called Ionic.

And we modified a WordPress plugin to display arbitrary consent notices on that website. And with this plugin, we conducted three iterative experiments between November 2018 and March 2019 in a between-subject study.

Before we could get started, we had to evaluate the available design space for the UI of consent notices. So we luckily still had a couple of consent notices laying around from our paper from last year. So we just sampled 1,000 of those and inspected them and identified the design space for consent notices.

Next slide, please.

So we identified eight different UI parameters of consent notices. Three of them are about the relation between the notice and the website,

like the position of the notice on the website, or the size, and whether or not it blocks access to the underlying website. And the other parameters are about the notice itself. So you have the text, whether or not it contains a link to a privacy policy, the general formatting, and then you have the choices offered by the website.

And this is also where nudging and dark patterns come into play, because often the available choices are not presented sort of equivalently, but some of them are highlighted, usually what the website owner wants users to click, like accept to allow cookies.

Next slide, please.

Once we had identified this design space, we designed our study. And in our first experiment, we evaluated the influence of the position of the notice on the website. In the second experiment, we looked at the influences of the different choices offered by the website and nudging. And in our third experiment, we tried to identify the influence of the presence of a privacy policy link, and whether or not the notice uses technical language. By technical language, we mean things like "this website uses cookies" versus "this website collects your data."

Next slide, please.

Our study setup looked as follows. So the user visited the website, and then they were shown one of n consent notices, with n being the number of notices in the current experiment. Our plugin then would log all interactions between the user and the notice, such as clicking an okay button or ticking a checkbox or clicking the privacy policy link.

If the user chose to interact with the notice, we replaced the content of the notice with another notice that mentioned that this is a university study and gave them the choice to either participate or close the notice for once and for all. And if they chose to participate, we just redirected them to our survey. But we were also interested in people who chose not to interact with the notice. And for that, after 30 seconds without any interaction, we automatically replaced the notice with the study invitation. And then, again, if they chose to participate, they were sent to another version of the survey.

Next slide, please.

So these are the results of our first experiment, location. In this experiment, we displayed a binary notice to encourage user

interaction, and we displayed it at six different positions.

Here and in the following, I'll only report interaction rates. In the paper, we have a more fine-grained analysis that shows what people actually clicked. And here we found that the position that yielded the highest interaction rate was in the lower left corner of the screen.

We had some theories where this might be the case. So the theory for top versus bottom was that on top, usually the banner is more likely to cover some less important parts of the website, like some header or a menu, while on the bottom you usually have some content and text. And the same argument applies for left versus right, because if you have text written in the Latin alphabet, everything is skewed to the left, so more important information can be -- so there's more information on the left versus on the right.

Next slide, please.

In our second experiment, we looked at the different options offered by the website and the influence of nudging. And here we had five different banners in terms of option. So we have one banner that doesn't offer you any type of option at all. The one on the very left, you just have a little x to make

the notes go away. Then the next banner is one that just has an accept button. And here you can see the non-nudging variant of all the banners. So this exit button is not highlighted. The next banner in the middle is the binary banner you've already seen. And here this is the non-nudging where each button looks like the other.

And then we had some more fine-grained options, one that allows you to check or uncheck different categories and one that has the same for different third-party vendors. And here the non-nudging variants don't have pre-ticked checkboxes, while the nudging variant would have pre-ticked checkboxes.

And here we saw a big influence of nudging. Because in almost all cases, the nudging variants yielded higher interaction rates. And the binary banner had the highest interaction rates. But combined with the qualitative data from our survey, we also saw that the category-based banner was also popular with users.

Next slide, please.

We then took a brief look at what people actually clicked. And here's just a quick example. So these are the selections people made on the

window-based banner. And you can see if you do not pre-tick the checkboxes, then there were only about 1 percent of visitors that actively opted in for one of the vendors, while in the case where you have pre-ticked checkboxes, we had about 10 percent who agreed to data collection by these third parties.

As for the results of experiment three, I don't have them on slides here because we didn't see any significant influence of either the presence of a privacy policy link and technical versus nontechnical language.

Next slide.

So then we took a brief look at what people wrote as answers into our survey. And we asked them why they chose to click or not click the banner. And among the reasons for not clicking -- for clicking, excuse me, we saw that one prominent reason was the expectation that the website would not work otherwise.

And this misconception was also present in other questions, like -- next slide -- when we asked what users expected to happen if they clicked decline or accept. The top reason or the top statement, what would happen if they hit decline, was that the website cannot be accessed. And this was named more often

than just mere functionality limitations, which would be much more likely than the website not working at all.

So what did we learn from all of this? We saw that the interaction rate is mainly influenced by position of the banner on the website and not the effects of nudging and preselections. We saw that users appeared to favor a binary or category-based approach versus more fine-grained ones, like vendor-based approaches. And there are widespread misconceptions about how consent notices work.

So one of them was that the site cannot be accessed without consent. So one recommendation here would be to inform users about the functionality limitations they can expect when they do not allow the use of cookies. And then from the survey, we also saw that people have some privacy by default expectations. So they expect no data being collected before they actively make a decision. And this is really not the case in reality. So this would be an issue that could be addressed by regulators because currently there are no incentives for companies to actively protect their visitors' privacy.

Thanks.

MR. WOOD: Thank you. I'd like to thank all

the panelists for excellent papers and really interesting research that they've contributed to this panel.

Now [audio malfunction] the way you would through email. And, hopefully, the slide that's available now is showing you where you'd send them.

But before that, I have a couple of questions for the panelists myself. And I'll start with Christine. So, Christine, your research found that the position of dialogues, the set of choices offered and other nudges significantly may influence consumer consent choices. How pessimistic should these findings make us about the possibility of mandating that firms obtain informed consent in other online contexts? Are those mandates going to be hard to make workable?

MS. UTZ: I would say that this depends on what this mandate looks like. Because in the case of the EU and GDPR, there was a lot of confusion about what informed consent actually means, because initially there was a big lack of guidelines how consent should be collected and what's actually free and informed consent. And only recently the EU has put out some documents that give a little more insight in that. But there are still lots of questions to be

answered.

And I think, yeah, if you want to introduce new mandates for firms to collect online consent, then there really should be some guidelines, along with a mandate that would help companies and anyone else who's collecting personal data to comply with the new regulations.

MR. WOOD: Okay, interesting.

I guess I'll rotate through the panelists. And my next question is for Jeff. So the privacy paradox, roughly stated, is that people report stronger preferences for privacy in surveys than they demonstrate with their actual behavior.

How well do the valuations for privacy you recovered in your research, Jeff, match up to valuations from consumer choice-based studies?

DR. PRINCE: That's a great question. I think the closest to us prior to our work was Savage and Waldman did some measures for the value of privacy with regard to apps. But there are some distinct differences. So I mean, obviously those are single apps. They looked at one-time payments. And theirs were kind of in the $1 to $4 range.

So are ours a lot bigger? I mean, technically, yes, since ours are monthly payments.

But again, you know, we're looking at major platforms rather than a single app. So people might look at that decision differently.

More broadly, I think, you know, with the privacy paradox, it's tough to say. I think, you know, this is one of the reasons why I think more quantification is valuable, because a lot of times, as people know, we ask people, do they value their privacy, and the answer is yes, maybe even a lot. But then that might not line up with what quantifiable metrics would be in terms of how much they value their privacy.

And, in fact, even with our numbers, we had outlets interpret our numbers as being large and we had outlets interpret our numbers as being small. So in some ways, it's a lot about perspective. So it's hard to align our figures with what people and people's qualitative analyses have revealed. But I think, you know, our hope is to contribute more to the quantifiable version of people's value for privacy.

MR. WOOD: Right. Well, so it seems to me like your numbers are going to become a standard for future research in the privacy space. So I thought they were really interesting.

DR. PRINCE: Thank you.

MR. WOOD: My next question -- like I said, I'm rotating -- is for Garrett, Garrett Johnson. So you talked a little bit about this, but it seemed like a lot of what you measured was the short-run adjustment to GDPR enforcement. Should we evaluate the effect of GDPR based on that? Or is the long-run adjustment what we should be interested in, or are both useful?

DR. JOHNSON: There we go.

It's an important question. My answer is emphatically that the short-run provides the best available evidence. But I want to make the three big points here. The first is that the GDPR is really confusing to study in this industry, because regulators have been slow playing the industry. So even as of today, the EU has not fined any of these websites or technology vendors. But at least three EU countries have criticized the industry for practices that do not comply with the GDPR.

So again and again, the EU keeps delaying enforcement and giving the industry time to adjust its practices. So the big question is, how are you going to study a law that hasn't been enforced? So we try to argue that the best time to do so is when the firms are most afraid of enforcement and change their

behavior accordingly. And our evidence does suggest that beliefs play an important role, like the fact that countries that face stricter regulators seem to keep their vendor use lower than those that don't, and, also, some evidence that also we talk about in the paper.

And the last thing I'll say is that this is an industry that moves very quickly, that has -- it's a fast growing market. And we see in general that the number of vendors increases over time. So given this fact, I think the best evidence we have is this big trend break we see right around May 25, 2018, which is the month right after the GDPR deadline.

MR. WOOD: Okay. Thanks for the clarification.

So again, let me encourage questions from the audience. If you want to ask a question, email privacycon@ftc.gov, and we're very interested in your questions.

Let me ask a question of Guy in the meantime. So, Guy, how do you think your results will generalize to other domains on the internet, beyond online travel?

MR. ARIDOR: Thanks. Yeah, I think it's a good question.

So when thinking about it, I think there's two things. So the first is that I think there's two dimensions to think about from the consumer perspective. The first is thinking about in a particular context what is the instrumental value of privacy and how do consumers perceive firms using their data. So in the context of online travel, at least anecdotally, consumers are more likely to be privacy conscious than they may be in other settings, because they think that consumers are -- or that firms are making use of this data.

And the second is that a lot of our analysis sort of hinges on understanding the impact of consumer histories. And so, for instance, if you compare this to a setting, like social media or something, where you see consumers much more often, it might not be as applicable. But in settings that are similar on those two dimensions to ours, we would expect the results to generalize. We would expect the sort of externality results to generalize to any setting.

Finally, I think it's also important -- and this is something that we spent some time trying to grapple with trying to contextualize our results in the context of the broader advertising ecosystem. So in particular, like the -- you know, the advertising

partner -- the intermediary we partner with sort of views itself as a competitor to Google.

And we think that our study sort of helps understand how a niche advertising intermediary gets impacted in terms of profitability in data observed by a smaller advertising intermediary. And we suspect that advertising intermediaries in other domains would be similarly impacted. And I think Garrett's paper sort of points to this, and there's a few papers that are pointing to the need to sort of think about the broader competition effects of GDPR on these things.

And so while we find that our results aren't wholly negative on the firm side, it would be interesting to think about how that would compare to, say, Google, who might not have been as impacted by GDPR as our advertising intermediary. So we suspect that our results on the impact to a third-party intermediary would be similar. And it's interesting future work to think about how that would impact itself in a broader competition between advertising intermediaries.

MR. WOOD: Cool. So let me ask you another question, Guy.

So part of what I found fascinating about your paper was that there were these externalities

between different types of consumers. How is the move from -- can you dwell on a little bit more and tell us how the move from software-based data obfuscation to GDPR opt-out is likely to affect the welfare of consumers who don't have strong preferences about privacy?

MR. ARIDOR: Yeah. So that's a good question.

So I guess as an economist, I have to state the caveat that we do a reduced form metrics exercise, so we can't directly say anything about welfare. But we do argue indirectly in the paper that if you think of consumer welfare as largely depending upon the quality of services they receive, this is largely dependent on the ability of a firm to do prediction. And so particularly in our context, we find that there, there's a marginal improvement in prediction. And this may lead to other domains and better personalization and ultimately improve consumer welfare.

There's obviously settings such as where firms are using this prediction to do price discrimination where, you know, arguably it would reduce consumer welfare. So the way we try to frame it is the effect for consumers really depends on the

alignment of preferences between firms and consumers in terms of how they use their data.

But I think it would be -- and again, we point this out the paper -- it'd be interesting to do a proper structural analysis to really decompose the welfare benefits to these policies.

MR. WOOD: Cool. Let me turn back to Garrett.

So, Garrett, while the California Consumer Privacy Act is sometimes compared to GDPR, there are some differences. Do you think the CCPA -- or if you want you could imagine a different hypothetical federal privacy legislation. Do you think that sort of legislation would lead to similar increases in concentration that you find?

DR. PRINCE: I think it would have different effects on competition. So I think the main difference is that the GDPR has a data minimization principle, and that places pressure on firms to limit their data-sharing partners. To my knowledge, the CCPA lacks this principle, which seems to be doing most of the work in our setting.

Instead, the CCPA operates on a notice-and-choice basis. So it basically tells sites you need to put some opt-out button on their site that

allows consumers to avoid data sharing. And we know from research, like Christine's and Guy's and my own works, that opt-out rates, at least according to our stuff, is like 5 percent to 15 percent when sites have to display this prominently. But the CCPA insists on this colorful language which is, "Do not sell my personal information" for the opt-out button, which I would speculate -- it would be interesting to hear what Christine would say about this -- I would think this could increase the opt-out rates.

So while I don't think this is going to have an effect on vendors, I do worry this could have an effect on the publisher side. In particular, large websites may have an easier time gaining consent than smaller and less recognizable websites. And our work examining over a thousand firms using Adobe Analytics data and the GDPR is consistent with this finding. We see a larger reduction in recorded web outcomes for smaller websites in our data.

MR. WOOD: Okay.

Christine, did you want to speculate on the effects of the CCPA's -- what's the exact wording, Garrett?

DR. PRINCE: "Do not sell my personal information."

MR. WOOD: Is that a good notice?

MS. UTZ: Actually, we already did some investigation of that. So we took a look at a couple of -- I don't know how many -- a couple of thousand of US websites, and we looked at how they implement the CCPA link. And we saw really, really big variance in how this link is named. Often, it's "do not sell" and "do not sell my info," "do not send my personal info." There were dozens of variants on that.

Yeah, this makes you wonder how the sites will deal with the rest of the CCPA requirements if they're already having kind of difficulties complying with this simple requirement like just put in a link that has this wording.

MR. WOOD: Yeah. My understanding -- and I haven't been following it super closely -- was that I think the California Attorney General might be producing guidance. I don't know if they have a deadline. But the guidance is -- you know, if they read your paper, the guidance might be great.

So let me ask a sort of very broad question of you, Jeff. What's the most important points about privacy policy that you think privacy policymakers should be taking from your research?

DR. PRINCE: Oh, wow, that is a big one. I

guess, you know, one of the takeaways for us was there was a very noticeable similarity in both the relative and, even in a lot of ways, the absolute preferences for the different types of privacy. I think many would have predicted ahead of time that Germany would come out with the biggest numbers, and they did. Although even with Germany, if you take away the financials, they're not that much higher than everybody else.

And we have a pretty wide cross-section. So I mean, obviously, we're influenced by our funding source in terms of where we directed our focus. So we had a lot of Latin American countries. But I also think we didn't know a lot about what was -- you know, different privacy preferences across those countries.

So, you know, that was one of the big takeaways, at least for me, is that when you think about privacy policy -- how people value privacy in a relative sense -- across countries and across different types, there wasn't that stark of a difference across countries, even though obviously there's vast cultural differences and other differences across those countries.

MR. WOOD: Yeah, that was interesting.

So it seems like -- just throwing this out

-- it seems like that sort of structural similarity makes me a little more confident in the survey-based approach. It seems like a method producing something -- the real preferences -- if we're finding consistency across different countries and across different groups of people.

But let me turn to back to Christine. So, Christine, you found that some users have misconceptions about how either they were -- how either what interacting or not interacting with consent notices would tell the website how they should use their cookies. How can we improve -- what's the best way to help these users have more accurate expectations about the very simple act of interaction?

MS. UTZ: Yeah, that's really an interesting question. And we've been wondering that, too. So I think one big step in the right direction would be to actually make websites comply with the user selection. Because there are other papers that have shown that many websites already start tracking before you've actually made a decision.

And the next step would be to just tell them which parts of the website they can expect to work and which won't work if they decline certain types of data

processing. Like, for example, in the case if you have an embedded YouTube video, you can say, okay, if you don't agree to YouTube setting cookies, then you just will see a gray box or something and not the embedded video, something like that.

MR. WOOD: Okay. So we did get one -- actually, two audience questions. And one of them is about Privacy Shield. And that's a big area, so we're not going to touch that yet.

But the other one is a somewhat more specific question for Christine. So we only have about a minute left. But, Christine, if you feel like you can do justice to this, how can data protection authorities, who often review consent statements for consent but not other factors, incorporate your research into their day-to-day work?

MS. UTZ: Okay. Yeah, I think they could maybe feel encouraged to issue -- to come up with some guidelines. I mean, we already have some guidelines. One was recently published by the EU. And we have some other guidelines by different EU member states. But I'm sure there are still lots of uncertainties what constitutes valid consent.

And then one big issue we still see -- this is something Garrett has already pointed out, which is

we do have the laws, but there's just a big -- they're not being enforced right now, or just a very, very small extent. So right now, there's really no incentive in many areas to comply with GDPR, because there's just a lack of enforcement.

And one big problem in this area is that data protection authorities just lack the funding and the personnel to actually enforce the law. But I hope we'll see some changes in that in the future so that the regulations can finally exist, not just in paper, but also actually in live systems.

MR. WOOD: Okay. Sounds reasonable to me, speaking purely for myself and not for the Federal Trade Commission.

I would like to thank you all again for a wonderful panel and for participating in this year's PrivacyCon. PrivacyCon will resume in about eight minutes. But for now there's a virtual coffee break. And when we resume, we'll do the last panel of the day on miscellaneous topics in privacy and security. So thank you. Thank you again.

DR. JOHNSON: Thanks, Dan.

DR. PRINCE: Thank you.

MS. UTZ: Thanks.

SESSION 6: MISCELLANEOUS PRIVACY/SECURITY

MR. HINE: Hi, everybody. Welcome to our final panel of the day, panel 6. This is sort of the miscellaneous panel. We'll call it the potpourri panel for today on privacy and security issues.

Just a couple of reminders. One, that you can send questions to the privacycon@ftc.gov mail address. And if you're tweeting, please make sure to use the hashtag #PrivacyCon20.

So we have four great panelists today to round out the day. We have Hana Habib from Carnegie Mellon University; have Ido Sivan-Sevilla from Cornell Tech; have Daphne Yao from Virginia Tech; and we have seen Yixin Zou from University of Michigan's School of Information.

So without further ado, we'll turn the floor over to Hana Habib from Carnegie Mellon.

MS. HABIB: Good afternoon, everyone. Today, I'll be presenting two research papers on the usability of online privacy choices on behalf of my coauthors at Carnegie Mellon and the University of Michigan. These papers are published at the Symposium of Usable Privacy and Security 2019 and CHI 2020.

As I mentioned, our focus in this work are privacy choices on the internet. And the next slide

has examples of regulation which mandate these privacy choices. And this includes the GDPR in the European Union, as well as the CAN-SPAM Act, COPPA, and now the California Consumer Privacy Act in the US. Additionally, groups like the Digital Advertising Alliance also work towards self-regulation in the advertising industry.

On the next slide are examples of three types of privacy choices that are commonly mandated by regulation and self-regulatory guidelines. And these include opt-outs for email communications, opt-outs for targeted advertising, as well as data deletion choices.

Next, I'll go over our research questions, which explore how these mandated privacy choices are provided in practice. We asked what choices related to email communications, targeted advertising, and data deletion do websites offer? Additionally, how are websites presenting these privacy choices to their visitors, and what are the potential usability issues?

To answer these questions, we conducted two studies. The first was a manual, in-depth content analysis of privacy choices on 150 websites. We followed up on this work by conducting an in-lab

usability study of a subset of these choices.

So next, I'll go a bit into more detail about our study protocols. To standardize the data recording for empirical analysis, for each website, we filled out an analysis template with 82 questions. An example of these questions included the location of the privacy choice, was it in the privacy policy, account settings, somewhere else on the website; the level of detail provided about each choice; the availability of links to the choice; as well as the path of implementation, for example, how many user actions were required to actually use the choice.

So next, I'll provide a quick overview of how we selected websites for the study. We randomly sampled 150 websites from Alexa's Global Top 10,000 list as of March 2018. All 150 of these websites were analyzed between April and October 2018, and half of them were reviewed by two researchers with an agreement of .82.

Next, a bit more about our user study. For our user study, each study session we conducted had three portions. First, we conducted a pretest interview to understand what users already believed about data collection and privacy controls. Next, we had participants complete two study tasks.

First study task, we identified a set of nine websites that had common implementations of privacy choice mechanisms that we identified in our empirical analysis. We gave users scenarios to describe this privacy choice task and asked them to complete this task as they would in the real world. For some websites, the scenario would require going to the account settings, while on others, it required going to the privacy policy. And the policy mechanisms could appear as links in the policy text or it could be described within the text as instructions.

In the final component of the study, we asked participants interview questions after the test to capture information about their experience and understanding of the study tasks.

So next, I'll go over a few of our results. But I encourage you to read more in our papers. So first, some good news. From our empirical analysis, we found that privacy choices are common. Almost 90 percent of websites that use email marketing or targeted advertising in our sample offered their respective opt-outs. And almost three-fourths of all sites in our sample provided a data deletion mechanism.

Next slide, please.

In our empirical analysis, we found that privacy choices were often provided in privacy policies. The downside of that, other than consumers largely ignoring privacy policies, is that the headings under which choices are presented are inconsistent from policy to policy.

This table presents bigrams and trigrams in headings of sections that describe these privacy choices. We noticed that some terms were evenly distributed, like "your choice." However, there were more unique terms for certain types of choices, such as opt out for email communications, third party for targeted ads, and your right for data deletion. Alarmingly, no single n-gram occurred in more than 20 of our analyzed policies. And this lack of consistency across websites could make it hard to locate choices in privacy policies.

Next, I'll go over another reason why offering choices through privacy policies is less than ideal. From our user study, here's an example of a privacy policy that users encountered on one of the websites in the study during the scenario in which they were trying to stop seeing ads for shoes that they searched for last month. So here are some

relevant information about how ad partners use cookies and beacons to decide which ads to show.

On the next slide, we see that the first link here is for opting out of Google Analytics. And participants often clicked that first when trying to disable cookies. But this link isn't that useful if the main goal is to disable cookies, so it's not clear why it's shown first here.

On the next slide, we see that the information about disabling cookies was presented underneath that link.

Next slide, please.

So from our empirical analysis, we noted that another reason why figuring out what to do could be difficult is that websites sometimes provide multiple tools for the same type of privacy choice on different pages of the website. So take Twitter's targeted ads for example. First, in the account settings you can find the opt-out provided by Twitter itself. If you navigate to it's About Ads page, it only shows opt-outs provided by the DAA, NAI, as well as Google. If you go to its privacy policy, only the ones provided by Twitter and the DAA will show up.

All of these links to multiple opt-out tools spanned across multiple pages of the website may cause

confusion about what tools should be prioritized and what their differences are. In fact, this is something that we observed in the lab.

On the next slide, we have an example of what privacy policy participants saw in one of their tasks. Participants who saw this had a difficult time understanding which of these three links would allow them to opt out of targeted advertising. While it was confusing when there were multiple links leading to different tools, when there were multiple paths to the same choice for information related to a privacy choice, we observed that it actually tended to be easier to find.

On the next slide, we have an example from the lab. Most participants who were assigned a data deletion task on RuneScape.com found the information they needed through searching the website support pages rather than referring to their privacy policy. And this led them to the Your Personal Data Rights page shown on the next slide, where they were able to see that the website offered this.

Another major result, which I present on the next slide, is that using these choices require high numbers of user actions. The user actions could include clicks, hovers, scrolls, filling out form

fields, or other types of interaction. For example, we see here that on average a participant took about 38 actions to exercise a privacy choice using a policy link.

And this average includes the reality that most users make some mistakes, like going to the wrong page, clicking the wrong item, on the way to the final correct action. When we collected data about the shortest path to each choice in our empirical analysis by performing the same tasks with prior knowledge of the location of the choices, it still required a high number of user actions. In the case of policy links, even if someone already knew exactly how to get to the final step and took the shortest possible route to get there, it would still take about 22 actions. In the lab, we uncovered some practices that required unnecessary effort.

On the next slide, here is an example of a part of a complicated form that some websites require to exercise a privacy choice. The example shown here is a form for deleting data on the New York Times website. Most participants dislike the number of similar-seeming options here. Further down the form, you had to select from a list of 22 different New York Times services, and you could only submit one request

type at a time.

The next slide provides another reason why exercising privacy choices might require unnecessary effort. So websites sometimes require users to submit written requests to complete actions, such as data deletion, when a simple web form would have sufficed. There were also participants who ended up writing emails to customer service to ask for help because they couldn't find a simpler way to do their task through the website itself. And it sometimes wasn't easy for participants to articulate what they wanted to do. For example, one participant who was given the shoe ad scenario I described before wrote this email to ask for help.

So after hearing so many issues, you might wonder, how do we improve the usability of website privacy choices?

On the next slide, we show that one way regulation could help improve visibility is to have explicit requirements that dictate parameters like the location of controls and the way that controls are presented. And the findings from our user study suggest that the CAN-SPAM Act has likely been effective in making email unsubscribing more usable. It mandates the look and placement of email opt-out

links in commercial emails, and users thus expect to find the unsubscribe link in that location.

Additionally, the next slide shows another way that policy could play a role, which is by standardizing policy section headings so that choices are easier to find. Such practice has been adopted by the US financial industry as a model privacy form to help financial institutions comply with the GLBA. Though it may not be perfect, it's definitely a good start, and research has shown that the standardization effort of the GLBA contributed to less ambiguity in privacy policies.

As summarized on the next slide, another way that choices could be made more usable is through unified settings. This could simply mean matching user expectations by always having privacy choices easily accessible within websites' account settings, rather than buried elsewhere on the website or its privacy policy. There is also the possibility of further unifying choices for users by offering more universal mechanisms, such as through a web browser that's able to parse privacy policies or use machine-readable privacy policies to help users exercise preferences across multiple websites with less effort.

Some of our group's recent research in the context of the California Consumer Privacy Act has also explored unified visual standards to help users find privacy options on websites.

The next slide provides an example this. The Privacy Options button here on the bottom right that we've tested is designed to convey the idea of choice and to serve as a central location for all privacy-related choices on the website. Ideally, this would lead to a dashboard that could also interface with automated tools to allow users to control privacy settings across multiple sites.

On the next slide, I wanted to quickly recap our work. We conducted an empirical analysis and in-lab usability evaluation of email opt-out controls, targeted advertising controls, and data deletion mechanisms. We found that privacy choices are prevalent but suffer from several usability issues. And our findings suggest that the standardization of choices through regulation could improve usability.

For more information about our ongoing work, feel free to visit this URL.

I'd also like to give a shout-out to my colleagues in the Usable Privacy Policy for their contributions to the research, as well as our funders.

Thank you, and I'd like to pass it off to the next speaker.

MR. HINE: Great, thanks. And that will be Ido Sivan-Sevilla.

DR. SIVAN-SEVILLA: Right. Thank you, Jamie. Good afternoon, everyone.

I'll be presenting our research today about the extent that third-party trackers in websites persistently identify users across websites or, more specifically, across social contexts. And this research was conducted with the research group in Cornell Tech, including Wenyi Chu and Xiaoyu Liang, two Cornell Tech Master's students, and Professor Helen Nissenbaum, Professor of Information Science in Cornell Tech. And we gratefully acknowledge support from NSA and NSF for this research.

Next slide.

Okay, so a little bit of background. If you think about the web, the web is an array of different social contexts. We go to the web when we want to look for information about our medical problems or express our educational aspirations or consume news. And advertisers take advantage of the fact that the web is an array of all these different things to conduct cross-context inference about individuals.

The fact that the web is an array of social contexts coming together is really profitable for this industry.

Think about, for instance, how advertisers can cross information about users' medical problems, educational interest, and news consumption habits. They become in a better position to know when a consumer can be turned into a purchaser and make purchasing decisions.

One more click.

And the fact of the matter is that we are never alone in the web. Embedding third parties in websites became an inevitable and disturbing social norm. According to recent statistics, there are 9 trackers on average per website and overall 33 tracking requests per page. And these trackers have the potential to undermine the integrity of our context and the way we browse the web and violate our privacy according to our approach of privacy as contextual integrity.

Next slide, please.

So this approach was also used in a previous paper presented by Madelyn Sanfilippo in this conference. And we argue that privacy is the appropriate flow of information based on informational

norms in a given context. Privacy is not about control, whether information is public or private. It's about how we use information.

So think about some examples of privacy violations according to this theory. So think about employment decisions based on religious affiliations. Think about the display of advertisements based on sensitive health information. Think about clinical tagging based on voice assistant data. These are all examples in which information was taken out of its original context, based on, without -- against the privacy expectations of the data subject that their information is about.

Next slide, please.

So what we're trying to do here is to apply this context-sensitive approach to online tracking. And when you look at previous studies, you see that online tracking was studied in bulk, across thousands of website, without distinguishing their social contexts. And our approach here was to apply a context-sensitive analysis to online tracking, comparing tracking across different social contexts of the web.

So what does it mean? Let's try to visualize what we're doing here. So one more click.

So for instance, if you go to WebMD.com, one of the third parties you will see there is DoubleClick.net. DoubleClick uses a user ID cookie and assigns you an ID, in this example the string, one to nine.

One more click, please.

Then you go to NYTimes.com, and you see the same third-party tracker.

And one more click.

The tracker assigns you -- or it uses the same user ID for you when you go to NYTimes.com and can potentially link information about you from both of these browsing sessions. This is the exact cross-context inference we're talking about.

So in this research, we're trying to label third-party trackers as what we call "persistent identifiers." So among all the third-party trackers out there in popular websites, who are those who persistently identify users across different social contexts against our privacy expectations?

Next slide, please.

Okay. So a little bit about our methodology. So we used an instrumented Firefox browser based on the Open Development Project from Princeton University to investigate the top popular 15

websites in three different contexts, in health, education, and news contexts.

One more click.

And as you can see from our chosen websites, these all embed very different dynamics or interactions for the users. Popular news websites, we consume news, we express our interest in news articles. For the healthcare context, we might express or share some information about our medical problem that we expect this information to be kept private. And, finally, in the educational context, we express our educational aspirations and maybe hint on our future career goals.

And we conducted six different experiments according to the different possible browsing sequences between these three different social contexts to realize whom among these third-party trackers persistently identify users across these different contexts. And it's important to remember that what we argue that what matters here is not only the amount of tracking within a website but also how those trackers choose to persistently identify users across the social contexts. We expect our information from the healthcare context to be used for health advice rather than for commercial purposes in other

websites.

According to user surveys, and some of them were discussed in previous sessions in this conference, people are not comfortable with trackers navigating their data between different contexts to get a better understanding of their profiles. And this is what we're trying to measure in this study.

Next slide, please.

So about our data analysis approach, so for each experiment, we were matching ID cookie among the contexts to realize which trackers use the same user ID for every context. So first, we observed all the third-party trackers that interacted with our browser. Then we detected all associated cookies of each tracker and grouped our data based on cookie name and cookie value pairs. Then we selected identical cookie values that appear in more than one social context. And. finally, we applied a known methodology to recognize among all these cookies with an ID cookie based on the uniqueness of this cookie and its length in the browsing session.

We did not simply assume that the presence of a tracker in two different social contexts means that they persistently identified the users across those contexts. Instead, we looked for valid

evidence, in this case the usage of the same cookie ID across these contexts, to assume this persistent identification trend. And we acknowledge that our results represent only a lower bound of these persistent identification instances. We are aware that cookie values are often hashed or encrypted when used by the same tracker.

We also acknowledge that persistent identification of users is happening in the server side as well, in ways that are more challenging for detection. So we expect our results to be considered as the lower bound of the amount of persistent identification that's actually happening in the web.

Next slide, please.

Okay. So now let's see some overview of our findings. Ultimately, we found that social contexts matter for trackers. We found a third of the studied third-party trackers use persistent identifiers among all three social contexts. Secondly, we saw that the three contexts that are linked by third-party trackers are linked to a different degree based on the website that is under study. And I will show this in a moment. And, finally, and maybe most interestingly, we found that third-party trackers are more likely to persistently identify users following users' visits to

healthcare websites. And this is especially alarming in our times of the global pandemic, when healthcare websites are becoming extremely popular, when users seek health information.

Next slide, please.

Okay, this figure is trying to start and capture and present what we found. So we've overall found that user IDs that were generated while browsing in healthcare websites are more likely than others to follow users to other social contexts of the web, to news or to education contexts.

And here you can see that 68 of the third-party trackers, when we visited healthcare websites first, were labeled as persistent identifier, which means that the user ID that was generated for you when you visited a healthcare website is likely to follow you by more trackers than in other experiments when health websites were visited after different websites. So the user ID that was initiated for you when you visited the healthcare websites is highly appealing for third-party trackers in other social contexts of the web.

Next slide, please.

And this is a complementary figure that shows that for news websites -- so when you visit news

websites after health websites, you see 69 persistent identifiers, 69 third-party trackers that are turning to persistent identifiers linking the ID that was assigned to you from a healthcare website. So this is very appealing for third-party trackers that operate in other websites to know that you were visiting healthcare websites as well.

Next slide, please.

Then we decided to zoom in. And we wanted to graph how this is happening between the websites, how persistent identification works between all the different websites that were under investigation. So we created an edge between two websites in case an ID cookie was used by a tracker that is present in both contexts. In this example, we saw in WebMD and NYTimes.com, Twitter, which is in this case the third party, uses the same user ID in both of these contexts, potentially linking our browsing habits from both of these contexts.

One more click.

And if there was more than one third party, our edge became thicker. So we're trying to understand the scale of this persistent identification trend between websites.

Next slide.

Okay, so how this looks when you look at scale at all of the websites under investigation. And I know this might be a little a little tiny, so you can zoom in on the upper right corner of your screen to get a better look of what we're trying to visualize here.

And what you see here is our browsing session moves from the healthcare to the news to the education context. You see that third-party trackers in news websites link user ID from healthcare websites, potentially violating our expected privacy norms from these websites. So you see that for each and every healthcare website, there is a link, to a varying degree, of persistent identification trends to a news website. So this is very appealing for trackers in a news website to link our ID and study about our behavior from these healthcare websites.

Next slide, please.

And here you can see this again. We're moving from health to education to news websites, and you see the dominance of persistent identification from healthcare websites. Trackers from each healthcare website identify users in each of the education websites and in the news websites, but in different volumes. So the thickness of each edge

means that a different number of trackers are actually following this trend of persistent identification.

Next slide.

Okay. Three takeaways from this study. So first, we see that users who consume their news or visit educational resources after browsing at healthcare websites are potentially more vulnerable for manipulation by the advertising industry. Like I said, this is especially alarming in times of the global pandemic.

Secondly, what matters for users' privacy is not only the amount of tracking within a given context, but also the extent that trackers link information about users between those contexts for potentially better targeting purposes.

And, finally, like I said, healthcare websites, which were regarded in previous studies as less dangerous for users' privacy because they had less number of third-party trackers that were following you, are actually the most alarming ones when it comes to persistent identification trends.

And next slide, that will be my last one.

So to conclude, we argue that this is a first modest step to apply contextual understanding to online tracking. We argue that this is a rather

unaccounted privacy violation. We should all work for keeping the integrity of our different social contexts when we go online, no matter how profitable their conflation might be for certain parties, in this case third parties and advertisers.

And, ultimately, this work is a call to apply a more context-sensitive analysis to online tracking in order to better understand this rather unaccounted privacy violation. So more, of course, is in the paper. And I'm looking forward for your questions and comments.

Thank you.

MR. HINE: Daphne, you have the floor.

DR. YAO: All right. Thank you, everyone. Thank you, Jamie.

I'm Daphne Yao from Virginia Tech. Today, I'm going to talk about payment card security. And this is work published in ACM CCS 2019 last year. And it was in collaboration with my PhD student, Sazzadur Rahaman, who just defended his PhD thesis yesterday, and my colleague, Gang Wang, from University of Illinois.

Next slide.

PCI stands for Payment Card Industry. The body behind the data security standard, the DSS, are

big banks, so Visa, MasterCard. They formed what is called the Security Council -- PCI Security Council. The standard, it started in 2004, many years ago.

A little bit of history. Before there was DSS 1.0, Visa came up with data security standards on its own. And quickly, many other companies -- MasterCard, Discover, American Express -- followed suit. And then it was so confusing. The payment ecosystem had so many intertwining components. The acquirer banks, the issuer banks, the merchants, they have to work together on transactions. And so it's very confusing to have one standard for Visa and another standard for MasterCard.

And so the big banks have formed the Security Council and decided that let's just unify all the data security standards. The current version is 3.2.1, which has evolved tremendously since its first version. The 4.0 version will come up in 2021.

So I got very interested in PCI.

Next slide.

So I got very interested in PCI DSS because of the Target data breach. I wrote an article explaining the details of the Target data breach. It occurred in 2013. And some of you may know the initial entry point of the attacker was this air

conditioner system. So [indiscernible] Fazio Mechanics. One of the employees there fell victim to a phishing attack. And, eventually, that person's credential was used to access internal Target networks because of the lack of network segmentation. And, eventually, malware, what is called the BlackPOS, was installed on point of sale devices in Target. Forty million credit card numbers were compromised.

So as I was reading about the Target data breach, Target was actually in compliance with DSS, the Data Security Standards, back in 2013. And that was one of the main arguments that Target's CEO Gregg Steinhafel used to say, oh, we are in compliance; we got breached; it's not our fault. But as you look into the standards, you realize that a lot of those measures were just a sanity check. It was just a baseline. So I will explain a little bit more about the exact measurement we did.

Next slide.

So as you look into a bit closer about the DSS standards, you realize that regardless of which merchant size you are -- I mean, you can be Walmart, you can be a mom-and-pop shop, 7-Eleven -- you have to satisfy this, what is it called, a quarterly scanning report. It is an external scan of your network, the

payment network of the merchant, and to ensure that all the system that touches the credit card has to be compliant with a set of standards.

For bigger merchants that have more than 6 million transactions per year, they have to follow additional requirements. For example, an auditor has to be on site and go through some internal design configurations to ensure the correctness and security of the systems.

And so we decided to focus on this ASV, this Approved Scanning Vendors, how secure it is, and from a scientific point of view, can we quantitatively measure it?

Next slide.

And so this is not really anything that people have done before in a way that has some quantitative measurement of commercial scanners. And part of it is to understand the requirement of DSS specifications and then be able to reflect it in some sort of a testbed that allows you to test the scanners.

Next slide.

And so we spent a long time designing -- you know, how should we set this up? And so we eventually decided to put together what we called BuggyCart Testbed. It is an e-commerce website. It's a web

application that sells electronics. It has a card payment system. It has different options for the user to design their purchases. And so we used this as a testbed and embed altogether 35 vulnerabilities. But only 29 of them can be scanned externally. So we only need a scanner to find out 29 of them.

And so we find out there's numerous -- more than 100 -- scanners to choose from. And, of course, from a scientific research group, we have a limited budget. But we tried to cover high-end scanners and low-end ones. And, luckily, some of them offer free trials. And so we selected a few of them to test.

The way that we tested the scanning services is we just do a baseline scan and see how many vulnerabilities they can find. And then we'll follow their instructions to fix some of them, but then only the minimum amount of fixes. And so, eventually, we'll have a version that all of the testbeds indicating the minimum fix and a testbed that can pass the certification.

So next slide.

A quick summary of our findings. I'm going to explain a bit more. Five out of six scanners knowingly certified vulnerable merchant websites. And this is somewhat expected, also somewhat

disappointing. And I'll explain more why it's somewhat expected. In addition, we also put up our own scanner, a lightweight one. We scanned a whole bunch of websites -- a majority of them are not fully PCI-compliant.

Next slide.

A quick summary of the findings on the scanners. We eventually settled six scanners. Two of them are advertised as two different products, but they use the same engine. Two other scanners, 3 and 6, are not approved. They are not approved ASVs. If you look at this, the last column is the most important one. Only one scanner, scanner 2, does not allow vulnerabilities knowingly to exist in a certified version, even though there are seven vulnerabilities it cannot detect out of 29. So this is a very disturbing result. The must-fixes has to be a vulnerability score greater than 4.0. And it was defined in the ASV scanning guideline to have to be automatic failure. You have to -- the scanner has to fail the website. But most of them don't.

Next slide.

More information about the certain type of vulnerabilities called application security -- and those are the typical cross-site scripting, cross-site

request forgery, SQL injection, the harder one, the harder vulnerabilities -- failed miserably. On the right last four columns, those are research products. They are top-of-the-line web scanners. Some are research products, some are commercial products. They also don't do very well. So this gives you a serious pause, what's going on here.

Next slide.

Good news is that when we use our scanner scanning websites, a majority of them, even though they are not fully PCI-compliant, some of the typical issues don't exist, you know, default MySQL username/password, weak hash in certificates, browseable directory. Those are dot, dot, slash, and you can go back. And those are gone, which is good. Animation, one click, please. However, we've seen wrong domain names, vulnerable OpenSSH versions, expired certificates. And those are the issues that PCI compliance prevents but that still exist.

Next slide.

And, of course, that's not very surprising. If you have inadequate scanners certifying insecure websites, you inevitably will have vulnerable websites.

And so this is a first quantitative study

measuring PCI scanner capability. But then the issue is much, much more beyond the PCI by itself. We tested web scanners; they don't do well. We tested research product; they don't do well either on certain types of more complicated web application vulnerability.

So what does it mean? And so if you can remember one thing, that's this slide. For all various stakeholders, everyone needs to improve. This is definitely not some work to say scanners, you know, you should be blamed. No, no, no. Everyone needs to improve.

The research community needs to have more deployable solutions. For cross-site scripting, a concept that's been around for a long time, there's no good open source deployable grade solutions. Regulatory authorities, how can we improve the specifications? And then part of it is to have a holistic measurement of system security as opposed to just one check, one check, one check, put them all together. Scanner evaluators, how to improve, more importantly, more robust testbed.

Next slide. Last slide here.

PCI specification, very comprehensive. We were very impressed about the completeness, but

enforcement is tough. Research needs to catch up. We also disclosed our findings with the Security Standard Council and got positive feedback. And this is a problem that needs everybody in the community to improve.

That's it for my talk. Thank you.

MR. HINE: Excellent. Thank you so much, Daphne.

Yixin, final presentation.

MS. ZOU: Thank you, Jamie. Hi, everyone.

My name is Yixin Zou, and I'm happy to present this research collaboration between University of Michigan and NortonLifeLock Research Lab on the adoption and abandonment of security, privacy, and identity theft protection practices. This paper was published at CHI 2020 with a best paper honorable mention and was sponsored by NortonLifeLock Research Fellowship.

Next slide, please.

Consumers need to know how to protect themselves online. Data breaches, hacking, phishing are but some of the threats they face. While there's lots of useful expert advice on how to protect oneself for privacy and security online, study after study shows that most consumers do not adopt best online

security practices, potentially leaving them at risk.

Next slide, please.

What we don't know, however, is whether this low adoption pattern also persists to other online safety practices, such as those for privacy and identity theft protection. Moreover, there's limited knowledge about what happens after consumers adopt advice, such as how often they abandon this advice and why.

Next slide, please.

For our research questions, we scoped them with security, privacy, and identity theft as three key dimensions for online safety. First, which online privacy and security practices are fully adopted, partially adopted, or abandoned?; second, what factors predict the level of adoption?; and third, why are certain practices partially adopted or abandoned?

Next slide, please.

We selected 30 expert recommended practices in all three domains from prior work. We included 12 security practices from Ion et al.'s 2015 study. We included 12 items from the US Census Representative Survey by the Pew Research Center for Privacy Practices. And we included six items from FTC's online resources for identity theft practices.

Next slide, please.

Here I want to give an overview of practices we examined. Security practices include two-factor authentication, antivirus, cautious clicking behavior, good password habits, and so forth. Privacy practices include a management of one's browser extensions and cookies, careful online disclosure, use of VPN and encryption, among others. Identity theft practices are a mix of services provided by credit bureaus, commercial services, and manual tracking of credit reports and statements.

Next slide, please.

As an overview of our method, we chose to conduct a survey and recruited 902 participants for the study on Prolific. All participants were U.S. residents, since some of our examined practices are specific to the U.S. context.

Next slide, please.

The main survey questions were about 10 practices randomly selected from our list of 30 to not overwhelm participants. With 900 participants, this resulted in about 300 data points per practice. Participants could select if the practice was something they always did, did it with exceptions, did in the past but abandoned, consider doing, rejected,

or were not aware of.

Next slide, please.

At the end of the survey, we collected information about demographics, technical background, and prior negative experience. All participants' gender and income distributions are representative of the US population, but are skewed to younger and more educated people.

Next slide, please.

On to our findings. First, what practices were adopted or abandoned the most?

Next slide, please.

We found high adoption of security practices indicated by the deep blue in this graph. Interestingly, the top two adopted practices had to do with cautious clicking, 95 percent for click links in emails and 93 percent for attachments in emails.

Next slide, please.

We found that the most abandoned practices tend to be privacy-related, though practices were not often abandoned, and the abandonment rate was below 20 percent for all practices. Looking at the light blue bars in this graph, practices with the highest abandonment rates were using anonymous systems like VPN, use fake identities for online activities,

and clean web browser cookies periodically.

Next slide, please.

By contrast, the adoption of identity theft protection practices were concerningly low. Looking at this large area of red and orange, most participants were either unaware of or rejected these practices. Credit freezes and fraud alerts, though strongly advocated by the FTC over the years, were among the top rejected practices, with more than 50 percent rejection rate.

Next slide, please.

On to answering our second question, what were the factors that influence a practice being fully adopted, partially adopted, or not adopted?

Next slide, please.

For factors related to the practice, we confirmed that adoption levels of security practices were significantly higher than the other two domains. Additionally, we divided practices into three Subcategories: manual practices that rely solely on user efforts, such as avoid clicking suspicious links; automated practices that, after initiated, require no user effort, such as running antivirus software; and, finally, assisted practices that use tools but still require regular user interactions, such as two-factor

authentication. What we found is that assisted practices were adopted significantly less than manual or automated practices.

Next slide, please.

We also examined factors related to the user that influenced adoption. Experts had higher levels of adoption than nonexperts. And we further unpacked this difference and found that computer science and IT expertise, more so than privacy and security expertise, significantly impacted adoption rates. Additionally, being a previous victim of identity theft made someone more likely to adopt protection practices across all three domains. Our paper includes more details about other findings related to demographics.

Next slide, please.

We analyzed participants' open-ended responses to understand why they partially adopted or abandoned certain practices.

Next slide, please.

Regarding reasons for partial adoption, the most common one is only adopting practices for certain sensitive sites, which is the case for private browsing. Another 10 percent of participants with selective partial adoption said the practice was

inconvenient and difficult to use consistently. For instance, saying if "I'm in the middle of doing something, I won't be able to install this software update," or saying that "it's hard to keep track of unique passwords for different accounts."

Next slide, please.

Regarding reasons for abandonment, 20 percent of participants who used but then abandoned a practice say they don't need the practice anymore, as it does not provide sufficient values to guarantee continuous usage. Like, "I have used it, but don't find it all that helpful for private browsing." Another 14 percent reported abandoning a practice when the perceived risk has diminished after a negative event. For example, "I had a credit freeze due to suspected identity theft in 2012." But after some years, they decided to not use the freeze anymore.

Next slide, please.

We discuss how our research has implications for how experts can provide online safety advice to consumers to increase adoption and reduce abandonment.

Next slide, please.

To bridge the gap that security practices were adopted much more than privacy and identity theft practices, it's important to show the synergy that

exists between practices, especially in cases when multiple practices could add additional protection layers. For instance, to combat phishing scams, avoiding clicking on the links is a common security tip. But this advice could be complemented by recommending users also actively monitor their financial accounts as an identity theft protection tip, but also an important mitigation practice after one has fallen for phishing.

Next slide, please.

Using an FTC's online article for identity theft self-protection as an example, this could be improved by giving more guidance as to which practices are most important to adopt and the connections between different practices and how they mutually benefit each other. For example, with measures for keeping your personal information secure online versus offline, we can illustrate how they work together and why it's important to do both. Moreover, it's important to identify the most effective and urgent actions to be prioritized so that consumers are not overburdened to take all actions at once.

Next slide, please.

We can also leverage at-risk situations for communicating advice given the finding that

experiencing identity theft drives the adoption of online safety practices. In case of a data breach, consumer-facing data breach notices can be a possible venue for education. Consumers reading these notices will be highly motivated to resolve the situation and mitigate future risks. And so resources that encourage and explain how to adopt protection practices will be most effective at that moment, though the advice must be actionable.

Next slide, please.

We discussed how current tools for consumer online safety protection can be improved.

Next slide, please.

We found that usability issues prevented the full adoption of practices across all three domains. This echoes previous research in computer security about 2FA, password managers, software updates, and encryption, et cetera. Though these are security practices, in our study, we also found evidence of usability issues with privacy and identity theft protection practices as well.

Next slide, please.

This calls for more systematic research to better understand what these usability issues are and how to solve them. And another potential idea, more

relevant to lawmakers, is to require usability testing for provided tools so that they are not made hard to use intentionally, which can reduce the burden on consumers.

Next slide, please.

As examples for requiring usability testing, we can think about requiring readability testing in data breach notification laws to ensure that breach notifications are readable and reduce the chances of them being lengthy and full of jargon. We can also think about auditing dark patterns in mandated privacy notices and controls to give consumers real autonomy in privacy and data choices.

Next slide, please.

To summarize, for our study we studied the adoption and abandonment of various online safety practices. We find different patterns of adoption and abandonment between security, privacy, and identity theft protection practices. This implies the importance of expert advice to emphasize that synergy exists between practices and, two, that more work is needed to improve the usability of privacy and identity theft tools in order to reduce user friction and encourage long-term adoption.

Feel free to refer to our paper for more

details and reach out to me if you have any questions.

Thank you.

MR. HINE: Excellent. Thanks so much, everyone. We really appreciated those presentations.

Let's move into Q&A. Just a reminder, if you have any questions, feel free to send them through the privacycon@ftc.gov address, and we'll try and reach out and get to some of those.

So the first question I have actually is for Hana. I wanted to first ask you, you know, one of the conclusions that you reach in your paper is about notice and consent, which you rightfully mention is sort of a dominant approach here in the United States. But you suggest that consent mechanisms have failed to provide consumers meaningful privacy protections. And so my question for you is whether your analysis justifies some type of an alternative approach. And the hard part of the question is, if there is one, what do you think that should be?

MS. HABIB: Yeah, I think going forward there still is a place for notice and consent. But it really needs to look a lot different from what it currently looks like now, which is typically you go to a website, you see a wall of text, and then you click a box that says, I agree, which doesn't necessarily

translate to meaningful notice or meaningful consent, because people don't really know what they're agreeing to.

We can potentially replace that with interfaces that allow people to make their preferences known up-front. Like I mentioned in my presentation, there is a potential for having tools built into the web browser, for example, where you set your preferences there, and those preferences are automatically communicated to websites without the user having to do anything, other than that initial step of setting those preferences to begin with.

And I think we should also consider what people should be consenting to. Is it specific uses of information? Is it what inferences can be made based on the data that's collected? So I think that's a space that needs to be explored in more detail. So in general, I don't think my work advocates for replacing notice and consent entirely, just maybe rethinking what that should look like in the future.

MR. HINE: So, Ido, if I can actually ask you the same question. I think that your research also suggests that consent mechanisms have failed to provide meaningful privacy protections. And I'm wondering if you agree with some of Hana's conclusions

or you think differently about that.

DR. SIVAN-SEVILLA: I totally agree with Hana's approach. I think consent became a meaningless term in our digital society. Users do not really understand what they agree for. They don't have a real alternative to choose from to get the service. Recent studies from Helen Nissenbaum and Kirsten Martin about what users actually think about information flows reveal that when users get aware of what's happening, they would never consent to what's going on behind the scenes of our favorite websites and mobile apps. And I think user awareness is critical to pivot around and change what's happening in this industry.

And one way to increase awareness is to visualize what's happening. There is an add-on to Firefox from Ghostery, a commercial company, to actually visualize what's happening, how many third parties are approaching you dynamically. And it's starting to get a sense of what's actually happening when you go to your favorite websites. So this is one step forward.

Users need to be much more aware of what's happening. And a complementary part of that is to require more transparency from these companies. How

do you actually use my data? How do you cross information about me? You can identify me in different contexts of the web, but what do you do with this information? That's what we call the server side analysis of things, which is going to be kind of a black box to understand how these companies are actually using our data.

This is our data. Remember, we are the data subject, and we have no idea what's happening with this data. So user awareness, one; more transparency on behalf of the industry, second. These are the first two steps to get us out of this disturbing path.

MR. HINE: So I just want to follow up. You mentioned a browser extension for Firefox. And Firefox is used by a relatively small portion of consumers in the country. So apart from sort of increasing the transparency about how the information is used, how do you think that we increase adoption or either creation of tools or tools that consumers can use to actually exercise their choice once they understand how their data is being used and they conclude they want to exercise choice to control or limit that usage?

DR. SIVAN-SEVILLA: Yeah, that's a great

question. First, we need to obligate service providers to provide alternatives for consumers. Firefox has done a very interesting step by preventing third-party cookies altogether. This is a great step for our privacy. But you see that the industry is now calling this the post-cookie area and moving to other ways to identify us and create fingerprints for our browser habits and operating system characteristics to know that we are the same person as we go over the web.

So the industry will always find sneaky ways to circumvent and go around. It's kind of a cat-and-mouse race for our privacy. So we need to make them more transparent about what they actually do to us. And then once we have this in place, it's for us consumers to decide what we actually want to do and weigh our options. But we have to have alternatives for the first place. And, unfortunately, we have no transparency and no alternatives. And the situation is not so encouraging.

MR. HINE: Great. I want to open it up.

Yixin or Daphne, do either you have any response to that?

DR. YAO: I agree.

MR. HINE: Okay. Sounds good.

So, Daphne, I want to pose the next question to you. One of the things that struck me so much in your findings was that scanners, at least in the PCI context, it appears, need some significant improvement. And it sounds more generally like some, either off-the-shelf or even open-source scanners, sometimes outperform ones that cost thousands of dollars, or for some things, like cross-site scripting or SQL injections for example, you may not really be able to find a reliable scanner to identify those vulnerabilities.

And I think about that in the context of the FTC and the work we do in the privacy division, where a number of the companies that are under order are required as part of those orders to engage in scanning and use tools to identify vulnerabilities as part of their assessments. So I guess the question to you is, are your findings more broadly applicable or are you just finding this to be a problem within sort of the PCI world?

DR. YAO: Great question, Jamie. It's definitely more broadly applicable. Some of the products that we tested are packaged as web scanners. So they have no mention of PCI, but then they still fail in some of -- a lot of the application level

tests.

And part of the struggle that we find is it just is so complicated. Because if you think about it, I would not -- now I will -- but a typical researcher would not say, okay, I have tenure to complete, I have a PhD thesis to devise, let's choose to build a deployable-grade cross-site scripting detector. No one in their right mind will do it, because the minute you submit the paper, you will immediately get rejected. The reviewer in most conferences will say, oh, this is not novel. We know about this attack. We know there is some way of, you know, conceptually how to detect it. Why am I reading this?

And the community, the research community, needs to change. It is changing slowly. I'm managing some conferences that try to push in this direction -- deployable and impactful security. You know, you consider it novelty, but then you also need to close the gap. There is this big gap between security theory and the practice. And then you need to reward researchers. Somehow, you want to encourage people's spending efforts, sacrifice of their time, and at the risk of not getting tenure to meet this, reduce the gap and meet the needs. And this is just a huge, huge demand.

And so a lot of the -- I think it's widely applicable. It's not just the web scanner. Many, many other aspects of security also need those kind of tools, the open-source tools that will be able to push the standards of the industry up. If you think about the profitability, you know, for-profit companies, the minute they put up a product, they will not list all the limitations, being against their interest. And so they will vaguely say, oh, you know, we cover this, we cover that. And then so it's only -- if researchers don't do this, don't do the measurement, don't provide transparency, no one will.

And then your security is something that there is no silver bullet, everyone knows, and there's no guarantee. And it's all, you know, the devil is in the details. And so you have to know what cases to cover, so what is the gap, what is missing, what is my attack surface. So that needs a lot of work.

MR. HINE: So I want to move over to Yixin quickly. But I want a quick follow-up, Daphne. I'm curious if you could very briefly talk about what the reaction from PCI was. Because if you've identified scanners and you believe that there may not be commercially available scanners to find certain types of vulnerabilities, how does an organization that

requires that type of compliance reconcile the fact that there may not be tools out there that reliably identify those vulnerabilities?

DR. YAO: Yeah, great question. So the person that I had a long conversation with from the Security Council fully acknowledged our findings, and I do understand their struggle. So basically, they have two testbeds, they test the scanners. But then, because the industry practice is a lot -- the level that we understand how to solve those problems is together collectively low, but then they have to certify some scanners, and so they have to reduce their bar to a certain extent.

And then PCI, they have built a very strong community. I really was very impressed that they help scanners to pass their tests. And so in that kind of thing, you know, they do have scanners. They said they kick out a lot of scanners out of their approval list. But then it's a problem that they have, they also struggle with, that if everyone fails the test, then this test is not very meaningful.

MR. HINE: Excellent. Thanks so much, Daphne.

Yixin, I wanted to talk about -- your research touches on usability. And it includes

recommendations, for example, practices and tools, to improve security, privacy, and identity theft. And so my question to you is, what do you think is driving this disconnection between users and interfaces? Is it just poor interface design? Is it just developer laziness or could it be consumers? Are consumers just simply unwilling to take responsibility for their privacy and protection on the web?

MS. ZOU: Thank you, Jamie. That's a great question.

So I guess my immediate response would be, I don't believe it's the incompetence of developers, designers, and engineers. I think we have people capable of doing this. The issues I see are probably threefold. First is the lack of understanding for usability issues. Like, in my work, I see this is well covered for security practices and for some of the privacy practices, but I have yet to see like comprehensive audits of major identity theft protection tools, even though our study has shown anecdotal examples from certain survey respondents. But we need a better understanding of what the issues are in order to solve them.

And then second is not to blame users, but we need to realize the fact that most consumers don't

have comprehensive understanding of technology, have limited knowledge, literacy, and also time. So for consumers, we need better education, more targeted, effective education to make them realize these are the available tools and how to use them, by giving very actionable guidance.

And then the third part, I think more for regulators, is to think about how to motivate companies to design usable tools. And things like what I mentioned in my presentation, of the audits or patterns throughout mandated privacy notices and controls that I think regulators are already working on this, this will be a very meaningful step to ensure companies are incentivized to solve their usability issues, not intentionally making them hard to use because that's for their own profits.

MR. HINE: Hana, would you also like to comment on this? I think some of your work touched on some of these issues.

MS. HABIB: Sure.

So yeah, as Yixin mentioned before in her presentation, I think the need for user testing is there. Like you can't really produce these tools and expect them to work great off the bat. Developers aren't their users, so unless they have the actual

tool in front and interfaces in front of real people who are using these tools as they would in their normal lives, they really have little to go on in terms of what problems people might encounter and what might be difficult for people to understand.

And, additionally -- and I wanted to make another point -- I don't think it's that people are incompetent in terms of -- or that they don't really care about their privacy and security. In fact, I think the opposite is true. And that's what the research overwhelmingly shows. I think it's more that security and privacy typically aren't people's primary tasks. They're usually -- people are using websites, using applications to do something else really, not really to come up with a strong password, for example, anything like that. So typically, privacy or security might be in the way of them doing their primary task.

So rather than putting the burden on users to make sure their privacy and security is taken care of, it should really be on the part of companies to have better privacy and security practices. And I think that's where regulation can have a major role.

MR. HINE: Okay. So there was one question from the audience. And that is to Hana. And we'll finish up with that.

And the question is, did you notice any patterns in terms of the type of site and how difficult or easy it was to find opt-out information? So for example, were e-commerce sites more challenging to navigate versus gaming sites or popular news sites?

MS. HABIB: Yeah. I didn't get into this in the presentation, but we provided a little bit of details about this in our paper. So the way we sampled the websites was that we picked the really popular websites from the top 10,000 list as well as some less popular sites and sites that really probably most people haven't heard of. So we call them top, middle, and bottom sites in our paper.

And one positive note that we noticed is that, across the three different categories there really wasn't a difference in terms of the number of privacy choices being offered. But how and where they were offered seem to vary. So for top sites, for example, they generally had controls within the account settings as well as somewhere else in the website, like a privacy policy, or even like an About Ads page, dedicated About Ads page for websites that had targeted advertising, whereas the middle and bottom websites relied more heavily on the privacy

policies to provide consumers these choices.

Additionally -- oh, I guess that's the time.

MR. HINE: Oh, no. Please finish your thought, sorry.

MS. HABIB: Okay. Yeah, so additionally, the way that these choices were provided in the case of targeted advertising opt-outs, for example, the more popular top websites tended to have their own implementations of these tools and a setting within that, a setting for that, whereas other types of websites relied more heavily on third-party opt-outs offered through like the Digital Advertising Alliance or the NAI.

MR. HINE: Excellent.

Well, on that note, I just want to thank all the panelists so much. This has just been an absolute pleasure to moderate. The research is fantastic. I invite everybody on the web, please go to the event page and check it out. Everyone from this panel is going to check out, and I have a few closing remarks in just a moment. Thanks again.

MS. HABIB: Bye, everyone.

MR. HINE: Thank you, all.

CLOSING REMARKS

MR. HINE: Well, that brings us to the end of our fifth PrivacyCon. We are so thankful for having everyone here today. As Elisa mentioned earlier today and was obvious to everybody who was here, we moved things virtual and I think it went really well. There were a lot of changes that happened moving from a live event to a virtual event. And we just want to thank everybody for your indulgence today. We want to thank all of the panelists. Please know that they worked really, really hard to deal with technology and to make this such a great event.

Just a few thank yous again. I know that Elisa thanked a lot of people today, but I just want to thank all of the moderators. I want to especially thank Elisa, who worked so hard with me to help put this together. I want to thank all of the support that we had.

There are people like Leah Singleton, who has gone out of her way to help make sure that all the slides were as perfect as they could be. Alex Iglesias helped with every technology issue that we couldn't figure out. Cheryl Thomas was on top of all of the Twitter today to help make sure that we sort of

got the word out on social media.

There are so many other folks behind the scenes at the FTC that had to work twice as hard to make this happen virtually. We want to thank all of you.

A few last remarks. We will have the video up on the event webpage on the FTC.gov in a couple of days. You'll be able to go revisit and see all the presentations again. On the agenda, almost all the papers are linked, so you can access all of those papers. If there are updated versions in the next few weeks or months, we'll update them accordingly.

And I just want to remind everybody that PrivacyCon is an event that we started several years ago to help create relationships with people that are doing amazing research. And that's not just about PrivacyCon. We have the privacycon@ftc.gov address. We also have research@ftc.gov. We encourage anybody who's doing interesting work that you want to share with the FTC, please send it to us. Send it to us any time of the year, we're happy to take a look at it.

In the next few months, we'll be announcing our sixth PrivacyCon, which will probably happen at some point next year, same time, maybe online, maybe virtual again. We'll see in the next few months. But

we want to thank everybody for participating. And we hope to see you next year. Thanks again. Take care.

(The workshop was concluded.)