>> Manas Mohapatra: We're gonna get started for the second part of the day. Before we get into the substance, let me just reiterate some housekeeping notes that we already heard earlier today. As most of you have discovered, anyone that goes outside the building without an FTC badge, when you come back in, you're gonna have to go through the entire security process again, so just take that into account.

>> Male Speaker: [Inaudible]

>> Manas Mohapatra: [Laughs] Sorry about that. In the event of a fire or evacuation of the building, please leave the building in an orderly fashion. Once outside the building, you need to orient yourself to New Jersey Avenue, which is across the street this way. Across from the FTC is the Georgetown Law School. Look to the right-front sidewalk. That's our rallying point. And everyone will evacuate by fours. You just need to check in with the person accounting for everyone in the conference center if you can. But hopefully, we won't have to deal with that. If you spot any suspicious activity, please alert security or one of the FTC staff. And with that, we can get into the substance. And it is with great pleasure that I'd like to introduce our commissioner, Julie Brill, to start off the day. Commissioner Brill was sworn in as a commissioner of the FTC in April of 2010 to a term that expires in September 2016. Since joining the commission, Commissioner Brill has worked actively on issues most affecting today's consumers. That includes protecting consumers' privacy, encouraging appropriate advertising substantiation, guarding consumers from financial fraud, and maintaining competition in industries involving high-tech and healthcare. Before she became FTC commissioner, Commissioner Brill was the senior deputy attorney general and chief of consumer protection and antitrust for the North Carolina Department of Justice. She has been a lecturer in law at Columbia University School of Law and also served as an assistant attorney general for consumer protection and antitrust for the state of Vermont for over 20 years. And so with that, Commissioner Brill. [Applause]

>> Julie Brill: Good afternoon, everybody, and welcome back. We're so glad that you made it back from lunch. Good to see you. This morning, we heard about the emerging uses of facial-

detection technology -- uses that until recently seemed the stuff of a distant future or a galaxy far, far away. But here and now, advertisers are using facial-detection technology to identify the age and gender of a face exposed to their ad space and targeting their marketing to the demographics identified. A 20-year-old woman might be shown an ad for perfume, a 30-year-old man an ad for shaving cream. We also heard about a mobile application that checks out who's at a bar so users can scope out the scene before they even arrive -- a new twist on see and be seen. Back in my day, you had to do a lap around the bar before committing to the optimal barstool. [Laughter] And now you can do it from your home. These advertisements and apps rely on facial detection, not facial recognition. While they gather general traits, often in aggregate, they don't identify specific individuals. But as the chairman's remarked this morning, if not now, then soon we will be able to put a name to a face. For me, this subject brings to mind one of my favorite songs. It's a song of The Beatles. And for those of you who are under 40, just to let you know, The Beatles were a rock group. [Laughter] They had a hit or two in the 1960s. The song is "I've Just Seen a Face." It's a classic about love at first sight. It all happens at the first glimpse. Paul McCartney tells us nothing about the object of his affections, not even her name, probably because he doesn't know it. I can't help but wonder if this song might have turned out differently if facial-recognition technology had been around in 1965. What if, when McCartney saw this face, he had instant access to more concrete information about its owner, like her name, her date of birth, or place of work, her e-mail address, online purchasing history, or other personal information? Now, let's assume that, like me, she wouldn't have minded if a young Paul McCartney had invaded her privacy just a little bit. [Laughter] But what if she did? Or what if, instead of one of the Fab Four accessing her information, it was an acquaintance that she wanted to avoid or an insurer deciding she was too high of a risk to cover based on a photo of her skydiving or a potential employer judging her character based on a picture someone put online from that party freshman year? Or what if the technology confused her for another person? We saw this morning how technology doesn't even get the gender of the person that it's looking at right. [Laughter] And, Brian, if you're back from lunch, so sorry. [Laughter] It's scenarios like this that we must bear in mind as we both guide and react to how these technologies change the way we buy, sell, and live. This afternoon, we'll talk about what we can already use facial-recognition technology to do, what we can expect in the future, and what this means for the policies that we put in place, both today and in the near future, to protect American consumers and competition alike. We'll hear from representatives from companies already using facial-recognition technology, like Facebook, Google, and face.com. We'll hear how Facebook and Google are using this technology to make it easier to tag photos of friends and families. We'll hear how face.com, which is a start-up in Tel Aviv, how it's using app development -- I'm sorry, it's giving app developers the capability to use facial-recognition technology in their own apps. And we're gonna hear from Chris Conley from the ACLU of Northern California, who will share his perspective of facial-recognition technology from the viewpoint of an advocate for consumer privacy. Along with learning about the commercial uses of facial-recognition technology, we'll also hear about a groundbreaking study to determine exactly how far the technology has already progressed. Can it identify previously unidentified photos? Even a bit more surreal, can it match students walking on a college campus to their social-networking profiles? We're pleased to be joined today by Alessandro Acquisti, a professor of information technology and public policy. He and his team at Carnegie Mellon tested whether an individual's partial Social Security number could be guessed accurately using only facial recognition and publicly available information. The results of this study were surprising, and we look forward to being surprised again by Professor Acquisti today. To conclude our workshop, the last panel will discuss the policy implications of the increasing commercial use of facial detection and recognition technologies and address two issues. First, what protections for consumers are currently in place? And, second, what protections should be in place? I'm gonna be particularly interested in the part of the discussion of that last panel about mobile applications. If a user takes a photo and tags friends in it with an app using facial-recognition technology, will the friends who are tagged be notified? Will they have to consent to the use of their photo or the use of facial recognition? And if so, how will we at the FTC and how will other regulators enforce these privacy provisions? Now, we're honored that our fellow privacy regulators from Canada and the United Kingdom have joined us today, both as panelists and attendees. First, I'd like to thank my friend and colleague, Jennifer Stoddart, Canada's privacy commissioner, for being here. Dan Caron from her office will also be on a panel later this afternoon, and we're delighted that Dan is back at the FTC. He had spent several months with us back in 2009 as part of the FTC's international-fellowship program. Also from Canada, we're pleased that Fred Carter from the Ontario privacy commissioner's office is here with us today. Fred, I think, was one of the panelists earlier this morning. And from the United Kingdom, Simon Rice from the information commissioner's office is joining us, and we'll hear from him this afternoon. We're delighted, as the chairman mentioned this morning, to have representatives from a number of organizations in the privacy-advocacy consumer, academics, and from industry. We value all of your input as we strive to protect consumers navigating the marketplace. And last, but very much not least, I want to congratulate the FTC staff who worked tirelessly in putting this workshop together. We and myself in particular are very, very grateful for your efforts. So, what better way to end my little afternoon opening remarks than to return to Paul McCartney? When he saw the face that he thought was "just the girl for me," he wanted all the world to see that we've met. In 1965, he did that by writing a song. Today, he could have just tagged a photo in Facebook. Tomorrow, who knows? I think we'll all have a better idea of what the future holds after hearing from our panels this afternoon, so thank you very much for being here. [Applause]

>> Manas Mohapatra: I think we're gonna get started straight through with our third panel of the day.

>> Amanda Koulousias: Good afternoon, everyone, and thank you, again, for joining us for our afternoon panels. My name is Amanda Koulousias, and this is Jessica Lyon, and we are both staff attorneys in the Federal Trade Commission's Division of Privacy and Identity Protection, and we will be moderating this first panel of the afternoon. Before we jump into the substance of the panel, I just want to go over a couple of administrative details. We'll begin the panel with some of our panelists giving presentations, and after they have presented, we will have a facilitator discussion exploring some of the issues raised in the presentations. For those of you here in person, there are question cards in the folders that you should have received when you walked in. If you have a question that you'd like to ask one of the panelists, you can fill out the card, and there will be FTC staff circulating. If you'll just raise your hand, they'll come and get the card and bring it up to the moderators, and we will try to incorporate as many questions as we can into the discussion. If you're watching on the webcast and would like to submit a question to our panelists, you can also e-mail it to facefacts@ftc.gov, and somebody will be monitoring that e-mail address and bringing the questions, as well, to the moderators. It's definitely possible that we won't get to everybody's questions, but we'll try to incorporate as many as possible. Moving into this panel, this morning we heard from a variety of panelists about the use of facial-detection technology, where consumers are not being individually identified, but their demographic characteristics or emotions are being

ascertained. On this panel, we will be taking it a step further and discussing facial-recognition technology that can and does identify individual consumers. We'll be hearing from both Google and face.com about how they are currently implementing this technology in their products and services. And as Commissioner Brill mentioned, we'll also be hearing about the recent study out of Carnegie Mellon that was using facial recognition and its resulting policy implications. And then finally, during our discussion, we'll be discussing the potential privacy concerns that may arise from both these current uses, as well as possible future uses of facial-recognition technology. And I would like to just take a moment to mention that we'll be focusing on commercial implementations of facial recognition, and we won't be discussing law enforcement or security uses. And with that introduction, I'd like to now just introduce our panelists. In order of how they are sitting, closest to me is Chris Conley from the ACLU of Northern California. To his left is Ben Petrosky from Google. To his left is Gil Hirsch from face.com. And finally, on the end is Professor Alessandro Acquisti from Carnegie Mellon. And on that note, I'd like to turn it over to our first presenter, Ben Petrosky from Google. Ben is a product counselor at Google, Inc., and he focuses on consumer products, including Google+ and Picasa photos. And in that role, he counsels product teams on Internet law, privacy, and intellectual-property issues. Ben?

>> Benjamin Petrosky: Thank you, Amanda and Jessica. My name is Benjamin Petrosky, and as Amanda said, I am an attorney for product counsel at Google. I'd first like to thank the FTC for the opportunity to participate on this panel and for hosting it with technologies that we're very excited about, and we're happy to be a part of the discussion. The title of this panel is of course "Facial Recognition? What's Possible Now and What Does the Future Hold?" To begin, there's three areas I'd like to talk about today and give a little bit of background on the research and development that's going on at Google around vision technology, in particular, the areas of pattern recognition, facial detection, and facial recognition. And, of course, just to sort of base line on those, for pattern recognition, I'm referring to the analysis of images to detect similarities in general patterns for facial detection, of course, referring to specific analysis of images to locate patterns that a computer believes correspond to faces, and, of course, for facial recognition, comparison of patterns that are determined to be faces by facial detection to determine similarities between them or possible matches. I'd like to make clear a bit, also, what Google's approach is to this technology. Throughout our research efforts, over the years, our approach has been to treat this very carefully.

And as we've said publicly, we don't want to deploy this technology until we feel like it's ready and that we have the privacy protections in place. To highlight the privacy principles behind the way that we think about this, I'd like to just go through them very quickly. These are the intersection of technology and privacy, and these are what sort of guide us in our deployment in building of products. First, we want to make sure always above all that we are building valuable products that users would like and that they find useful. Of course, those are built with strong privacy standards and practices in place. And there's got to be transparency around how the products work, meaningful, relevant choices that give users control and security. At the start of the talk, we talked about what's possible now and what does the future hold, so I'd like to start with what's possible now. The first example that I've got on the slides here is a screen capture of two search-engine results pages for Google image search. I imagine most of you will recognize this. This is an example of how Google has deployed facial detection. What you may not be able to see from where you're sitting is that these are both queries for the word "rice." Now, many of you can't see what the specific images are in each. You can at least tell that on the left-hand side, the images that are coming back in that query are very different than those on the right-hand query. Google offers a number of refinements in the image-search product. Among them are color of the photo, size of the photo, whether it's clip art, line drawings, for example, and one of them is also face. For that, we simply refine the queries to only return images that facial detection believes contain a face in them that would have otherwise been returned, and, obviously, they get promoted higher. So, in this case, on the left-hand side, you see photos that are I think exclusively, actually, of the grain rice. And on the right-hand side, you see pictures of Anne Rice and Condoleezza Rice, and Damien Rice, also. The next example I'd like to talk about is from Google Street View. This makes use of both facial detection and pattern recognition. As you know, the cars in street view drive down the street in the daytime, and throughout the course of that driving, they may, of course, capture images of pedestrians on the street, and they may also capture images of other cars on the street. As Fred Carter mentioned earlier, this is one of the examples of where this technology can be used to blur imagery from these captures. In these, we identify areas of the images that we believe to be faces or to be license plates, and then algorithmically, we blur them so that this helps to protect the privacy of those individuals. And, of course, technology is not perfect, so if something is missed, there is a report link, and users can submit those, and those are manually reviewed and then blurred. The next example is an area where we use facial recognition. Now, you may not have seen this yet, but you may have heard of it. It's from the Ice Cream Sandwich operating system, which is the newest operating system from Android, and it's available on certain Android devices now. What you're looking at is a screen shot of a feature called Face Unlock. This is the low-security entry point for users that essentially most likely will have no security on their phones. As you know, a lot of users just simply have their phone sort of swiped to unlock, have nothing there. And this is a good entry point to help someone who may not be ready to start typing in an onerous password or deal with that in some way, and this gives them something more than what would just be the default on the telephone. The way that this would work is the user goes into the settings, takes a picture of herself. This picture is then stored on the phone as the key image. And then when the phone is locked, the user would like to unlock it, she essentially performs the same action. It takes another image and compares the area that is detected to be her face, recognizes that, and compares it between the key image. If there's a match, the phone is unlocked. If it's not a match, it's not. The next example that I'd like to talk about is another example of facial recognition, and this screen shot is from Google's Picasa desktop client. This, in case you're not familiar with it, is Google's desktop photo-management and editing software. And what you're looking at here is actually a screen shot from my Picasa. That is the face album for me. And what you see is a bunch of photos of me, as well as a bunch of photos that Picasa thinks correctly are me that I have not yet accepted that are there. You see the check mark and the "X" underneath those. With Picasa desktop client, we allow users to locally manage face models on their computer. These models can then be used in turn by the user to help to organize their photos. Obviously, we're in a time when the proliferation of digital cameras, they're everywhere. People are taking photos. They're getting CDs and photos from friends, e-mailed photos. This is way for a user who wants to organize these on his computer to put them all together, organize them by faces, see whose image albums, kind of quickly go through them. Also to, of course, tag the people in the photos if they'd like to. The next example that I'd like to talk about, the final example of what's possible now, is this feature called Find My Face in Google+. As I explained at the beginning of my comments, we like to build these products at Google with the privacy principles firmly in mind. A marquee example of how we've done that is our Google+ network, which has baked in privacy from the beginning and has been developed in that manner. Earlier today, we rolled out a brand-new feature called Find My Face, and this feature allows a user to bring some of the benefits of facial recognition from the desktop, the version that you saw through Picasa, to Google+ in a privacy-sensitive way. Now, obviously, since it launched this morning, I'm sure that none of you are familiar with it, so I'm gonna walk through briefly what exactly it does. So, when a user in Google+ accesses a photo or uploads a new album, he's gonna see this notification. And the notification, which I'm guessing you probably can't read it from where you are, is actually reproduced on a handout that we have which we'll make available probably on the table outside if you'd like to get a copy of that to read it more carefully. But this notification will be seen. It explains the feature and core language to the user and gives the user a choice to opt in to turn this feature on or not, to say no thanks. There's also a "learn more" link that will take them to a help-center article with a little more detail, if that's something that's interesting to him. For the user that chooses to opt in and turn this on, we then use existing tags of that user to build a face model for that user. This face model that's built is dynamically created with existing face tags on the service of that user, and the model can then in turn be used to suggest name tags on photos both for that user and people that that user is connected to. While desktop computers obviously have a lot of great storage and management capabilities, we see that a lot of people are clearly moving photos into the cloud, and they want to share them on social networks, and we think that this feature is a great way to help with that. For example, as Commissioner Brill noted in her opening remarks, you've got the example of the party freshman year. The photos are uploaded. Now, if you turned on this feature and chosen to opt in to it, if your friend uploads an album, she's gonna be suggested hey, do you want to tag Ben in that? And if she tags me, then I get an e-mail these photos are now online, and I know that this album is there, and I can take a look at it, and I'm aware of what's going on. It's a good way to do reputation management, and you can see photos that you might not have otherwise have known got uploaded, and it gives you a little bit of visibility into the online reputation. The feature, of course, preserves all existing defaults, Does not change the sharing settings of any albums, and it doesn't change any of the settings that you may have made about which tags are auto-approved. And, of course, you always have the ability to approve and reject any individual tags. So if you get a notice that somebody has tagged you in an album and you don't want to be associated with it, you can just simply delete that tag. Although, of course, if you realize -- I'm turning it on -- that more than just deleting that tag, you don't want to be using that feature anymore, you can, of course, turn it off. And another great benefit of Find My Face is that the user who has opted in can turn it off at any time. She can simply go to the top of the settings page, toggle the switch on or off, and at that point, the face model is deleted, and no more tag suggestions are made. Existing tags, of course, still stay on the photos, but those, of course, can also be deleted individually if the user would like to. Since the feature just launched today, I'm sure that none of you had a chance to tinker, but I hope that through the description, you can see the thought that went into this and the way that privacy has been baked in from the beginning with this feature. And the second part of the title of today's talk is "What Does the Future Hold?" And I'd like to just briefly finish my remarks with a couple of examples of ways that computer-vision technology, pattern recognition, facial detection, facial recognition might be used in some interesting ways coming down the pike. Obviously, the engineers that work on these offer a lot of promise, and this type of thing could make great contributions. Two specific areas, perhaps extensions of some existing projects that are already done, could be things that come around. As Chairman Leibowitz recognized this morning, child exploitation is a serious problem. In 2008, Google engineers teamed up with the National Center for Missing and Exploited Children and help to develop pattern-recognition technology that could be used to run across large collections of images of child exploitation to help collect those images together, select particular groups of them, and then to further investigate child predators. Another area this might work is in disaster relief. Looking forward, perhaps imagine using facial recognition to help with disaster relief and recovery, providing, of course, that there are strong privacy protections. Given the way that technology is already used for these helping in the aftermath of these events, you can see that at scale, this could be particularly significant. For instance, and this is existing technology, in the aftermath of the Japan earthquake, Google's People Finder helped to manage more than 600,000 records of missing person. That, of course, has no facial recognition involved in it whatsoever. But you can imagine if you were to apply this to a thought with images, for example, someone could upload a photo of a loved one into a service or provider that has images of news imagery from a devastated region, and facial detection or recognition could be done on those images to determine if there's any matches and possibly some more information about a loved one. The examples I have discussed have not yet been deployed or built, but they just give you an example of what might be done. Finally, to wrap up, taking it back to reminding that our privacy principles are always the guideposts around developing and building this technology. This is an area that we can continue to research and continue to develop very carefully. There's gonna be a lot of interesting ways, as I described, with just two possible things, but then you can just imagine beyond that that this could happen. And, of course, the privacy principles would be the matrix, delivering valuable features, delighting the users, having strong privacy practices and standards, offering transparency, relevant, meaningful choices for control and security. Thank you. [Applause]

>> Jessica Lyon: Great. Thank you, Ben. That was really informative, and I'm sure, you know, we'll have an opportunity to talk about it a little bit more when we get to the discussion section. For right now, I'd like to turn to Gil Hirsch. Gil is the cofounder and C.E.O. of face.com, the largest provider of face-recognition technology for Web and mobile services. Before founding face.com, Mr. Hirsch held multiple positions in research and development, product and sales, executive roles at Amdocs, Xacct Technologies, and Telrad Networks. Today, he'll be talking about his work with face.com.

>> Gil Hirsch: Thank you. You are a tall guy, aren't you? All right. [Laughter] Hi. I'm Gil Hirsch. I'm the C.E.O. and cofounder of face.com, and I'm here to talk about what it is that we do and how we offer it. I'd like to join my colleagues in thanking the FTC for setting up such an impressive workshop, and it's very thoughtful, and we hope that we can contribute to this discussion. So, who are we? We are a provider of face-recognition technology. But unlike many other vendors that offer this technology, we offer it as a service, so you don't have to install. It's easy to integrate. And our focus is on the online world, so not for shops or, you know, really the physical world, but rather, you know, we cover the online world. And it makes it a lot easier to integrate if the service itself comes in from the online world, as well. It works essentially the same, but because it's online, it has its own characteristics. But I want to share with you first is what we offer and then how our clientele has been using it. So, we offer two main features that you've heard of, two main technologies. The first one is face detection, and face detection is locate faces in the photo. Then the other one is deeper analysis, so alongside the detection value, we can also add gender, mood, and glasses on. We had a client that asked that. So we can recognize that. And then we also offer facial recognition. We offer it in a very specific way. And we started offering this a year and a half back. And back in the day, there was not even a discussion around privacy of face recognition, so we had to think a lot about how we foster responsible use. Because there's one thing that we're gonna talk about a lot here. And I've heard a lot, you know, in talks is how do you avoid what Alessandro is gonna be talking about, right, which is de-anonymizing people on the street or in a photo. And we do get those e-mails, you know? "Here's a missing child. Can you

find out who it is?" Or, "Here is, you know, a date, you know, from last night. Who is she?" Right? Our answer to all of these, by the way, is the same. We can't do that. Right? We can't do that. And the reason we can't do that is because we have rigged privacy from day one so it's limited to very specific context. And that's what I want to share with you. All right, so, face detection. We've seen a lot of different examples today. It's pretty amazing. Actually, I learned about a few new ones. I didn't know it was this advanced, and it's pretty awesome. You'll find that our stuff is a little bit more lightweight. We have released smile detection or mood and emotion detection. And one way we see this being used is by a bunch of hackers that came up with this solution called Emotional Breakdown. Emotional Breakdown scans The Guardian's 24 hours in photos and produces this pie chart of the emotions break down for that day. So was it a good day? Was it a bad day, happy, sad? And we've seen this being used in multiple different fun ways that go along those lines. Much heavier use of face-detection technology is for filtering and moderation. I guess the best example is on the far right here is chat services. Have you ever used Chatroulette? Raise your hands if you have. All right. Not a lot of you have, have you? All right. Let me tell you a little bit about that service. It started off as a bang. You open up this application. You immediately get matched up against another person that is sitting in front of a camera, and you only have this one button, "next," right? And you have a chat window to go along with it, so you randomize who you're chatting with. This was an amazing service. People loved it. Paris Hilton was on it. It was a big thing. But then it was abused, so sexual-oriented material started appearing, and before long, it crashed because nobody wanted to use it anymore. These video-chat services are looking to moderate the use of video. And the way they do this is they send us photos, like frames from the actual feed, and ask a very simple question -- "Is there a face there?" Okay? This is to make sure that there is a face in the feed and that it's big enough so that, you know, it doesn't leave much room for anything else to fit in. [Laughter] I'm trying to be careful here. [Laughter] So, we get pounded by requests, 'cause this is a video feed, and they sample quite often to make sure this works. And we have over 10 different chat services operating their services using us. Last, but not least, here's me with a fake mustache on top. Movember was a big month for us. If you don't know, it's a charity where you grow a mustache for charity during November. And we have four or five different services that allow you to put a mustache on your face. We also have a group of girls that developed the exact opposite, so remove the mustache from a face. They don't like that the way hipsters are wearing it right now, so they're all against it. So that's another way of using it. Glasses, I buy direct. A client of ours has used it to automatically place glasses from a catalog on your face, so you can have this try before you buy type of an experience. They have seen conversion go way up, just because you can automatically fit it on your face. Yeah, it looks great, yes. Up goes the sales. So that's a commercial use for it. But really, anything that can go on your face, we've seen it, right -- clown hats, your team colors, makeup, all of it, all right, on -- it's huge on the Web if you've seen Photo Booth and the likes. Face recognition. Face recognition, by far the largest use for it on our platform is photo tagging. For those who don't know the economics of photo tagging, photo tagging allows you to effectively share those photos so that you don't only post it for five seconds on your feed, but you actually get notified. Your friends get notified that there's a new photo of you. All depends on your privacy settings, of course and on the specific platform. This changes from platform to platform. But this is an excellent way to share. It's also a great way to organize your data so you can recall all of the photos at your end or one of your friends. Another use of this technology, which is almost the reverse of it, is something that Ben mentioned today. It's Find My Face. So, we've had that as a feature on Photo Tagger, one of our applications, and now we're working with two separate reputation-management companies, one in Europe, one in the U.S., to offer a service that scans photos that you have access to, your friends usually, your personal paparazzi group, right? They will post anything about you, sorry. And allow you to be notified on those photos that may be you, okay? It will suggest some photos. It's only high-confidence photos, but it says something like, "Is this you?" Okay, and so you can take action. This was a huge success. For anyone who used Photo Tagger, 75% signed up for this, and they found this very useful. So good choice on feature right there. And then look-alikes. So, a by-product of face recognition is that we'll come up with a few options. And the wrong options are people that are maybe similar to you, all right? We did not think of that ourselves. We were not even pushing this into the market. And we're saying, "Hey, we have this great look-alikes engine." We don't. But people are using it that way to match you up with celebrities, to match you up with -- Okay, so here's one site. It's called findyourfacemate.com. And what they're using is people that opt in, and then they're looking for people who are biometrically similar to them from the opposite sex. Now, I don't know if you buy into that, but, you know, that's what they're selling, and they're using our technology in order to do that. This is what it looks like. I don't have a demo, but this is, you know, an illustration of what our system would spit back when you send us a photo. So, this is me without this thing on my face. So we do the glasses analysis, the gender, obviously, how sure we are that this is a face in a photo. And then you can see the results. We're 82% in that case that this is Gil Hirsch. Then a number of my friends are there, as well, so I can put those -- one of them, actually, is a little bit similar to me. He's not me though. All right, how it works, very quickly. Face detection -- easy. You send a URL, and we'll spit back the answer. We don't provide you with any more information, aside from the mental data that we extracted. You won't get anything to hold, but we'll give you that information. With face recognition, we have to provide is reference photos of the people that you were looking for. So you can't ask, you know, "Here's a photo. Let me know who it is." You have to provide both, all right? "Here's a set of photos. Here's the gallery," right? "Here's the photo. Please do your comparison and let us know if one of those was found there." Nice work, Greg. So, I want to talk about privacy. This is about privacy, so... Since day one, we asked ourselves how do we avoid the one use case that everybody fears, which is to de-anonymize people. So, the first thing we said, "You don't know everything and everyone." That is, we're not going to hold a huge database of photos for you. We're not gonna hold a huge database of people for you to look for. You're gonna have to know both. So the input into our system is both the photos and the people that you want to have identified, and that will give you back the answer. So, in fact, you can never identify people you do not know. That's our mantra, right, is this one thing that we wanted To make sure that doesn't happen. And in addition to that, we've limited the scale, so you can only do that up to a certain amount. You can't flood our system with data. You can only set a certain scale. So the gallery is only that big. And within social networks, we have also added a concept of friends, so if somebody is a friend, we'll add that additional layer where we try to validate that. "Am I a friend of Greg before I can even ask for Greg to be identified?" Right, so I cannot identify anyone who is not my friend. But, again, keep in mind that we are operating a service. This is the things that we have applied under the hood. So, in addition to that, we are required by our terms of service that all our clients perform all the other actions that Ben was talking about and how the people get notified if they're tagged And that the use will be as expected by the system. We have, however, figured out that we've got to do our best to narrow down the usage to those areas where we believe that, you know, this is the stuff that we can do. Let's do it. So you cannot identify people who are not your friends on social networks or people that you have not been able to recognize before. And, again, like I said, our system is built in a way that you cannot do that, all right? The technology is built in a very specific way. You cannot do that. There's no database scan or one of those things. One of our employees, for instance, cannot do that either, all right? So, again, we don't support "Who is this?" but only, "Who is this amongst known people?" And by known, we ask you to prove it. In addition to that -- and, again, I'm not a lawyer. I'm not deep into the legal terms of how these things should be culled. We have our own names to certain things. But when we come to transparency and choice where we force all of our clients is to place a "Powered by face.com" notice. This is not to drive traffic to our site, because we're already high up on the SEO. This is so that users will know that there's face recognition behind the scenes. It's clickable. You can find out about what it is that we're doing and read more about it. We also offer an option for select social networks, like Facebook and Twitter, for you to place your I.D. only -- not your face. Just your I.D. -- and then no one can ask for you to be recognized in a photo anymore. So, again, we don't own the users. We don't interact with them directly, but we do have an option straight on our site. We want to keep the very minimal amount of data, both on our developers and whatever we need to operate. We don't repurpose data. It's only built for one purpose, which is to provide this service. It's not sellable in any other way. It's such a binary derivative that you cannot do anything else with it. And it has to change with every improvement that we do, so you have to provide us with data again and again for this stuff to work. And, of course, beyond applying security, whatever it is that we keep it's just, again, derivative data. It cannot be reverse-engineered into -- let's see -- the original photos. All right, stuff like that. So, again, we have put a lot of effort into this, but one of the reasons we're -- Somebody asked me, you know, very straight-up question, "Why are you here?" You know, very simple, if there's a framework an agreed-upon framework of privacy, we can apply that to our clientele, right? We want them to do that. And there will be a framework that they can follow and guide, and we're in a pretty good position to drive it. That's it. Thank you. [Applause]

>> Jessica Lyon: Thanks, Gil. Again, you know, we look forward to talking about these more during the discussion portion. For our next presentation though, we're going to be hearing from Professor Alessandro Acquisti. Alessandro is an associate professor at Carnegie Mellon University's Heinz College and the co-director of CMU's Center for Behavioral Decision Research. Mr. Acquisti researchers the application of behavioral economics to the analysis of privacy, decision making, and the study of privacy risks and disclosure behavior in online social networks.

>> Alessandro Acquisti: Thank you very much. And, really, I also would like to start thanking the FTC for taking leadership in organizing this event, which I believe is so important. I would like to describe some experiments we did recently and then after presenting the results talk about the extrapolations and the implications we can derive from the results. So, with Ralph Gross and Fred Stutzman. We started thinking, "If you take technologies which are available now, and you put them together into a big mixer, what can you end up doing?" And the technologies we were interested in combining were online disclosures, the increasing amount of information that we disclose online which can be found about ourselves online and especially photos. The continuing improvements in face recognizers, which have been discussed in the first panel this morning, the increasingly powerful computing ability that we have thanks to cloud-computing cluster, ubiquitous computing, which allows devices such as my smartphone, which does not have the computing power of Cray supercomputer, but can access the cloud cluster and do the operations for the cluster. And finally, statistical reidentification, which allows the combined data sets of sensitive and not-sensitive data creates something even more sensitive, or reidentify or de-anonymize data sets. So, we wanted to put all these together into one big mix and see whether, given these technologies, we can already do face recognition with a goal of reidentifying people online and offline, in real time, on a massive scale, in a peer-to-peer fashion. And by that I mean it's no longer the National Security Agency or maybe the largest corporations that are able to do that. It's any of us. Plus, can we also do sensitive inferences merely starting from a face? So, we did a couple of experiments. The first one was about what we call online-to-online reidentification. I realize that this is the version, the PDF, not the PowerPoint, unfortunately, so I cannot show the slides as they were meant to be. I wanted to show you the transition of this work. But it would be less clear in this fashion, but I hope you can still follow me. We will start from identifying images -- sorry, un-identify images from a popular dating site in the United States. Then we use identified images coming from Facebook. We use a primary images which are accessible through social engines, so we did not even need to go onto the Facebook. We simply use what, of a profile name and photo is accessible from the API or search engines, and we use face recognition to find matches between the two databases. If we did, then now we could connect a Facebook profile, which has a name, to the image on a dating site which is usually under pseudonym, so we could reidentify the dating-site profile that people wanted to keep private. In our experiment, we were able to decipher 1 out of 10 of members of the dating site, in the geographical area where we did experiment. Not nationwide,

a specific area. The second experiment was offline to online. We set up a desk on CMU's campus. And we asked people passing by -- they were prevalently students -- whether they wanted to participate in an experiment where we would try to find their face online and tell from that their name. If the subject agreed, he would sit in front of our laptop, would take a webcam shot of the student, and then we would upload his shot to a cloud-computing cluster. In essence, the student had a new way to fill out the survey. The survey was about usage of unlicensed or metro dating sites and so forth. While the subject was filling out the survey, on the cloud-computer cluster, we were doing the matching between the webcam shot and the Facebook images. By the time the subject had reached the last page of the survey, the page is being dynamically populated with the best-matching photos that the recognizer had found, and this subject had to indicate, "Yes, I see myself in this photo," or, "No, I do not see myself in this photo." This was our ground proof. And using this approach, we were able to identify -- that means find an identified Facebook profile matching the person for one out of three of our subjects. In this case, we have about 93 subjects participating in this experiment. So, so far, what we had done was to show that we can start from a face, an anonymous face in the street or online. We can use face-recognition to find a matching identified face, for instance, from a Facebook profile using PittPatt, our face recognizer, which was -- I wasn't, in the sense, from CMU. It was developed by other researchers at CMU. But two years ago, and some of you may remember this study, what we had done was to show that we could start from a Facebook profile or from just generally, broadly speaking, public-available demographic information about people, simply date of birth and state of birth. We could interpolate this data with, also, publicly available information from the Death Master File containing the Social Security numbers of people who are dead. And by using statistical processes, we could end up predicting the Social Security numbers of the people alive. So, if you do one plus one, combine the two studies together, you see where I'm going with this. Starting from a face, from a face finding in this case a Facebook profile likely to be matching the face, from the profile finding personal demographic information, passing this onto the algorithm to predict Social Security number, predict the Social Security number connecting the number to the face. We did these, and we had a 27% accuracy. We were four attempts for the first five digits of the SSN. But, anyway, We can also predict the last four, but statistical significance requires much larger populations than the population we had in this experiment. Now, mobile-ly speaking, the process I'm trying to describe here -- and, now, imagine literally that this is another PDF by PowerPoint presentation where the first thing to appear on the screen is on the top-left, then the top-middle, top-right, and then we go down following these arrows, okay? So, as I said, we lost all the dynamics here, sadly. I get an anonymous face, in the street or online. I find the matching face. Facebook's just an example. LinkedIn, organizational rosters -- there are so many sources of identifying faces online. I would bet, and from my experiment -- I mean, my test, as I was in this room, I would say that that's the case, that for most of you, of us in this room, there is at least one identified face online. Now I have a match, I have a presumptive name, presumed in the sense that it is probabilistically correct. From this name, I can find out their online public information -- demographic, maybe their interest for the LinkedIn network, for instance. And now with the public information, I can try to infer much more sensitive information -- the red box, and therefore I can connect the red box to the anonymous face -- a process of data accruation. Similar to how capital accrues over time, data about you can accrue as you combine more and more databases. And then we said, "Okay, let's show that we can also do it in real time," so we developed a smartphone app which takes a picture of a person on a mobile device, sends the information up to a cluster. On the cluster, it does exactly what our third experiment did, only the case was asynchronous. Instead, in this case, it does it in real time. It tries to find a matching face. If it does, uses the name to find the person's information, just simple demographics. If it does find them, uses the demographics to predict SSN and then sends all of that back to the phone, overlaying this information on the face of the person. So the story of augmented reality, the online populating and penetrating the offline world through our devices. So there are a number of implications here. We are very much interested in the ideal angles of these studies. Your face is a veritable conduit between your different online personas -- in fact, between your offline and online persona. It's much easier to change your name and declare "reputational bankruptcy" than to change your face. Your face creates the link between all these different personas. The emergence of PPI, Personal Predictable Information, what people may infer about you simply by looking at you, not using just feelings, not using what our natural evolution through hundreds of thousands of years gave us the ability to bloom, but using algorithms. The rise of facial visual searches, maybe one day search engines offering what they do now for text-based searches for face searches. The democratization of surveillance. And I'm not using this term necessarily in a positive sense. I'm referring to the fact that this power is now in the hands of all of us. What we did, anyone else can do. You don't need an incredible amount of money or resources. Social networks are becoming de facto I.D.s, because so many people are uploading good photos of themselves and using their real names. Therefore, what does privacy even mean in this kind of future? So, I was asked to discuss is this worrisome? If so, what are the scenarios I should be worried? Well, technologies such as face recognition can be used, as many other technologies, for good and for bad. So, you have at the center in the street that you recognize as a person that you met at a party. Good purpose, right? Especially a conferences such as this. We all would like to know, "Of course. Yes, Mark, how have you been since six months ago or two years ago when we met? Good." Or maybe the stalker will see you in a bar and finds information about you in real time and can know where you live. The brick-and-mortar store, which can recognize you and greet you, make you feel at home in the store. Good. Or the "Minority Report" story, which many of the panelists through today have brought up. And I would like to argue that as science-fiction as that story may sound, and it teases you in a sense that cures your devices, that's not yet allowed, the kind of future that "Minority Report" imagines as 2054. On the other side, there also are examples where we are really thinking too small. We are not thinking big enough. Yes, you walk through the mall, and the device knows that you are John Anderton, the name of the character, and it shows you a Lexus advertising with it saying, "Hey, John Anderton, this is for you." Let's push the envelope a little bit. Lexus has a Facebook app, which you downloaded, so it has access to your network of friends. So it knows who you're friend to. In fact, it knows if you're a guy which kind of girls you interact most with and therefore can infer who you find attractive and creates in real time some synthetic composite of who you find hot. [Laughter] And that's the person which appears on this screen telling you, "Maybe this car is for you." In fact, maybe also can do emotional detecting, and he knows that you're sad or happy and therefore chooses the picture, high-pitch, low-pitch, to send you that advertising. Where I'm going at is the fact that privacy is much less about control of a person's information and much more about the control that others can have over you if they have sufficient information about you. Technologies which should empower us can be used to control us, to influence us. And as we are talking about influence, there is another scenario, the large-scale, real-time surveillance. Good usages -- stop criminals, stop terrorists. Shady usages. You might have seen this photo. It was taken in I believe it was Vancouver before a hockey game. Famous photo, because it's a composite, actually, of multiple photos. But the incredible thing is that you can zoom into the photo, and you can get to the level of the photo in the bottom part of the screen. Those are the same photo. You can zoom in into a square, a street with about 60,000, 80,000 people and go down. With the level of definition, it can show each single person. And you know what? In some years, with sufficient computational power, you cannot just identify these three people and all of them, all the rest, but, in fact, in real time see whether they are connected or not in LinkedIn or Facebook, by what degree of separation they are connected. In real time, you can overlay the online connection to their physical disposition on a place. It can be used for good -- avoid criminal or terrorist attack. It can be used as a way to control, also, your right of free speech or be anonymous as you go through a political event. Now, this, of course, is not happening yet. Currently, we cannot do face recognition of everyone, everywhere, all the time. There are a number of challenges -- how many faces, partial images, are really available to you Versus to big corporations versus the government. If we start using databases of your hundreds of millions of images rather than hundreds of thousands, which is what we did, well, you start having big problems in terms of accuracy and especially false positives. I used to joke about the fact that as you start working with hundreds of thousands of images, you realize that you are not a unique and beautiful snowflake. [Laughter] There are many people who look like you. And face recognizers, a problem with this, underperform humans. We used comparative subjects. Our students, our subjects, were sitting in front of us. In the street, you cannot really stop people and ask them, "Hey, can I recognize you? I'm a stranger. But, you know, stop for three seconds while I take your photo." Computational costs. The bigger database, the more computations you have to do, and even cloud-computing costs start being quite expensive. However, current technological and business threads do suggest that all these limitations are not systemic. They will fade over time. So consider how many images are now existing compared to what they were 10 years ago. Imagine what will be 10 years from now. In 2000, this extrapolation based on a paper, publishing site, academic journal would not be your photo shot worldwide. Unless you are a celebrity, your photo didn't going online. Only a miniscule percentage of this 100 million went online. In 2010, only on Facebook, only single 2.5 billion photos uploaded online. So you can imagine this future in which these databases are constructed by downloaded, publicly available images, as we did, hacking into databases or maybe private-sector databases which either sell data to other entities or simply become identification or identity-provider services. Let me give you an example, and I use Facebook because not it's the only company -- as I said, there are LinkedIn and organization rosters -- but because of the size, the sheer size. Facebook, as you know, makes primary profile photos public by default and wants people to use their real first and last names. According to our surveys, about 90% of American users do, in fact, use their first and last names, and your name is also public by default. If you want to use your face for your primary photo, of course you're free not to do so. But we wanted to estimate how many people do use their faces, and Ralph just passed me the numbers a couple of days ago. We randomly sampled out of the publicly available Facebook directory, without even needing to log in Facebook, about 2,000 profiles. About 49% of them had a unique face. About 60% of them had at least one face. So focusing on those with a unique face were arguably likely that's really the person. And knowing that 90% on average have real first and last names, out of 800-plus million users, these suggest about 330 million unique-identified faces accessible through their public directory. Accuracy -- I will go quickly here, because, as mentioned this morning, also, by Dr. Phillips, the accuracy improves by about an order of magnitude every four or five years. So comparing 1997 to 2010, dramatic increase in accuracy. And, in fact, researchers are very well-aware of all these problems with lighting, facial position, facial hair that make computer face recognizers still underperform human face recognizers -- us. In fact, not only that, but consider the following issues. This morning, we were talking a lot about facial identification, but we humans do much more than that. We do people identification. When we recognize each other and we avoid the problem of false positives, we're using not just facial features. We use the head, the shape of the body. Because we recognize people in real time as they move, like in a video, we also use the way they move. We use the way they dress. We use holistically all this information. Now, most facial-recognition research focuses only on face, as the name says. No doubt, 5 to 10 years from now, face recognizers will incorporate all this additional metadata. Where will this data come from? Can you guess? Online social networks, mostly. And finally, about accuracy and the possibility of doing these with devices. Well, when we went public with our results in August, we were discussing with Ralph, "Should we talk about this tonight? We did it with a smartphone, okay -- his smartphone. We took pictures in real time. Get back the information in the picture on the phone. Said, "Well, we can also say that you can use sunglasses." Fair enough. There is the example of the Brazilian police which apparently, allegedly, is working on that for I guess the Olympic Games in a couple of years. And then I said, "Why don't we say, also, contact lenses?" And Ralph said, "Let's not do it, because it will kind of decrease the credibility we have with all of our results. It will sound so much like science fiction, that you can contact lenses for face recognition." And although it is still science fiction, let me show you what happened a few weeks ago where a team of researchers from different countries developed the first contact lens which contains not just the LED, but the wireless antenna. So this contact lens can connect through some other server, through Wi-Fi, and has a LED to project information which could come through the antenna. It has not been tested on humans yet. That's not a very airy, human eye. It's a rabbit eye. But the rabbit did survive. [Laughter] So 5, 10 years out, you can imagine what sounded crazy to Ralph and me just three months ago, which is your eyes looking at people and highlighting information in real time. Plus, there are these business trends. I will go quickly here, because we have many representatives of the companies who are involved in face recognition, and we have already seen how hot and stimulating the business environment is in this area. So the short is, currently, what protects us -- if you feel that we have to be protected. I would grant you that we can debate about this. If you feel that we should be protected from the kind of future that I alight, where anyone could look at you and predict sensitive information about you, what protects us are mostly false positives, so the scale of the problem, and regulatory self-constraints. So the issue is, for how much longer this will act as protection? Let me quote and say I didn't -- here is completely my fault I couldn't put -- the slides, I gave them too late, because we were crunching numbers. We were on supercomputers, and they only produce the results this month. I'm kidding, I'm kidding. The numbers are not -- they didn't need supercomputers. They simply were back-of-the-envelope estimates, extrapolations that we did. Think about this. Currently, we are protected by self-regulation in the sense that, for instance, "Gee, very eloquently notice how protective of the usage of facial information they are. Now, compare now to 10 years ago, how much information about you was available to others, to corporations, for instance, years ago. You would be surprised if you could go back in time and think about what we have now. So extrapolating years out, what will be more accepting in terms of information known about us. And now consider the technology. Today, with the cloud-computing cluster we were using, which costs $2 an hour, we were able to compare a pair of images in 0.0000108 seconds. If you wanted to do what we did, rather than how we did it, which was just hundreds of thousands of individuals, if you wanted to do it nationwide, say 300 million people, we would never be able to do it in real time. It would take four hours for each face that we tried to match -- four hours. Impossible. Imagine 10 years out, 2021, the population of the United States will be about estimated to be 337 million. Consider just the population which is 14 years and older. Let's assume there is Moore's Law, and therefore a certain improvement in cloud-computing power over the time. Let's assume simply the only pre-massaging of data we do is that we split male images from female images. 10 years from now, we could compare one random shot taken here to everyone in the U.S. in five minutes at $2 an hour. If you want to spend more or if you assume that competition will bring the cost of cloud computing down, you could spend $60 an hour, and this comparison can be done in 10 seconds. My point is that we can get really much closer to this future than some of us may have believed. And to conclude, how do we -- if we feel, once again, that this is a problem to be solved, how do we solve it? It's very simple. Oh, we want to have facial recognition because the technology can be used for many good purposes, but we want to split the bad from the good purposes. Well, how do you define bad and good? How do you define what is criminal and what is not if social norms change over time? And once you find it, how do you distinguish? I will probably stop here, because I guess I went down a little long. Thank you for your time. [Applause]

>> Jessica Lyon: Thank you, Alessandro. Your presentation I think raised a lot of really great issues that, you know, we're here to discuss today, and I think they're a great lead-in to the discussion portion of this panel. Thank you, also, to Ben and Gil. I think you provided us with a great backdrop to start discussing some of these issues. To start off, I think we'd like to begin with the topic of consumer awareness and consumer control when it comes to facial-recognition technology in commercial applications. And I know this was touched on briefly in some of the presentations already, but I'd like to turn this over to Chris Conley from the ACLU of Northern California to get us started. Chris is the Technology & Civil Liberties Fellow at the American Civil Liberties Union of Northern California, where he focuses on the intersection of privacy, free speech, and emerging technology. His current focus is the Demand Your dotRights campaign. It's a multifaceted campaign to protect individual privacy rights in a technology-rich world. So, Chris, we'd love to hear your perspective on whether consumers are aware of how facial recognition is being implemented in commercial applications and whether the current commercial uses tend to be transparent enough that consumers understanding what is happening, what the potential implications are. And then, just in addition, since we haven't heard from you yet, if you'd just like to give sort of any other thoughts or your perspective on sort of some of these issues surrounding facial recognition.

>> Chris Conley: Sure. So, I'm somewhat glad I didn't prepare slides, because they'd be pretty much redundant with all the wonderful presentations already, in this panel and the previous panel.

But there are a couple of topics I really want to touch on. The first is consumer awareness, as you asked. And the reality is, consumers are always behind technology in terms of what it can capture, how much information is out there, how easily it can be aggregated. You know, step away from facial recognition. The previous panel talked about an easy method of tracking people -- a supermarket loyalty card. So, how many people in here have a supermarket loyalty card? Lots of us -- well, actually, far fewer than in most rooms, I expect. So, for most consumers, this is basically a coupon. I swipe it, and, hey, I get discounts. You know, a few think of this as "Okay, well, I'm telling them what I'm buying. But, you know, it's just a card, and I lied to them when I filled out the application, so they don't know who I am." You know, It takes a lot more thought and a lot more experience to realize not only is this an aggregate of not just what I thought, but when I bought it, how frequently I buy things. You know, every third Friday, you buy six cases of beer. What's going on? [Laughter] You know, but it's also the kind of information that can be aggregated and can be linked to other information. You know, I paid for that with a credit card once. You swiped it, there's my name. You know, there's my credit-card number that could possibly link to other purchases. So the kind of information that's available about users through technology like loyalty cards of facial recognition is much robust than consumers realize. That's only part one. The second part, of course, is at least consumers realize when they're using the loyalty card. And with facial recognition, it's entirely possible that people don't know when a picture's being taken, because it's a stranger on the street who looks like they're on their cellphone, and they push a button, because it's taken from 10,000 feet up in a satellite that has incredible resolution. There are all sorts of pictures of us that we have no idea are being taken, that we forget about. We forget about we're on webcam right now, and this is being logged and recorded. And with facial recognition, if you connect this to a picture of me through an automated tool, it is now part of my permanent record. And this kind of -- you know, the advances in facial recognition, the advances in the ability to connect one photo to my identity presents a lot of threats. Alessandro went through several of these. But there are kind of two different contexts that get blurred here. The first is the context of, you know, I don't know you, but you can look at me. What can you know about me? And in that case, you know, traditionally, you think, "Oh, you don't know much about me. You know I'm talking up here. You have my name tag. That's about it." But if you can use a picture of me or my identity to link to my Facebook profile, to link to my purchase records, to link to whatever other data out there is available without me realizing it, it really changes my ability to control context, to go to a bar and be anonymous, to go to a protest, to go to, you know, a support group, to walk into an Alcoholics Anonymous meeting worrying that, hey, somebody took a picture of me outside the door, and they know who I am and what I'm going in for. And that turns into the second context, which is the fact that I don't necessarily want everything I do and everywhere I'm photographed to be knowable by everyone else I know. You know, I want the ability to control who sees which rallies I go to, which meetings I go to, where I am and where I'm not. And if facial recognition means that anyone in the world can just say, "Hey, I want to know where else Chris has been. Show me all the pictures of Chris," that's a little bit frightening. And I want to get a little bit more broadly into who can use this, because, as Jessica and Amanda said, we're talking about consumer facial-recognition technology, but the fact that it's developed for consumers doesn't mean it's only used for consumers. And certainly, if we have issues with the government or other third parties who are not authorized users of this service, but are still taking advantage of it to track our whereabouts, to log our interests and our passions and, you know, whatever we do, this raises serious concerns from a civil-liberties perspective, from a consumer-privacy perspective, from all sorts of perspectives. So, what I really want to know is, what I want to work with, is how do we actually enable consumers to understand what's going on and to be aware of what's going on. You know, the first thing we've talked about is transparency, and both Ben and Gil really talked about how they've tried to make their service as transparent as possible so that people know what's going on, trying to make things opt in so it's not, you know, "Hey, you took my picture. You linked it to my Facebook I.D. Was that okay?" No. Too late. That doesn't help very much. So we really need controls that give people real choice in the transparency so that if my information is being linked to my photos, I have the choices, I have the control. Digging even deeper, you know, can I opt out? Can I delete my picture? And if so, does that actually delete my face print or my Eigenface vector or whatever you want to call it, depending on how technical you're getting. What exactly am I controlling, and how am I able to see what the service knows about me, what it's doing with that information? Can I choose could who can find my face and who can't? That's actually one of my questions for Ben about their new service is can I choose? I can choose, yes, I opt in or opt out, but can I control with more fine granularity? And these are some of the things that we really like to see. And so I'll wrap up my little presentation by saying we thought about this. We're far from the only ones. Obviously, the Canadian, Ann Cavoukian from Ontario, thought a lot about privacy by design. But we have a publication called "Privacy and Free

Speech: It's Good for Business." Because especially in areas like facial recognition where there's a real threat of the "ick" factor -- this is creepy. This is scary -- you know, it is something that businesses need to think how are they going to respect privacy while building a product. Because if they don't, consumers are going to be outraged when they found out that, "Hey, you just showed me an ad, and, oh, there's a camera behind there. Oh, and, by the way, you remembered me from the last time I came, because the same camera took the picture. And you linked it to my name because i swiped a credit card, and you linked it to my Facebook profile, and so now you're selling me ads based on the wall post I put yesterday." That's not gonna go over well if you aren't transparent and if you aren't really thinking about how to build user trust so that you can leverage that information while protecting user privacy and giving individuals the ability to control their own information. So, I have more copies of this. There are a few out front if you want it. And, of course, we'll be around for questions. But I want to make sure there's plenty of time for conversation.

>> Amanda Koulousias: Thanks, Chris. You know, I think you touched on a lot of the issues that we were hoping would come up in the discussion. And, you know, given, like you said, how facial-recognition technology can be used in a less than transparent way, I think consumer control is particularly important. And we've gotten a lot of great questions from the audience, so I want to try to incorporate as many of these as possible. And a lot of them do focus on the issue of consumer control and consumer awareness. Ben, one of the ones that we've gotten is in reference to Google+'s new Find My Face feature. And somebody wants to know, "Will I be suggested to people who have me in a picture, even if I haven't added them to one of my circles?"

>> Benjamin Petrosky: Okay. And I think that that was a question that sounded like that Chris was also wondering about in terms of the scope of who you're gonna be suggested to as a possible face tag. And I think that just as a base line to make sure that everybody understands what we're talking about with Google+, the basic sharing model of Google+ is based around there are circles, and you decide to put people into which circles, and that's a good way to sort of control and disseminate which information you want to share with people. At launch right now, the feature is just simply on or off. It's an opt-in completely, and at that point, we use who we believe you know. And that is working on a dynamic model along the lines of what we use for something like social search I've been using for a while, and we think that that presents the most value to the user in this

case. For instance, obviously, if the user doesn't want to opt in, they don't have to. But With the example it would look at things, for instance, like a circling relationship or a bidirectional, circling relationship, or perhaps if you had e-mailed somebody in Gmail and had, you know, a series of conversations with them in that, it might use that information, as well, so those kind of vectors. And, obviously, as people start to use the feature and if, for instance, we start to hear from users that this is something that would be very valuable, obviously, we're continuing to develop it. And, you know, we always look for feedback and are interested in hearing those things, so we definitely welcome any comments.

>> Amanda Koulousias: So, then, exactly who right now would be suggested to users? I mean, I know you said, you know, social, but can you be a little bit more specific?

>> Benjamin Petrosky: So, the way that it's described to the users is people that you know or that we think you know, and so it's essentially starting I think with looking at, like, the circling relationships and then bidding out of that. There's a number of different types of affinities. For instance, if you are posting frequently and plus mentioning somebody, if you were continually sharing albums with somebody, if you're always sort of plus-oneing somebody's content, those are the kind of things that could play into this factor over time as the future develops.

>> Amanda Koulousias: So, then, would that go both ways? So, for example, if I am frequently sharing albums with a particular person, but they are not necessarily sharing albums with me, you know, maybe they have a different view of our relationship than I do. [Laughter] So would they be suggested to me in that instance?

>> Benjamin Petrosky: I know, and that's exactly the kind of work that's going into the development, because that's one of the benefits, I think, of doing a slightly dynamic model is that it doesn't have to necessarily be a bidirectional sort of if I get suggested to you, you would be suggested to me, and it would allow for allowances to that type of thing. If you see somebody who's just, you know, aggressively plus-mentioning somebody and is never being responded to, that might be a signal that that person isn't actually connected to that person.

>> Amanda Koulousias: So is the somewhat evolving nature of this explained to consumers on the Google+?

>> Benjamin Petrosky: I think that the way that we tried to describe this is making it clear in the sense that it is not just limited to circles, for example. The future, obviously, as I mentioned, is a complete opt in. So if it's something that there's any concern about or any worry, the users don't need to engage with it, but using the description of people that you know, and that's the goal that we have, of starting with a small group of people who we have strong confidence that there's a relationship that you would, you know, most likely want to be shared with there and then sort of developing that as time continues.

>> Amanda Koulousias: Okay. You know, you just mentioned that users don't have to opt in to it if they don't want to. A couple of our audience members have actually posed the question, if a user chooses not to opt in to this, is there some other easy way for them to scan the service to determine all of the photos that they themselves are tagged in in order to remove those tags?

>> Benjamin Petrosky: Yeah, absolutely. If you go into the photos tab of Google+, there's a "photos of me" section, and you can click that, and that will link to photos that are linked to your profile. And, of course, you can just click on those and remove any tags that you'd want from there. And just also to mention, I don't know if people -- I see a few laptops around here. I don't know if people have tried to look for this. But as with several of our features, it's a rollout over a course of time, so I'm not aware that it will necessarily be available on everybody's account today. So that, should be, I think by early of next week should be available to all accounts.

>> Amanda Koulousias: Okay, great. And, Ben, not to pick on you, but I know we have some questions for Gil in a second. But, you know, I think one more question about this Find My Face feature, actually, a couple of the audience members and some people who are watching on the webcast have raised this. Are there safeguards that exist that would prevent somebody from using somebody else's face and then being able to use Find My Face to find all of the photos of that person, as opposed to themselves? So, for example, if I uploaded a photo of Jessica?

>> Benjamin Petrosky: Right. So, if the model requires that the faces be tagged to your profile, so assuming that you have, you know, a number of photographs that are tagged to yourself and you, you know, tag the president in a photo as yourself, or you tag a piece of a tree or something, obviously, the statistical algorithms are gonna look at those and are going to have more heavily weight to the information from the photos that are sort of the group there.

>> Amanda Koulousias: Okay.

>> Benjamin Petrosky: And thank you for the feedback. This is great. And, of course, we always, as I mentioned before, are interested in hearing and do appreciate any of this kind of feedback, so we would welcome it. Thank you.

>> Jessica Lyon: Excellent. Thank you. So, we have an audience question for Gil, actually. So, the member of our audience was wondering if a consumer requests that you delete all images of him or her, do you honor that request, and how does that process work?

>> Gil Hirsch: Oh, absolutely. So, there's a way to opt out, again, out of an existing social network. The one feature we do not support currently, because it's technologically not possible, is for you to upload your face. It also has privacy implications on itself. But to upload your face and ask us to search for that face in every photo so that we can rule you out, that doesn't work. So, instead, what we allow you to do is identify yourself using either a Facebook connect or Twitter connect for now. We'll be happy to do this for Google+, as well, once the API is there. And then you only save your I.D., and, then, what that means is every piece of data that is associated with that I.D. not only doesn't get saved, it doesn't get served. We ignore it completely.

>> Jessica Lyon: Okay, so I load a photo I.D. of yourself, I.D. it as yourself, and then thereafter --

>> Gil Hirsch: But no need for photos.

>> Jessica Lyon: What?

>> Gil Hirsch: No need for photos. You all know Facebook Connect, just as an example, right? So, you click a Facebook Connect, and your I.D. on Facebook is the only information that it will keep. It's like a blacklist of people that if these people are being asked for, they're ruled out automatically.

>> Amanda Koulousias: And is there a way that somebody could delete any information face.com might have about them if not connected in any way to one of the social-networking services?

>> Gil Hirsch: So, we haven't found a reasonable way to do that yet. We do, however, require that if you're using our system, not through an existing identification system, like Facebook, Twitter, later on Google+, then you should enable your users to do that.

>> Amanda Koulousias: And, so, is that in contracts that you have with developers who might use your service?

>> Gil Hirsch: Sure.

>> Amanda Koulousias: It is?

>> Gil Hirsch: First of all, it's a mind set, right? I mean, the legal stuff is less interesting. We are after every operator of a service that we think is inappropriate. We haven't found any yet that are significantly as such, but when we do, it's very easy to shut them down. It's one of the benefits of operating from a cluster, not on a device. Can just shut them down.

>> Amanda Koulousias: All right. I think it looks like we're running a little short on time. I think we have some leeway here, because we're heading into a break afterwards, so I think we have a few more minutes to go through some more questions and maybe follow-up. I know one of the things that we wanted to touch on on this panel is the possibility of facial recognition being used to identify populations that might be more vulnerable to harm than other populations, such as either children or possibly a domestic-violence victims. And, Chris or Alessandro, I was wondering if

either of you have any thoughts on, you know, any privacy concerns that might be raised by facial recognition being used to identify children?

>> Chris Conley: Well, facial recognition -- I mean, one of the concerns that does raise is things like false positives, because, obviously, if you're using a tool that is intended to highlight some incredible wrongdoer and it instead flags the wrong person -- and, again, we have the question of as you increase this in scope, how accurate is the tool going to be? That's certainly a concern that we have to draw. The other concern, really, with tools like this is, you know, as Alessandro said -- and I should probably let him say it, but I'll say it for him anyhow is once you have a tool that has one purpose, it's very easy to repurpose it. If you have something that is built so that you can take an enormous number of photographs and scan for a particular image, you can say, "I'm only going to use it for purpose 'X,' well, purpose 'Y' is another good purpose, too, and, well, we might as well do a purpose 'z,' too, and suddenly you have a general-purpose facial-recognition platform that, you know, may carry more negatives than positives. So if there are ways to constrain it so you really can get the benefits without having the possible consequences, that would be great, but that's hard to do and certainly requires a lot of design and forethought before you roll something out, rather than saying, "Well, we've got this, and we can do it for this. That's great. We'll figure out the rest later."

>> Alessandro Acquisti: Extending on what Chris was mentioning, and a broader issue here is the fact that I do believe that these systems based on inferences will become increasingly accurate, but will never be 100% accurate. There will always be some element of noise. However, as the accuracy gets smaller and smaller and smaller, we tend to trust the system more and more, and arguably it will be possible that the errors, although less likely, will become more and more damaging, precisely because we put so much trust in the system being correct that we trust it, also, for more and more sensitive decisions. So this is troublesome. The other potential concern -- and I was referring, hinting at that showing the contact-lens scenario -- is that you can imagine a future where all communications will be computer-media communication in a sense that nowadays, we think of computer-media communication as what happens for your laptop, your cellphone, through your iPod. But imagine that you have these contact lenses which in real time tell you what is the political, sexual, religious orientation of the person in front of you, their credit score, their favorite

color. And now you're using all this information and making decisions about the person based not on what your gut feeling is telling you, but on this machine, on this algorithm. In a way, incredibly exciting from a side. Also incredibly creepy from another side.

>> Benjamin Petrosky: I just had one point as I was thinking through it. It occurred to me in responding to the question about the person who had asked about tagging faces that were not your own in your model, and I think that what the questioner might have been getting at was the idea of instead of having a collection of your own faces of just simply using the profile to tag, for instance, you see a picture of somebody in a club or on a train, and you take a single picture of them, and then you create this profile to use that. The way that the system has been architected with the necessary social connections in place makes that impossible, so you're not going to be able to just -- this is not gonna be the equivalent of running a query across a massive set of information in order to return a result to identify that. Just wanted to make that clear.

>> Amanda Koulousias: Okay. Thanks. On that note, you know, I think I'd like to give each of the panelists just a brief maybe one-minute each to just, you know -- We are almost out of time here, so, you know, any final thoughts you have on, you know, either anything that was raised by any of the other panelists or just any final thoughts that you have on these topics generally that you'd like before we finish up? We can start I think furthest away. Alessandro?

>> Alessandro Acquisti: Perhaps the only point I can add very briefly in less than one minute is this issue of control, notice and consent, consent as a form of control. I believe that it's useful here to bring in the distinction between sufficient and necessary conditions. Control and notification are necessary conditions, in my view. We need to give control to users. We need to notify users. However, they are not sufficient condition for privacy. If you go back to the OECD guidelines for information principles, they were not simply notice and consent or control. There were six others. And we sort of lost track of them focusing only on the first two. And this is the challenge with notice and control, not that it's not good. It's not enough.

>> Gil Hirsch: Okay. [Siren wailing] What is that? [Laughter] All right. Your face. Time's up? So, we've been discussing a lot of different things. I think one more note to add is we're seeing an

incredible amount of data already out there, all right? There's a lot of public information already out there. Alessandro has pointed that out. I think one more area where we can look at to add control or privacy or at least think about those are the uses, okay? So it's another approach to how we deal with data versus not only "Is that data there? Is that data not there?" 'Cause in many situations, it's already there. But, rather, what are proper uses, what are not proper uses, you know, even without any specific consent? Because, again, it can always be abused. So what is that abuse line? How do you not cross it, you know? That will be very interesting.

>> Amanda Koulousias: Okay, thanks. Ben?

>> Benjamin Petrosky: I just want to say thank you again, and we look forward to continued discussion on this.

>> Gil Hirsch: Good one. [Laughter]

>> Chris Conley: Sure, make me look bad. [Laughter] So, I will echo that I think notice and control are necessary, but not sufficient. And I think what Gil touched on is something important, that it's not just what information is collected. It's also how long is that information stored, in what format. Is it reverse-engineerable? You know, are you retaining whole pictures or just the face prints or, you know, computational values you need to identify it later? It is certainly about use -- who can use the information, how can it be used, how often can it be used. Face's idea of throttling -- you know, you avoid abuse by saying you can only query so many times a day, if I understood that correctly -- is an idea. And there are different ways you can put use controls in and hopefully make them user-centric use controls -- or rather consumer-centric. Even if I'm not a user of the service, I have ways of limiting how it can be used to identify myself. And then Disclosure and sharing is a very important part of the control and notice, as well. How is this information being shared with third parties? How is it going to be handled if you get a demand for information? Do you have security in place to avoid breaches? You know, these are all different parts of notice and control that are essential. And then last and not least, since we're talking about the FCC enforcing these promises and making sure that they stand up is very important, and it's both in terms of if you are, you know, a company that has a privacy policy that has very clearly stated what notice and

control you're giving and you've reached that, there have to be consequences. And, as well, there have to be protections in the back end so that you can't be forced to breach those. One of our big concerns is with electronic-communications privacy law. And if the law says that someone can come in without a search warrant and just demand information from you and you have to comply, as a company, you can't do anything. And so we want to see stronger laws so that as this information is protected on the front end, you make your promises, you don't have to say, "Well, except if we're forced to disclose on the back end," because we don't have the security, we don't have the privacy law that we need to protect that. So that's what we would like to see.

>> Male Speaker: Hear hear.

>> Jessica Lyon: Thank you, all, for your thoughts and for coming here today and presenting to us. I think we all enjoyed it and learned a lot. We're going to take a short break right now and return at 3:00 P.M. for the final panel of the day, which will address the broader policy implications of both facial detection and facial-recognition technology. So, again, please return to the room by 3:00 P.M. thank you. [Applause]