SPRING SYMPOSIUM:

DATA PRIVACY & TRANSPARENCY IN PRIVATE AND GOVERNMENT DATA COLLECTION♦

INTRODUCTIONS

EDITOR-IN-CHIEF*

Good morning and thank you everyone for gathering here today for the *Cardozo Arts and Entertainment Law Journal's* ("AELJ") 2014 Spring Symposium. We welcome our esteemed panelists, scholars, professors, and distinguished guests. In this room are some of the most influential scholars, practitioners, and advocates in the fields of data

♦ The following compilation is composed of edited transcripts from *The Cardozo Arts & Entertainment Law Journal's* spring symposium, which took place at the Benjamin N. Cardozo School of Law on April 4, 2014. Permission is hereby granted for noncommercial reproduction of these transcripts, in whole or in part, for education or research purposes, including the making of multiple copies for classroom use, subject only to the condition that the name of the journal and speakers, a complete citation, and this copyright notice and grant of permission be included on all copies. © 2014 Cardozo Arts & Entertainment Law Journal.

* Francesca Montalvo, Editor-in-Chief, CARDOZO ARTS & ENT. L.J. Vol. 32, J.D., Benjamin N. Cardozo School of Law, Class of 2014.

privacy and information law. The dedicated editors and alumni of AELJ are also here with us today. AELJ would not be the distinguished journal that it is today without all of your help and commitment. A special thank you to our Acquisition Editors, Christina Noh and Mark Pellegrino, who have helped to develop today's topic. And especially, thank you Pamela Grutman, our incredible Managing Editor who organized today's event and helped bring together our distinguished panelists. Thank you for going above and beyond to making this symposium a success.

We would not be here today without the help of our faculty advisor, Professor Felix Wu, who helped us develop the notice and disclosure in data collection topic, and also helped us gather our notable panelists and moderators. We are delighted to have Professor Brett Frischmann and Professor Jonathan Manes from Yale Law School moderate our two panels. Thank you both for offering your expertise in guiding today's discussion. And last not but certainly not least, thank you to our honored and distinguished panelists. We are grateful that you are here today to offer your insights and opinions on how notice and disclosure practices play a role in government and private data collection practices and policy. All of the people who I have previously mentioned have helped make AELJ what it is today: a top law journal that has been cited three times by the U.S. Supreme Court as well as by the high courts of Canada and Australia. And, of course, AELJ remains a staple in the district courts and court of appeals. We are proud to maintain our rank as the leading journal in the nation for arts, entertainment and sports law and as one of the best journals in the State of New York for intellectual property.

AELJ remains at the forefront of recent developments in data privacy and information law. It is an honor that our journal can come together with our notable speakers and leaders in the field to discuss the issues and concerns involved in government and private data collection programs. "Data privacy," "data mining," and the "National Security Agency" ("NSA") are terms frequently heard in the news and permeate current legal literature. After Edward Snowden's controversial leaks that went viral last summer, lawyers and policy makers are left asking what safeguards can be put in place to balance the benefits of both private and government data collection with privacy concerns. Some advocates have called for federal rules and regulations to control data collection. Others propose certain limits on how much and what kind of information both the government and private companies can collect. Another theory, which we will discuss further today, recommends disclosure and notice to the consumers and the public whose data is being collected. Until recently, the federal government had entirely

prevented private companies from letting the public know that subpoenas had been received because of national security concerns. New guidelines permit limited information to be disclosed. To address concerns among their users and the general public, companies like Google and Microsoft have published transparency reports in order to advance their users and the general public's knowledge of how often they release their users' data to the government. Recently, major companies like Amazon and Twitter have gone as far as to fight for their users' privacy rights in court and with congress.

Today's panelists are legal scholars, advocates, practicing attorneys, and professors who each bring a unique perspective on the data privacy concerns in today's "Big Data" world. Our first panel will discuss data collection by private entities and their disclosure and notice procedures. It will focus on the impact of big data on consumers and if disclosure practices, new or old, could actually help alleviate data privacy concerns. Our second panel will discuss the government's surveillance program and the growing concerns raised by its data collection programs. In particular, the panel will focus on the NSA's surveillance gag orders and the disclosure and notice practices that have been implemented by major online service providers as well as the practices that these providers have adopted to make their users aware of the government's requests. The panel will explore the role of transparency in the national security context, what types of government disclosures should be required in this context, and whether transparency can serve as an effective safeguard for civil liberties. With that brief introduction, I turn the conversation over to the Professor Brett Frischmann who will guide our first panel.

PANEL I: DISCLOSURE AND NOTICE PRACTICES IN PRIVATE DATA COLLECTION

LORRIE CRANOR*

BRETT FRISCHMANN*

RYAN HARKINS*

HELEN NISSENBAUM*

BRETT FRISCHMANN: Good morning everyone. I also want to thank AELJ for putting together an excellent conversation. The issues both panels will be confronting are very important and timely. In fact, my view of privacy is that it's much more foundational to our society and who we want to be or become than conventional privacy law or theory would suggest, especially as it is tested in the emerging context of big data, wearable surveillance devices, ubiquitous distributed sensors automating financial media and other supposedly smart systems, as well as the Internet of things. So there's much to discuss and we're going to have plenty of opportunity to engage directly with the audience. Let me give you the plan for the first panel. I'm going to introduce our speakers very briefly. Each of them is incredibly accomplished and frankly, I could spend most of our time introducing them. Their bios are available on the AELJ website. But after brief introductions, each speaker will have fifteen minutes to talk. I've told each of them that at the fifteen-minute mark, if they continue to talk (and they've more or less consented to this because they've been informed), I'm going to start revealing personal, sensitive information about each of them. We'll see what happens. After the speakers, we'll open it up to Q&A. The student organizers have framed our panel in terms of "the role of transparency, disclosure and notice practices during private data collections." I've suggested to the panelists to talk about private data collection and the promise and peril of informed consent-based approaches to dealing with privacy concerns. Though the students focused mainly on the informed portion of that, we may venture a bit more broadly to talk about the relationship between various transparency, disclosure and notice

* Professor of Computer Science and of Engineering and Public Policy at Carnegie Mellon University.

* Professor of Law and Director of the Intellectual Property and Information Law Program, Benjamin N. Cardozo School of Law.

* Attorney, Regulatory Affairs, at Microsoft.

* Professor of Media, Culture and Communication and Computer Science at New York University.

practices and various contested conceptions of consent and/or privacy. So after the presentations, we're going to have plenty of time for Q&A.

I'm honored to introduce three incredible experts on privacy, each of whom brings a different perspective to the discussion. First, we will hear from Lorrie Cranor. Lorrie is an Associate Professor of Computer Science and Engineering and Public Policy at Carnegie Mellon University. She brings the computer science perspective for us to hear. Her research interests include usable privacy and security, technology and public policy. She's been incredibly active and you can look at her bio for more of the details. Ryan Harkins is an attorney working at Microsoft. The simplest way to introduce Ryan is to say that he is what all of the students in the audience want to be. So we'll hear more about that from Ryan. He's bringing the perspective of a practicing lawyer who has thought a lot about the privacy issues we're confronting. And then we have Helen Nissenbaum, the director of the Information Law Institute and professor of media culture and communications and computer science—lots of stuff—at NYU. She's going to contribute more of a philosophical perspective on these issues. Honestly, before I knew I was going to monitoring the panel, I've got her book and I've been reading along with this big data book [Professor Frischmann shows audience two books]. But if you haven't read this, you really must if you're interested in thinking about privacy in a careful—in a careful and sophisticated way. And so without saying anymore, let me pass the mic on to Laurie and we'll get started.

LORRIE CRANOR: I wanted to mention that besides being a computer science professor, I codirect a master's program at Carnegie Mellon in Privacy Engineering, where we are trying to train technologists in privacy. I'm going to be talking about privacy, notice and choice in practice today. And I think there has been a lot of discussion about the theory of privacy, notice, and choice, but not much empirical data on what actually happens in practice. So I hope to shed some light on that.

We all know that nobody actually wants to read privacy policies. I think that's pretty evident. Some of you may have seen a paper[1] that I wrote with Aleecia McDonald a few years ago where we look at the question of—well, hypothetically if people actually did read privacy policies, how much time would they spend reading them? And we came up with just this ridiculous amount of time, 244 hours per year if you had to read them all just for the websites you visit. So clearly, if the vision of privacy notice and choice that everyone being informed by

[1] Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 I/S: A JOURNAL OF LAW AND POLICY FOR THE INFORMATION SOCIETY 541 (2008).

privacy policies were actually to come about, it basically borders on ridiculousness.

So another form of consumer notice—not in the privacy space—is nutrition labels. This is something that has been mandated in the United States for some time now. And we've all become pretty comfortable with nutrition labels as way to be able to make comparisons between different items that we might want to purchase, and to get information that is useful to us. And so about a decade ago in the privacy community, people started saying, "if only we had nutrition labels for privacy." You could imagine privacy facts instead of nutrition facts. And so I started working with my students on trying to come up with what exactly would this nutrition label for privacy look like. So we looked at the nutrition label literature and asked what were the important elements. It needs to be standardized so people can learn how to use it. It needs to have standard terminology. We didn't all know what saturated fat was when they came out with nutrition labels, but we learned because it is the same terminology. It needs to be brief and it needs to be linked to an extended view where you can get more information.

So through an iterative design process over a couple of years, we developed some examples of privacy nutrition labels and we tested them in focus groups and in lab studies and in online studies. And we came up with something that looks kind of like this, which can be seen on the slide [shown in Figure 1]. We have a few different versions. But basically, it's a table where you have down one side different types of information that might be collected and across the top different uses of the information. Then it's color-coded so you can see whether this information and use combination is going to happen as a mandatory thing or whether it is opt-in or opt-out or whether it won't happen at all. And you can kind of see at a glance based on the color of the notice how much data collection activity is actually going on.[2]

[2] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, & Lorrie Faith Cranor, *Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach*, *in* CHI 2010, ACM Press 1573–92 (2010), *available at* https://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab09014.pdf.

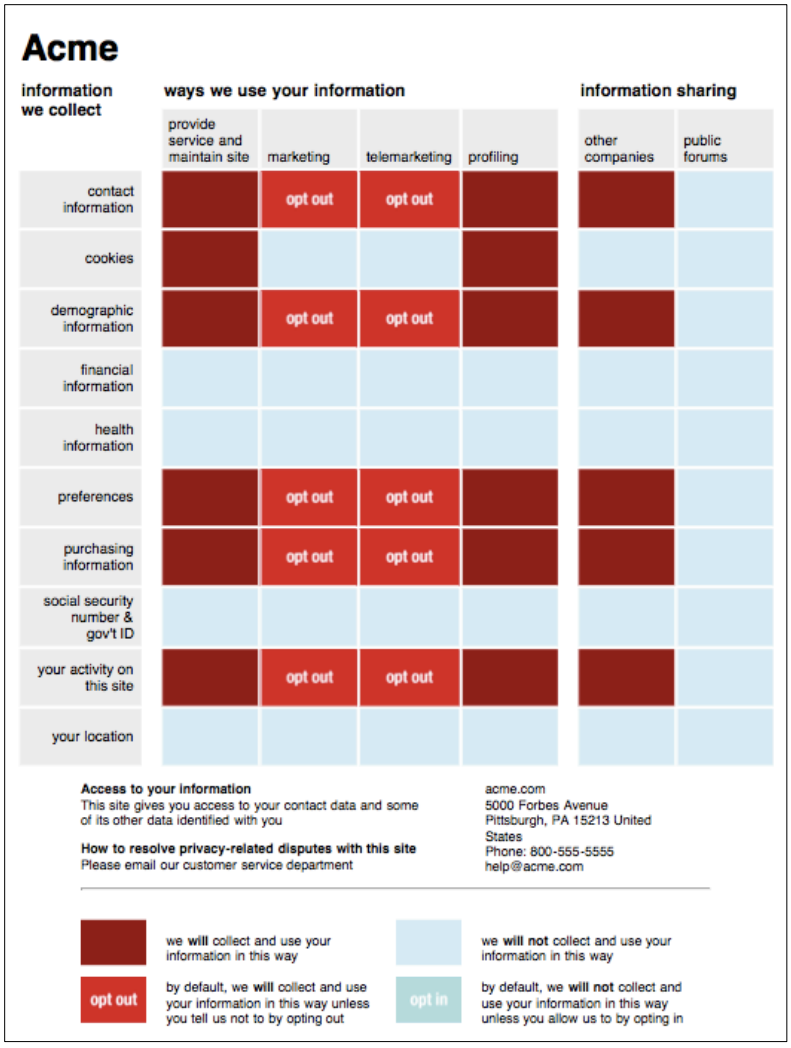

FIGURE 1: An Example of a Privacy Nutrition Label Developed at Carnegie Mellon University.[3]

Another approach that people have looked at is privacy icons. This is a set of icons that was developed by Mozilla. The best I can tell it is not actually being used at the moment. But it is actually really difficult to come up with privacy icons. It's an easy thing for people to throw out and say, "Oh, we just need some icons for this." But if you have ever sat down to do it, it's hard. And I've sat down with some really good designers and tried to figure out how can we represent privacy concepts

[3] *Id.*

and icons. Basically, what we've concluded is that you can come up with icons but they're not going to be intuitive. People are going to have to learn them and that might be okay. If we use them consistently over time, people would learn them as long as there are not too many of them. Take a look at these icons on the slide. If you didn't have the words underneath, you probably wouldn't understand them. This one looks like road with a policeman on one side of the road and a ticket on the other. It's really hard to tell exactly what these mean.

Android phones, smartphones, iPhones—this is another area where people have started looking at privacy notices. The Department of Commerce had a multi-stakeholder process all of last year looking at transparency for smartphones. What kind of notice can we provide about app data collection? In my view, that process was rather unfortunate the way it proceeded, where they were trying to develop something that was ultimately a consumer-facing notice and they conducted the entire process with absolutely no input from consumers and no user testing. At the end of the process some of my students actually conducted their own user test and we found that the end result was fairly confusing—not only for random consumers but even for the experts that we had actually go through the test.[4] So that does not seem like the best way to proceed.

We've done a variety of work on what might be a useful sort of notice. So one of my students came up with this idea of a privacy facts notice on your smartphone. In order to do these tests to find out whether these are useful or not, you need to actually put people in the context of trying to use them. So we came up with a study where we told people, "Someone you know, your friend, has just gotten a new smartphone and would like your assistance in selecting some apps. Here are the kind of apps they want." And we showed them different apps with different privacy information and we looked to see in what cases were they influenced by the privacy information and in what cases did they not even notice the privacy information. And so that's one way that these things can be tested.[5]

Another approach to notice and choice is to say, "Well, nobody is going to read privacy policies, so let's provide them in a format that the computer can read." Then your web browser, your smartphone or whatnot can read these notices automatically and do something useful

[4] Rebecca Balebako, Richard Shay, & Lorrie Faith Cranor, *Is Your Inseam a Biometric? Evaluating the Understandability of Mobile Privacy Notice Categories*, CMU CyLab (Technical Report: CMU-CyLab-13-011, 2013), http://www.cylab.cmu.edu/research/techreports/2013/tr_cylab13011.html.

[5] Patrick Gage Kelley, Lorrie Faith Cranor, & Norman Sadeh, *Privacy as Part of the App Decision-Making Process*, *in* CHI 2013, ACM Press 3393–3402 (2013), http://patrickgagekelley.com/papers/android-decision.pdf.

for you. So Platform for Privacy Preferences ("P3P") was the standard that the World Wide Web Consortium ("W3C") adopted in 2003 that was supposed to do this. I was heavily involved in that. A big problem was adoption and incentives for adoption. Microsoft actually built this into the Internet Explorer web browser back in 2002 and used it to make cookie-blocking decisions. The default setting, which 99.99% of the world has because nobody even knows they could change it, is that if you visit a website that has third-party cookies that do not use this standard format of P3P, those cookies will be blocked. So that was incentive for companies to adopt P3P so that they wouldn't have their third-party cookies blocked. The problem is that Internet Explorer didn't actually check to make sure that these P3P policies were in the right format. So it turns out that the format looks like a bunch of these kind of three and four-letter symbols and there's a dictionary of which ones are allowed. What the Internet Explorer system did is it looked for bad symbols—symbols that indicate a bad privacy policy according to Microsoft. It turns out that if you put a symbol that Microsoft doesn't know about—one you made up—it can't be bad. And so it lets you through. It doesn't block your cookies. So Amazon just used AMZN. All right? That's not on the list of symbols. There's nothing bad there. It also is completely meaningless to the software, but their cookies could get through.[6]

Facebook actually made their policy say, "Facebook does not have a P3P policy, learn why here." Those are all symbols that actually don't make sense to the web browser. So therefore, they are not bad symbols and their policy therefore was considered good and could get through. So this is basically routing around the system. It's circumventing the system designed to help people protect their privacy. There was a whole kerfuffle about it. Google and Microsoft issuing press releases and yelling at each other, but basically nothing came of it. At this point, Microsoft still has this built into their browser. Google and many other companies are still circumventing it and consumers' privacy is not being protected.

We have "Do Not Track," which you've probably all read a lot about. That has been going on and on and on, the process of trying to define what exactly is Do Not Track. There are also lots of tools that the ad industry and others have come up with to try to help consumers filter cookies and tracking—trying to prevent tracking in their browsers. We did a study to find out what people actually understood about the

[6] Pedro Giovanni Leon, Lorrie Faith Cranor, Aleecia M. McDonald, & Robert McGuire, *Token Attempt: The Misrepresentation of Website Privacy Policies through the Misuse of P3P Compact Policy Tokens*, CMU CyLab (Technical Report: CMU-CyLab-10-014, 2010), http://repository.cmu.edu/cgi/viewcontent.cgi?article=1072&context=cylab.

tracking that happens in your web browsers. We found that people didn't understand it. They didn't realize how it works. And they didn't recognize this little advisors choices icon—that little blue triangle thingy—which is supposed to be informing them. We also found that when people wanted to make decisions about trackers, they did it based on what they knew about the brands. So people told us things, like, "Oh yeah, Google can track me. They're a good company. Microsoft advertising—they're pretty good too but I already bought a computer. I don't need another one. So I don't need any advertising from Microsoft. And AOL—oh, yeah—I used to have their Internet service. It was kind of crappy so, nah, I don't need their advertising." So they're making decisions that are not really based on an informed sense of what is actually going on here with this tracking.[7]

We also looked at some of the tools to see whether people could actually use them. I'm not going to go into detail here. But just to point out even some things that seem really basic. This industry association says all you have to do is come to this website and opt out of our tracking. We showed this to users and asked, "Okay, where do you click to opt out?" And they couldn't figure it out. It turns out you're supposed to click on that checkmark. This wasn't obvious to anybody. Some websites—if you wanted to opt out, you got information in languages that were not English. We had people who then tried to go use Google translate to figure it out. Some tools gave you all sorts of jargon. Frames, embedded and object tags, prevent redirections. Right? This makes no sense to most people.[8]

So we decided to actually dig a little bit deeper into this icon problem.[9] So here's that blue icon much bigger and you can see it in the corners of the ads. This is how it appears on the Internet. We did a study with over a thousand people to see whether—as the industry claims—this is actually revolutionized consumer education and choice. We were a little bit skeptical.

So we had 1,500 participants and they went to our survey online. We showed them the screenshot of the *New York Times* and this little icon appears several times on this page. We asked them if they noticed

[7] Blase Ur, Pedro G. Leon, Lorrie Faith Cranor, Richard Shay & Yang Wang, *Smart, Useful, Scary, Creepy: Perceptions of Behavioral Advertising*, *in* SOUPS '12, ACM Press 1–15 (2012), https://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab12007.pdf.

[8] Pedro G. Leon, Blase Ur, Rebecca Balebako, Lorrie Faith Cranor, Richard Shay, & Yang Wang, *Why Johnny Can't Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising*, *in* CHI '12, ACM Press 589–98, (2012), http://www.blaseur.com/papers/CHI2012-opt-out-usability-final.pdf.

[9] Pedro G. Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, & Guzi Xu, *What Do Online Behavioral Advertising Disclosures Communicate to Users?*, *in* WPES '12, ACM Press 19–30 (2012), http://www.blaseur.com/papers/wpes2012-obaicons.pdf.

anything about privacy and the vast majority of them did not notice it. Then we showed it to them again and this time we pointed it out, and we had a few different versions. We had the triangle. We also had this asterisk man version, which had been considered and rejected. And we had different wording. So generally the industry presents it with the word, "AdChoices." But there are a bunch of other terms that they can use and we came up with a few others. And so people were in different conditions and they saw different combinations of these things. And then we asked them questions. One thing we asked them was to what extent, if any, does this combination of the symbol and phrase placed on the top right corner above the ad suggest that this ad has been tailored based on websites you have visited in the past. So they should all say, yes, this is what this means. But, in fact, for the people who saw AdChoices, only 58 percent of them understood that that's what it meant. On the other hand, we saw that when we asked them whether clicking on this would take them to a place where you can tell the company you don't want to receive tailored ads, which is also true, only 27 percent realized that, but 45 percent thought they could go to a page where they could buy ads, which is not true. And 56 percent thought more ads would pop up, which is also not true. So this is not communicating at all.

So summarizing what we've learned here about notice and choice through these studies. Privacy policies: how effective are they? Well, nobody reads them, so not very effective. Privacy nutrition labels: interesting research but nobody is really using them. Privacy facts for Android: again interesting research. Nobody is using it. P3P: well, that's being widely used, but it's also being widely circumvented so not very useful. Do Not Track: we still don't know what it means. Tools to opting out of tracking: they are out there, but they're actually pretty difficult for people to use correctly. The ad choices icon: nobody knows what it means and they're afraid to click on it because they think they're going to get more advertisements. So this is really not looking very good for notice and choice.[10]

So the one glimmer of hope that I want to leave you with is financial privacy notices. So you've probably seen these sorts of notices [as shown in Figure 2][11] recently in your credit card bills and your bank statements. The Gramm-Leach-Bliley Act of 1999 ("GLBA") mandated

[10] Lorrie Faith Cranor, *Necessary But Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice*, 10 J. ON TELECOMM. & HIGH TECH. L. 273–308 (Summer 2012), *available at* http://www.jthtl.org/content/articles/V10I2/JTHTLv10i2_Cranor.PDF.

[11] Lorrie Faith Cranor, Kelly Idouchi, Pedro Giovanni Leon, Manya Sleeper, & Blase Ur, *Are They Actually Any Different? Comparing Thousands of Financial Institutions' Privacy*, *in* WEIS 2013 Appendix I (June 11–12, 2013), http://weis2013.econinfosec.org/papers/CranorWEIS2013.pdf.

that we have financial privacy notices. But initially they were really long and full of lots of legalese. And then in 2009, eight federal agencies came up with a model privacy notice that standardizes them

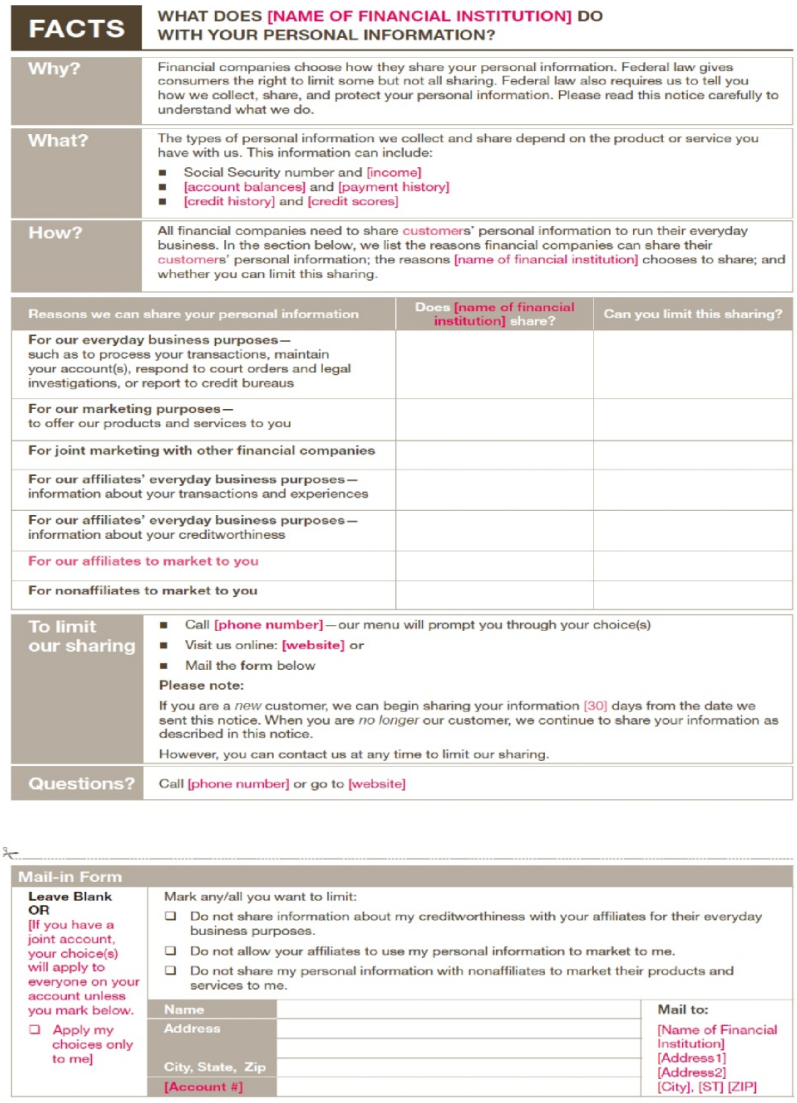

FACTS — WHAT DOES [NAME OF FINANCIAL INSTITUTION] DO WITH YOUR PERSONAL INFORMATION?

Why? Financial companies choose how they share your personal information. Federal law gives consumers the right to limit some but not all sharing. Federal law also requires us to tell you how we collect, share, and protect your personal information. Please read this notice carefully to understand what we do.

What? The types of personal information we collect and share depend on the product or service you have with us. This information can include:
- Social Security number and [income]
- [account balances] and [payment history]
- [credit history] and [credit scores]

How? All financial companies need to share customers' personal information to run their everyday business. In the section below, we list the reasons financial companies can share their customers' personal information; the reasons [name of financial institution] chooses to share; and whether you can limit this sharing.

Reasons we can share your personal information	Does [name of financial institution] share?	Can you limit this sharing?
For our everyday business purposes— such as to process your transactions, maintain your account(s), respond to court orders and legal investigations, or report to credit bureaus		
For our marketing purposes— to offer our products and services to you		
For joint marketing with other financial companies		
For our affiliates' everyday business purposes— information about your transactions and experiences		
For our affiliates' everyday business purposes— information about your creditworthiness		
For our affiliates to market to you		
For nonaffiliates to market to you		

To limit our sharing
- Call **[phone number]**—our menu will prompt you through your choice(s)
- Visit us online: **[website]** or
- Mail the form below

Please note:

If you are a *new* customer, we can begin sharing your information [30] days from the date we sent this notice. When you are *no longer* our customer, we continue to share your information as described in this notice.

However, you can contact us at any time to limit our sharing.

Questions? Call [phone number] or go to [website]

Mail-in Form

Leave Blank OR [If you have a joint account, your choice(s) will apply to everyone on your account unless you mark below.]
❑ Apply my choices only to me]

Mark any/all you want to limit:
❑ Do not share information about my creditworthiness with your affiliates for their everyday business purposes.
❑ Do not allow your affiliates to use my personal information to market to me.
❑ Do not share my personal information with nonaffiliates to market their products and services to me.

Name
Address
City, State, Zip
[Account #]

Mail to:
[Name of Financial Institution]
[Address1]
[Address2]
[City], [ST] [ZIP]

FIGURE 2: The Gramm-Leach-Bliley Act Model Financial Privacy Form 2009.[12]

into the sort of tabular format. And this is the standard format and the parts you see in pink are kind of the fill-in-the-blank that each company would fill in for themselves. So this is not a computer readable format.

[12] *Id.*

But it is online in PDF or HTML.

So my smart students wrote scripts that would go and collect all of these, and parse the PDF files or the HTML—extract all the data and put it into a database. We now have thousands of these in a database.[13] We're working on putting thousands more. We had to "FOIA," ("Freedom of Information Act") the Federal Reserve to get a complete list, but we now have it. This allows us to actually see, if all these privacy polices are the same or if there are differences for different banks. What we found here—just a little taste of it—is that if you look at all banks overall, there are a lot of banks that don't share your personal information—all this green stuff on the left. But if you look only at the 100 largest banks—the banks that we've all heard of—there's a lot less green. They're actually sharing a lot more of our personal information. And so what we see is that consumers do have choices in banks according to privacy, but you may not be able to find the choices because you don't know about all the little banks. You know about the big banks. Basically what we found is that there is some promise here that we could actually have some meaningful notices and some opportunities for choice, but we're missing the ability for consumers to actually find things. So one of our current projects is actually building a bank search engine that will hopefully allow people to actually search for banks based on what kind of financial products and services they want and privacy simultaneously.

RYAN HARKINS: Good morning. So as Brett mentioned, I am Ryan Harkins. I'm a privacy lawyer at Microsoft. I appreciated Brett's introduction of me. I think if my wife were here, she might tell you that in some ways I might not be who you want to be. But I appreciate the remarks nonetheless. I also have to say just how incredibly excited I was when I learned I was going to appear on a panel with Helen and Lorrie. Most of you probably know this, but Helen and Lorrie are two of the best and brightest minds we've had working in privacy for some time. And they're both incredibly accomplished. They're both incredibly impressive and I'm not entirely sure how this obscure privacy lawyer from Seattle lucked into appearing on a panel with the two of them, but I'm incredibly happy to be here. Before I begin, I'd like to briefly thank Cardozo and all the students and everyone involved in putting this together today. I know how much hard work goes into putting on events like this and I very much appreciate it.

I'd like to spend some time this morning talking about notice, consent and "big data." What does big data mean for notice and

[13] *Id.*

consent? What are some of the unique challenges that big data presents for notice and consent? And what are some possible ideas that we at Microsoft have started thinking about in order to address some of those unique challenges?

So first, I think it makes sense to begin with talking about what is big data. And it turns out there is no single uniform definition of what big data is. Generally speaking, it refers to the burgeoning ability to store, aggregate, and process massive volumes of data from diverse data sets and to do so at incredibly rapid speeds. By doing that, it allows us to glean surprising insights or new correlations from those data sets—things that may not necessarily have been contemplated at the time data was collected. Big data is principally the product of a couple of developments. First, you have a number of technological developments. So things like the rise of ubiquitous computing and online services. Things like cloud computing, social media—you have the development of the Internet of things. And a new generation of sensors that will increasingly collect more and more data about our offline activities. And then of course, you have the rise of incredibly cheap data storage. And all of that adds up to mean that there has been an explosion in the volume and in the variety of data that's being collected about each and every one of us. Second, you have the rise of much more powerful and incredibly sophisticated techniques to aggregate, to analyze, and to mine data.

And so we've seen plenty of anecdotal evidence over the past few years about this data deluge. In 2011, a number of outlets reported that 90 percent of all data in human history had been created in just the two preceding years. Facebook last year reported that they hold over 250 billion photos with over 350 million more being added each and every day. In 2010, The Economist reported that scientists had used big data techniques to reduce the time it took to decode the entire data genome from ten years to only a week. I'm sure it's even faster today. So a number of people have raised concerns about what big data means and what it means for privacy law in particular. Big data presents unique challenges for privacy law largely because for the past four decades, privacy law has been governed by something called the Fair Information Practice Principals ("FIPPs"). These are basically a set of principles that are designed to ensure that when data is being processed about individuals, it's done fairly and with some protection for privacy. At the heart of the FIPPs has always been the notion of informed consent—the idea that individuals have the right to receive notice when data about them is being processed and that they have the right to make a meaningful choice about the data being processed. But people have recognized that the consent edifice has been cracking for some time.

That's in part because it obviously places much of the burden of privacy protection on individuals.

So while technology has continued to evolve in ways in which more and more data about each and every one of us is being collected, it still remains the case, by and large, that individuals are expected to read privacy notices. Individuals are expected to understand privacy notices and individuals are expected to make informed decisions about complex data processing activities based on privacy notices. And that has become exceedingly difficult for individuals to do. You probably all notice this as you browse the Internet that we're all faced with an overwhelming flood of privacy notices from virtually every retailer, every service provider, and every other entity we interact with online. Lorrie mentioned research that she and Aleecia McDonald conducted in 2008. That is astounding; 244 hours a year to read every privacy statement you encounter on the Internet. That's an indication that something isn't working here.

But this whole notice and consent edifice—it doesn't just put pressure on individuals. It also puts companies in an increasingly difficult position. That's because—I can say as someone who is on the front lines rolling up his sleeves and working with engineers to try to draft privacy notices—you wind up trying to, in an effort to comply with the law, craft notices to cover virtually every piece of personal data you might collect and virtually every use you might make of that data. That can be really, really hard to do. On the one hand, you want to present choices and notices at a time and in a context in which they'll be relevant and in which they'll resonate with users. But on the other hand, you don't want to overwhelm them, you don't want to annoy them, you don't want to lose their interest. On the one hand, you want to be clear and comprehensive, but on the other hand you want to be comprehensible and concise. So sometimes, you know, as you're working with your clients and trying to do these things, it feels as though you're being asked to reconcile irreconcilable goals. So while the notice and consent edifice may have been cracking before, I think it's fair to say that big data threatens to obliterate it altogether because big data will mean that there will be even more data being collected. It'll be overwhelming and it will make it extremely hard for individuals to provide effective consent or make informed decisions about all of the data that's being collected about them and about all of the prospective uses of data.

The rise of the Internet of things, which Brett alluded to, is one example of how this challenge will be compounded. With the Internet of things you'll have an increasing number of devices that will be connected to online services and collecting data, but may not have

screens or user interfaces ("UIs") or other opportunities to provide people with effective notice and consent experiences. So all this is very challenging and it's concerning in a way as well because privacy law's heavy emphasis on notice and consent at the time of data collection may preclude certain uses of data—uses that could help unlock incredible value. There are plenty of examples of the promise of big data: it could help us predict drought and famine and plan for food shortages; it could help urban planners plan our communities in better ways to reduce congestion; it might help us cure diseases.

One example I wanted to share with you is that Microsoft researchers were recently able to apply analytical methods they had developed to combat spam to help doctors understand the way HIV mutates. If you talk to medical experts focused on HIV, they'll tell you that HIV can be extremely challenging to address because it's constantly mutating and it's doing so in order to avoid the ways in which our immune systems might attack it. Well, as it turns out, email spammers program their spam to constantly mutate so that they can avoid anti-spam filters. So by applying some of the same analytical methods used to understand the way spam mutates, doctors have been able to get a better understanding about the way HIV mutates and perhaps understand ways in which our immune systems may be able to address it.

So what to do? How can we unlock the potential of big data while providing protection for privacy and while addressing some of the other concerns that scholars have raised with big data? For example, some have raised concerns that it could result in much more powerful and invasive profiling of individuals. It could lead to surreptitious discrimination—discrimination in areas like applications for credit, application for insurance, applications for employment. And that's true because even though we may have anti-discrimination laws on the books today, big data derives its conclusion by using proprietary algorithms, which among other things will make it extremely difficult for individuals to know whether they've been discriminated against. So it's still early days, I think, in addressing some of these unique challenges.

We at Microsoft have started to put forwards some possible solutions that we think are worth exploring in further detail. Number one—and perhaps most importantly—we do not think that notice and consent should be abandoned. Sometimes you hear this from some people. In contrast, we think notice and consent should be strengthened and adapted to the world of big data, adapted in ways that will make it more effective, adapted in ways so that it will actually apply in scenarios where decisions can be informed and can be meaningful.

And so to determine what scenarios in which informed consent might be meaningful, we could ask ourselves three questions. Number one, would a particular data use be consistent with user expectations? Number two, would a particular data use create a significant risk of harm either to individuals or to groups? And number three, what would the societal benefits be from a particular use scenario? And so if you ask yourselves those three questions, you could conceivably bucket uses into three categories.

On one end of the spectrum, you might have uses that are widely expected or understood by users, that present high societal benefit and low risk of harm to individuals. So, for example, if someone goes to an Internet retailer and purchases a product and enters his or her mailing address, I think it's fair to say that everyone understands that the retailer will use that address in order to ship the user that product. So including information like that in a privacy statement is unnecessary and it actually serves to obscure other disclosures that wind up being more important. So ironically, including that sort of information may actually decrease transparency.

At the other end of the spectrum, you might have uses that present a great risk of harm to individuals or are not expected or present low societal benefits. So you might say that uses in that bucket should flatly be prohibited. Using big data to illegally discriminate against vulnerable communities would be one example of that. And then, of course, you have this great big bucket in the middle. And that's where we think focusing on notice and consent could be most effective and could be most helpful in terms of protecting users' privacy. So in addition to focusing on notice and consent, we also think it would be worthwhile bolstering other principles in the FIPPs. For example, you'd continue to include certain security requirements to make sure that data is protected effectively. You might tweak data minimization so that it requires certain de-identification techniques and might prohibit re-identification under certain scenarios. And you might include things like transparency and integrity requirements. Integrity requirements that would apply not just to data itself, but perhaps also to the processes by which decisions are made about individuals.

We also think it would make sense to focus on some other things. You might propose technological solutions to help address privacy concerns with big data. One example might be something called differential privacy, which is something that Cynthia Dwork at Microsoft has done a lot of research on. You might include standards and other things. So I guess I'd just say, in wrapping up, big data obviously presents big challenges. It presents big and unique challenges to privacy law. It's early days in addressing those, but we look forward

to a vigorous public discussion about how we can unlock big data's potential and still protect users' privacy.

HELEN NISSENBAUM: Thank you very much for including me in this event. Thanks to Pamela Grutman for keeping us all well informed and, to the Editorial Board and the Staff, and thanks, of course, to Brett and my two colleagues on this panel. We do have points of disagreement, but it's really interesting to see the ideas weave around each other. Zappos' privacy policy is an example of privacy notices that we are all concerned about.[14] As I told Brett, because I haven't given this talk before and didn't know how long it was going to take me, I thought I would start my presentation slides by giving the conclusion to the presentation first so you would know where I'm going, in case I do not reach the end in the time allotted.

The question, "What is the problem for which Notice and Consent is a solution," is inspired by Neil Postman, who was the founder of the NYU department of media and culture and communication (of which I am a faculty member). The question he recommended asking about new technologies is: what is the question to which this (new technology) is the answer? So I find myself asking (often, when I'm getting depressed about the staying power of notice and consent), "What is the problem for which notice and consent is the solution?" And my conclusion is that it might be a solution to some problem, but it doesn't seem to be a solution to the privacy problem. Perhaps if I weren't a philosopher talking, here, in a law school, I might have not concluded in quite such a dramatic way, that "Notice and Consent is a Sham," but let me see if I can persuade you of this conclusion at least to some degree. My fellow panelists have already mentioned several important concerns with notice and choice; their respective work has made ongoing contributions to research and practice in this area.

Before proceeding, a word on terminology: the organizers of this event have used the terms transparency and choice. I have typically used the terms, notice and consent. Although some might see difference in the two phrases, I intend none but find that notice and consent are more easily integrated into discussions surrounding "privacy notices," and the consent paradigm in other fields.

Returning to answer the question of what problem notice and consent solves, the common premise is that notice and consent solves privacy problems. What makes this premise believable is a presumed definition of privacy as control over information about oneself. With

[14] *See* Privacy Policy, ZAPPOS.COM, *available at* http://www.zappos.com/privacy-policy (last visited May 13, 2014).

this in the background notice, too, makes sense as a protocol for facilitating control with particular conventions for expressing notice drawn from Fair Information Practice Principles that Ryan had mentioned.

As an aside, I would like to endorse a point that Brett made in his introductory remarks, namely that what standards guide proper consent is sadly underexplored. Further to that, there is a gaping absence of discussion of what counts as true consent versus coercion and where actually to draw the line because, in my view, much of that is presented as consent cannot be counted as genuine consent because the stakes for individuals not consenting are way too high. Although I'm not going to develop this particular argument here, I would love to engage in further conversation about it and would like to see more work on it.

FIPPs-guided notice and consent, as you heard Ryan mention, is at the heart of so much U.S. privacy regulation whether legislation that was directly about privacy or required accompanying privacy rules to be written. Examples include the GLBA Model Financial Privacy Form 2009, GLBA Financial Privacy Rules ~2000, the Health Insurance Portability and Accountability Act Privacy Rules ~2000, the Family Educational Rights and Privacy Act 1974, and the Video Privacy Protection Act 1988. The common message in all these contexts is: "Here are some rules that define privacy expectations. But if you provide adequate notice, you can ask users to forgo these expectations." In other words, even though there are laws or rules that impose substantive constraints, there always is the loophole, the ability to provide notice and get consent for something other than the substantive constraints expressed in the law.

As we know, the model of notice and consent has migrated online, embodied in the ubiquitous privacy policy. Fellow panelist, Lorrie Cranor's amazing user studies have shown repeatedly how problematic these are, particularly notice. But, online, as off, consent is also problematic as users are not exactly presented with a choice. Further, a fact I have hoped my friends in the field of law might answer is why consent can be presumed even if one simply lands on or navigates to a website, and goes no further. Information about your visit can be collected under the privacy policy even though you have not even tacitly accepted the policy by going further with your visit to the site. There's already an exchange of information to which you haven't consented.

Notice, however, has remained the most discussed element of notice and consent, as I have elaborated in my Daedalus article and will briefly review with you here, today. In my view, the status quo in 2014 continues to be that the protection of privacy is tossed over the fence to

us, individual data subject. The logic begins with privacy as control over information about ourselves; it continues with an incomprehensible policy and the handoff: "Well, here you go! This is how you're going to protect your privacy: take it or leave it." The status quo clearly hasn't worked (for us) because here we are today worrying about the amount of information being collected, the way it's used, and all the staggering array of online tracking means and mechanism, a situation that many—including the White House and the various consumer agencies consider to be highly problematic. So whatever the policy has been, it isn't working. Ample evidence in excellent work by Joe Turow, Chris Hoofnagle, Kirsten Martin, Lorrie Cranor, respectively, and collaborators—some of which we heard today—shows that people simply do not understand privacy policies.

There are some people—I call them "critical adherents"—who would argue that the problem isn't with notice and consent *per se*, but the fact that it's poorly implemented. Here's a short excerpt of the privacy policy from the website *Obama For America*, the campaign website. They say, "Oh—you know what—it should have been better than this!" I've been looking at Massive Open Online Courses ("MOOCs") and am a little horrified that their privacy polices look very much like those you might see in a commercial website and in particular—I don't know if you can see the bottom line—"be sure to return to this page periodically to ensure familiarity with the most current version of the privacy policy."[15] A colleague of mine at the New York University School of Law, Florencia Marotta-Wurgler, conducted an empirical study between 2009 and 2012 of 276 websites studying their privacy policies focusing on 67 terms per policy. She found that 59 percent of sites had altered their policies during the period of study, and among these, on average, they changed 20 times during the period of study. I have not yet seen a policy that does not end with the caveat something to the effect, "Be sure to return this page periodically to ensure familiarity with the most recent version of the privacy policy." If we were to takes this exhortation seriously, an analysis published by Lorrie and collaborators reveals what the cost of this would be in terms of time, and dollar amount opportunity cost. From a societal point of view, we would not want individuals to do what these privacy policies say.

I don't, however, share the critical adherents' hopefulness. Yes, I agree we can improve and do need to improve privacy policies, but as long as these policies are pegged as the full and final defense of privacy,

[15] *Privacy Policy, Coursera,* COURSERA.ORG, *available at* https://www.coursera.org/about/privacy (last visited May 14, 2014).

no matter how improved, they will fail in their mission. One reason for this, I have discussed in an article on a contextual approach to privacy online under the label, "transparency paradox" meaning that if you simplify a privacy policy so that a non-expert user can grasp its meaning (even expert users, for that matter), you no longer can present crucial information that individuals need to know adequately to grasp the implications for privacy. The transparency paradox does not, of course, apply in all situations. Take the simple privacy policy you find on the Alcoholics Anonymous website—a model policy by my count. The policy statement works because it maps well onto the policy itself, which basically states, "although we might collect information, we don't share it with anyone, period." This is a simple privacy policy, and its simplified presentation is perfectly right. But when you're Facebook or Google, or one the countless websites that allow ads from networks that track users and deliver behavioral advertising, then offering simplified privacy policies for immensely complex, and possibly exploitative practice these simply, plain language policy statements are downright disingenuous, certainly not able to communicate what information is collected or what is being done with it.

I was very pleased to learn that Lorrie and her group are studying financial notices because when, recently, a notice from American Express came in the mail recently, explaining what it does with "personal information" I was impressed with the way the form had been simplified and tabulated, with columns clearly statin: Why are we collecting? What are we collecting? And how—and with who we are we sharing this information? That was my first impression. Closer scrutiny reveals a less happy situation: first of all, one learns that while in some areas customers do have the capacity to limit sharing, there are many others where we actually cannot. In other words, the scope of consent is quite limited. Further, my cynicism rises when I learn that American Express is offering this format because regulation requires this. Second, if you want to know not only what information American Express is sharing and with whom they're sharing it, but what information they're sharing with whom—something that really matters to me—you're out of luck. An instance of the transparency paradox—the mandated, simplified template becomes a liability as complexity of what is crucial to know ratchets up.

Our critical adherents might say, "Oh, people are just too un-analytic, too simple-minded to grasp this level of complexity and so we have to dumb it down." A study by Kirsten Martin of online commercial interactions suggests otherwise.[16] Though not yet published, the study

[16] Kirsten Martin, *An empirical study of factors driving privacy expectations online*, paper

was presented at the 2013 Privacy Law Scholars Conference, but even the preliminary findings are immensely important (and, based on Ryan's comments, comports well with what I understand to be about Microsoft's commitments). Martin's study involves many subjects, many conditions, and a complicated factorial analysis, to which I'm unable to do justice here. Because, here, I can offer only a brief and pointed glimpse of one thread of the analysis, I strongly recommend reading the article in order to appreciate the full context of the work and the diverse range of results. The finding thread I wish to highlight aims to understand the way people calibrate stated privacy policies with their own expectations or privacy practices. In the first step of this thread of Martin's study, subjects are provided privacy policies of fictional banks, search engines, retail merchants, etc. In step two, they're shown vignettes such as the one below:

> You are working on an online banking website that you have used infrequently for about a week. On the online banking site, where you clicked and looked on the page is collected by the website and will be stored for a month. The data collected also includes a unique identifier for your computer. The website then uses the information for future ads targeting your friends and contacts.

From this point, the study follows two forks: in Treatment A, subjects are asked whether the practice described in the vignettes meets their privacy expectation. In Treatment B, they are asked whether it complies with the respective privacy policy they had been asked to read. What Martin found was that when subjects in Treatment A respond, "No, this does not comport with my expectations," subjects in Treatment B are saying, "No, it does not comply with the privacy policy." A surprising result, one might think, when, in fact, every vignette does comply with the privacy policy.

What this study shows is that that people don't arrive at websites as *tabula rasa*; instead they arrive with substantive expectations about appropriate and inappropriate information practices. Lorrie, Martin's findings are fascinatingly compatible with your subjects' brand expectations. Their expectations are not shaped by what they read in privacy policy statements—at least, in the format that has become commonplace—but by norms, or standards of good behavior that they bring into an interaction. Expecting individuals to read carefully and take in the implications of complex privacy policies, provided as unilateral contracts, is unrealistic, disrespectful of our time, and even exploitative. Moreover—and this is important—it places organizations

presented at Privacy Law Scholars Conference 2013 (May 23, 2013), *abstract available at* http://privacylaw.berkeleylawblogs.org/2013/05/28/kirsten-martin-an-empirical-study-of-factors-driving-privacy-expectations-online/.

in the untenable situation of contriving, to the last detail the relationship they're going to have with individuals expressed as "gotcha-proof" privacy policies. The system is not workable, or efficient, except for the few businesses able to afford armies of lawyers both to write and defend themselves against opportunistic privacy policy trolls.

Some of you will recognize these signs from the fences of Washington Square Park.

FIGURE 3: Signs in Washington Square Park.

The connection I'd like to draw with the statements of notice, and consent is this: When any of us go to Washington Square Park and you see "No bicycling," "No skateboarding," "Do not feed the pigeons," we don't conclude, "Oh. Great! Look how many things I can do. I can throw litter on the ground. I can trip passersby. I can pull plants out of the beds." I know you're lawyers, so we should set aside questions of legality for a moment; still, there are many things most of us would not do because we have a common sense of what it means to be ethical, considerate, respectful users of public space. Sensibly, the signs do not specify all prohibitions but those applying to behaviors that might, or might not be acceptable in a park, behaviors about which visitors might be uncertain. There are good reasons for such parsimony. We depend on laws, norms, and conventions to do most of the work so that specified prohibitions will be salient and effective. The regulatory regime of notice and consent flouts these longstanding practices and holds individuals responsible for internalizing tens, no scores of terms of service contracts whose contents are virtually unbounded by standards,

despite Martin's profound demonstration that they are presumed by most users entering online interactions.

At the beginning of my talk I set out to show that notice-and-consent (transparency-and-choice) simply cannot shoulder the full burden of protecting privacy. The ubiquitous privacy policy statement, though it may give regulators and lawyers cause for action when behaviors are not consistent with statement claims, is little more than empty ritual when it comes to privacy itself. While agreeing with critical adherents that the statements themselves are incomprehensible, my argument goes further, pointing to the hopelessness of an ecology containing countless privacy policies with which individuals must reckon each day.

The final challenge to notice and consent, however, beyond privacy policy statements, is revealed in the practices known as big data. It may manifest in the briefest, most insignificant encounters of daily life, for example, a store cashier asking, "Can you give us your zip code?" We oblige; we may even think, "Oh it's so nice, they're asking for my zip code; they must be interested in serving me better!" And, in the spirit of notice—we have been asked—and consent, we oblige. What we don't know is all the other information that is keyed to the zip code and what else we are sharing when we share it. Here is the question: did you consent to providing all that other information to the store? Although discussions of big data have come to the forefront only in the past few years, I anticipated this question in my 1998 article, "Protecting Privacy in an Information Age: The Problem of Privacy in Public."[17] In general, the practices of data aggregation and analysis mean that much of the information we explicitly share functions merely as a hook into these vast troves that may be inferred. Consent, it seems, is required only for the initial sharing. In an article with Solon Barocas, "Big Data's End Run Around Anonymity and Consent," (which will be published in June 2014), we argue that big data practices allow inferences to information with or without your consent because you can be placed in clusters of people who are similar to you. If a small, but representative sample of these people have consented to sharing information that is of interest to the parties with whom you interact, this makes you vulnerable to discovery with a high degree of probability.

This, finally, is why I have concluded that notice and consent is—putting it starkly—a sham. The research and regulatory community have wasted too much time on notice and too little time on exploring substantive limits on the practices of data collection, sharing and use.

[17] Helen Nissenbaum, *Protecting Privacy in an Information Age: The Problem of Privacy in Public*, 17 LAW AND PHILOSOPHY 559 (1998) *available at* http://www.nyu.edu/projects/nissenbaum/papers/privacy.pdf.

We have, irresponsibly, thrown the problem over the fence to individuals, ill equipped with relevant expertise and power. Although notice and consent does have a roll to play and work on improving both elements remains important, the urgent need is to articulate and impose substantive constraints, which are specific to social context. This is what I've tried do—what I've tried to suggest anyway—in the theory of contextual integrity. Thank you.

BRETT FRISCHMANN: All right. So we're going to audience Q&A in just a moment, but before we do that—two things. One, I want to give the speakers a chance to react to each other's talks. So if there are comments or questions that you want to raise, we'll start with Lorrie and go down the line. Then I may have a question or two.

LORRIE CRANOR: So I actually agree with most of what my fellow panelists said. Thinking about what Helen just said about notice and choice being a sham, I mostly agree with that. Even as I've spent time trying to improve notice and choice, I think that ultimately what problem can notice and choice solve—I think that I'm mostly in agreement that it doesn't directly solve the problem of increasing privacy. I think what it might be able to do if it were more useful and usable is to highlight where we don't have privacy, call attention to the problem, and therefore, maybe motivate people to come up with real solutions.

RYAN HARKINS: I think I also agree with what Helen and Lorrie are saying in terms of the incredible strain that's on the notice and consent regime now, and the increasing strain on that regime. I'm not sure I would go as far as Helen does to say that notice and consent is a sham because we do think it plays an important role. I wonder whether Helen would even acknowledge it, at minimum, provides individuals with a way to ensure that their data is being processed or transferred appropriately. But, you know, with that said we clearly need to rethink the way we're applying notice and consent particularly as we move forward and there are a lot of questions and a lot of ways we could do that. The devil is always in the details and Lorrie sort of alluded to this or pointed to this in her discussion of things like P3P and the way different proposals that could have been solutions wound up being implemented.

BRETT FRISCHMANN: Helen, since you went last and they were both responsive to you, I'll just start by asking you a question. Explain why it's a sham and why maybe consent versus coercion is important.

Suppose we had theoretically perfect information or transparency about data collection and use practices. Assume that. It's just a given. Assume Lorrie fixes all of the problems about conveying to consumers or users or citizens the kinds of data collection practices and uses that are occurring. What would be meaningful consent and even if people are consenting, when would you say that informed consent would lead to the wrong answers?

HELEN NISSENBAUM: The point of the last comment, about big data, is that with or without notice, your consent doesn't matter as to whether the information in question is going to be revealed. It is odd to hang onto a mechanism for protecting privacy, which in certain circumstances is irrelevant to whether others will have access to information in question. This is why we see this mechanism as a failure, if we are concerned about privacy and not merely the enactment of empty ritual. This point is different from the concern over identifying when consent is closer to coercion. This was impressed on me, in a small way, when reading the American Express notice and realizing how little choice I had, in fact, regarding so much of it's data practices. The situation is coercive because I cannot choose to come back to American Express and say "I don't like your practice; please make the following changes," it's to accept the policy in its entirety or not to use American Express. The trouble is, all other credit cards follow the same practices—so unless you want to do without a credit card, you really do not have a choice. It would be wonderful to see more work done on the many "choices" we encounter that are not really choices. This knowledge seems crucial for a regulatory regime that relies so much on choice.

LORRIE CRANOR: Right. And then you might say, "Well, are all banks like American Express? Are they all that bad? What if I switched to MasterCard or Visa? Are they as bad?" So what we've seen is that all of the banks that you've heard of and would probably think of to switch to, are all about equal. But what if you go to the agricultural bank of North Dakota—I just made that one up—banks you've never heard of. There are some that actually have much better practices, but how would you find them?

HELEN NISSENBAUM: And maybe you would be too scared to put your money there.

LORRIE CRANOR: Yeah.

RYAN HARKINS: I would just pick up on something Helen said about how problematic notice and consent is—especially in the world of big data. I referenced this a little bit in my remarks that placing such heavy—I mean, in some ways, it could be effectively impossible to get consent for all of the ways data will be used in the world of big data because that is, in a lot of respects, precisely the point. Data might be used in unexpected ways—ways that you didn't foresee at the time in which you collected it. I guess it just further emphasizes or underlines the point that we really need to rethink the way in which notice and consent is applied.

Frankly, there are other FIPPs that we need to rethink as well, which I didn't mention in my remarks. Just as one example, two FIPPs include data minimization and purpose specification. So the idea that you should minimize the amount of data you collect about someone to only that which you need to further a legitimate business purpose and you should hang onto it only for the minimum amount of time you need it. And purpose specification, of course, referring to the fact that you should disclose the purpose for which you'll use data when you collect it and limit the uses to only those purposes. Well, those are, to a large extent, antithetical to the whole concept of big data. Big data is all about collecting more and more data in the hopes that you might be able to use it in unexpected ways.

One other example I didn't talk about is Bing search queries. Microsoft Research has worked with Dr. Russ Altman at Stanford to analyze Bing search queries to help identify a potentially lethal diabetic reaction between two seemingly unrelated drugs, Paxil, an antidepressant, and Pravachol, a cholesterol drug. They wound up discovering by analyzing large volumes of de-identified Bing search queries that users who searched for both of those drugs were much more likely to also search for certain diabetic symptoms, much more likely than users who would search for only one of the drugs. So it's just another example of how data might be used in ways that we don't foresee or portend when it's collected.

BRETT FRISCHMANN: Great. One last question and then I'll throw it out the audience. The question is this: in the world of big data in particular—informed consent, even if it worked well, would be challenged and maybe it would be a sham. Is the concern about what big data will enable others to learn about you or you as an individual? Or is it about how big data will drive and automate systems that affect you, even if it isn't necessarily identifying you? In other words, automating financial systems, say, micro-payments because of things you use, or automating media or advertising systems that service you as

opposed to learning about you in particular. I just want to ask if that's something that you've considered; it seems to be implicit in some of what you guys were thinking about.

LORRIE CRANOR: I think that there are a lot of different types of privacy harms. There's not just one privacy invasion. There are lots of different ways that privacy can be invaded. So I think the concerns about big data include everything that you just said. They include both my specific data that is going to have a direct impact that you know this about me, as well as the fact that decisions will be made because of my data being in there or because the database exists, even if my data isn't involved, decisions will be made about me. And I think all of those are privacy concerns that people are worried about.

HELEN NISSENBAUM: This question connects to the next panel about uses of these (big data) techniques by the NSA or law enforcement agencies. We need to not have to accept that because big data has benefits we simply jump to perform cost-benefit analysis. This is a mistake many have made. It's possible for us to be more careful about specific uses and specific types of practices and demand that accountability for those specific uses not to mention, rationale for collection in the first place even before we get to the stages of aggregation, analytics, and so forth. It is a conversation that we need to have.

BRETT FRISCHMANN: Great. Let's open it up to the audience. Plenty of questions. Just raise your hand and say who you are and then ask your question.

AUDIENCE MEMBER: [inaudible question]

RYAN HARKINS: I think you're putting your finger precisely on the problem that we're highlighting, which is applying notice and consent in the world today and applying consent moving forward, given the way technology has evolved and will continue to evolve, is a real challenge. So I think it's absolutely true that we need to rethink the way we're applying informed consent and there are all sorts of ideas on how to do that. But I still think it's early days, at least in terms of trying to evaluate the way in which informed consent and notice will work as we move forward. But I think the question you're raising is a legitimate one.

HELEN NISSENBAUM: Over time, I've looked to other environments,

such as bio-medical research to learn from their models of informed consent. Parenthetically, there is a good dose of humility in these environments where researchers and practitioners constantly seek to improve it. One important difference, also discussed in the paper with Solon Barocas, is adherence to background expectations, often shaped by professional responsibilities, unlike in the general case of information flow where the full load is place on the informed consent juncture.

LORRIE CRANOR: And I think, I usually don't use the word informed because I think for the most part it's not informed. I think we have consent, but it's not informed consent. And I think getting true informed consent would be actually very difficult. It would be very time-consuming and I think that isn't necessarily what we want. I don't think we want to spend all our time being informed so that we can consent. I think we want things to just work the way we want them to work.

AUDIENCE MEMBER: [inaudible question from audience asking what the person should do about a credit card company that misused her personal information]

LORRIE CRANOR: I should probably give the advice because I'm not liable for anything [laughter].

RYAN HARKINS: Yeah, I'm surprised that they—obviously credit card companies use all sorts of data points for anti-fraud purposes, which I think we generally recognize as a good thing, but I'm surprised that they would actually call her friend and try to track her down in that way. I obviously don't know all the details.

AUDIENCE MEMBER: [inaudible question]

RYAN HARKINS: It's a great question and I think we would say that those sorts of principles—data minimization and purpose specification are implemented today and we would point to the privacy disclosures that we have and the privacy notices. But of course, then you get into this whole debate about, you know, whether privacy notices today, across the industry, are actually effective? Are they actually accomplishing the purposes for what they set out to accomplish? And I think that's a part of what we're talking about today and I certainly think it's very fair to raise that point and to recognize that we think transparency and control are important, but we clearly need to rethink the way in which this is working in practice.

BRETT FRISCHMANN: We've got about three questions in about four minutes. I'm going to take all three questions and then we'll have the responses to the three questions.

AUDIENCE MEMBER: [inaudible question]

BRETT FRISCHMANN: That's a question. That's one.

AUDIENCE MEMBER: [inaudible question]

BRETT FRISCHMANN: Great. And then a third.

AUDIENCE MEMBER: [inaudible question]

BRETT FRISCHMANN: Three provocative questions. Panelists? We'll start with Helen and come back this way then. Great.

HELEN NISSENBAUM: Right now, the burden of consent is centered on the individual data subject. When you click through terms of use, you are held accountable for them, however long and abstruse. To answer Ira's question, I do think that there is a limited role for notice and choice, but the ideal reforms that I would like to see is that there would be substantive rules—let's call them contextual rules, similar in form, for example, to those we have for, video privacy—that say "this is what you can and cannot collect. This is how you can share it. This is how you can use it," and so on. At present, there is the inevitable: you can do something different as long as the subject consents.

I would like to strike those loopholes off. To override these substantive rules would require showing an urgent, countervailing need, perhaps a powerful benefit to the individual or even a broad societal need, not simply consent, where we are back with the insurmountable problems. Beyond these cases, there still is room notice and consent, where people may have different tastes and legitimate preferences. As a social matter, some prefer French dressing, others like oil and vinegar. Some areas of information sharing are clearly in that category and it is in these, where privacy involve choice or a preference where I see a circumscribed role for notice and consent. As to the idea that we don't have any privacy anyway, I flinch whenever I hear it because it reveals a bizarre understanding about privacy and its worth. A quick anecdote illustrates the point: Recently, I was invited to participate in a workshop organized by one of the large data companies whom many consider the "arch nemesis." The invitation stated that Chatham house rules would

prevail. That, if nothing else, is a privacy rule. We believe that living in a democracy under rule of law means that information collected by the police inappropriately can't be used in a court of law and banks do not play loosely with their customers' PIN numbers. These are all privacy rules. Identifying these appropriate flows of information in a social life managed and mediate by digital technologies, and a resolve to protect them, constitutes privacy worth fighting for.

RYAN HARKINS: Quickly on the burden of proof question—I think that is a really interesting question and frankly most privacy laws don't expressly address that question. There are some exceptions. There's a relatively new law in Canada called "CASL," the Canada's Anti-Spam Legislation, which expressly provides that the company will have the burden of proving that the user provide consent. But I think obviously, if I were to recall my old litigator days, if a plaintiff is bringing a case and alleging some sort of privacy violation, the plaintiff would have the burden of proving his or her case under the law today. As far as Ira's question, I'm not sure we would say that notice and consent would operate in a diminished capacity. I think we would say that notice and consent would operate in a strengthened and adapted capacity, but we certainly don't think that notice and consent should be abandoned. We think it should be adapted and frankly we think that should be coupled with the bolstering other aspects of the FIPPs to help provide more privacy protection.

LORRIE CRANOR: So I don't think we should abandon notice and consent, but I think the problem right now in the U.S. policy is that in many domains that is the only form of privacy regulation we have. We've seen repeatedly that there are calls for doing something about privacy and the solution is notice and choice, and sometimes it's regulation that requires it and sometimes it's, "Oh, we'll have an industry self-regulatory program of voluntary compliance of notice and choice," and everybody congratulates themselves and says, "Yes, we've solved the privacy problem in this domain" and they stop. And I think that we have to realize that notice and choice is not solving the problem there and I think it can supplement some other privacy solutions. I also think that when we talk about notice and choice, I think the notice part is actually more important than the choice. As much as I would like to have informed consent and choice, I feel like that's actually not really all that realistic because we don't want to spend 200 hours a year being informed so that we can have that. I think having the notice so that in the event that you want to find out what's going on, you can do that. I think that's really important. It would be nice if we had actually

complete notices, so that we know not only that they're sharing, but actually why they're sharing and things like that. Thinking that people are making informed choices here with these notices, I'm not sure that that's really realistic.

BRETT FRISCHMANN: I want to thank our panel for a fantastic discussion. We could go on for hours, I think, continuing this conversation. But we're going to have a five-minute break to grab coffee or whatnot, and then we'll start the second panel. So let's thank the panelists.

PANEL II: BALANCING NATIONAL SECURITY AND TRANSPARENCY IN GOVERNMENT DATA COLLECTION

NATE CARDOZO*

MARIKO HIROSE*

JONATHAN MANES*

IRA RUBINSTEIN*

CHRISTOPHER WOLF*

JONATHAN MANES: Welcome all to the second panel of today's symposium. My name is Jonathan Manes. I will be your moderator for the second panel. I am an Abrams Clinical Fellow and a clinical lecturer at Yale Law School. I teach in the Media Freedom and Information Access Clinic and am a fellow in Information Society Project. Full disclosure: my clinic has a few cases seeking greater transparency on national security and surveillance issues, both in the Foreign Intelligence Surveillance Act ("FISA") court and in the district courts. So I am certainly interested in these issues.

The title of today's panel is *Balancing National Security and Transparency in Government Data Collection*. Many thanks to Pamela for putting together such a wonderful group of practitioners and scholars who have thought very deeply about these issues. I will first give you a short introduction to the topic and will try to get out of the way as quickly as possible, because the panelists have much more interesting things to say than I do.

We are all here discussing these issues because in June of last year *The Guardian* began publishing news stories based on the Snowden disclosures, at which point we realized how little we knew about the government's surveillance practices. The revelations have pointed out some very serious privacy concerns with the government's collection programs, and also concerns about the way that our transparency system works, or does not work. Many people agree that much of the material

* Staff Attorney on the Electronic Frontier Foundation's digital civil liberties team.
* Staff Attorney at the New York Civil Liberties Union.
* Clinical Lecturer in Law, Associate Research Scholar in Law, and Abrams Clinical Fellow of the Information Society Project at Yale Law School.
* Senior Fellow at the Information Law Institute at New York University School of Law.
* Partner at Hogan Lovells, Washington, D.C. office and director of Hogan Lovells' Privacy and Information Management practice group.

that has been disclosed should previously have been made public by the government, one way or another. Some people disagree, and there is significant controversy around the edges. But, at the very least, the disclosures revealed how much in the dark most of the public was. In today's panel, we are going to discuss the privacy issues that these disclosures raised. We will also discuss various mechanisms for achieving transparency on government data collection and such mechanisms should work or how they could work better.

The mechanics of today's discussion: because there are four panelists, I have asked them to limit their individual remarks to ten minutes, which will leave forty minutes after for discussion. Like the previous panel, I will try to provoke the panelists to talk to each other because I think we have a great range of viewpoints represented here. Then I might ask a few questions and open it up to the audience. I will briefly introduce our panelists. I won't be able to do each one of our panelists justice, unfortunately, but I will give it a shot.

The first person speaking today will be Mariko Hirose. She is a staff attorney at the New York Civil Liberties Union and, like many American Civil Liberties Union ("ACLU") lawyers, she has a fascinating and varied docket of cases—everything from LGBT rights to privacy and surveillance to anti-discrimination lawsuits. She previously worked at the ACLU's Speech, Privacy and Technology Project. She will be giving us a brief overview of what issues are at stake in these government data collection and surveillance programs.

Next up will be Chris Wolf. He is the head of the Global Privacy and Information Management Practice at Hogan Lovells in Washington D.C. and is a leader in the field of privacy law and on issues of Internet law and computer law generally—one of the pioneers in that field. He is the cofounder of the Freedom of Privacy Forum, which I am sure that you have come across if you are involved in these issues. He has leadership positions in many other civil society and non-profit organizations, and has published widely in his field. A number of those articles are in your materials today. He is going to be talking about company transparency reports and providing us with a comparative perspective on how the United States and other countries compare on these data collection and surveillance issues.

Following that will be Nate Cardozo. He is a staff attorney at the Electronic Frontier Foundation ("EFF"). Nate is on EFF's digital civil liberties team and works on free speech and privacy advocacy and litigation. Like Mariko, Nate has a very interesting and varied docket—everything from automotive privacy to hardware hacking rights to anonymous speech. Of course, EFF is right in the middle of these issues, both on the policy and the litigation side.

Finally, Ira Rubinstein is a senior fellow at NYU's Information Law Institute. He researches and writes on online privacy, electronic surveillance, big data, and Internet security—essentially all of the issues that we are talking about today. He was for seventeen years a lawyer at Microsoft, most recently as associate general counsel. So he has also seen these issues from the perspective of the companies who are asked to disclose—or forced to disclose—information to the government. We are very much looking forward to his perspective. His remarks will provide a bigger picture view of how to think about transparency, what makes for effective transparency, and whether transparency is effective in the context of government data collection.

So I'll turn it over first to Mariko and I look forward to all of our panelists' remarks. Thanks so very much.

MARIKO HIROSE: Thank you, Jonathan, for the introduction and thank you to the *Cardozo Arts and Entertainment Law Journal* for putting together this panel. We've had an extraordinary year in the debate on privacy and transparency in government data collection and I'm honored to be on this panel with so many experts who have been thinking about this issue for a long time. I have worked on a number of privacy issues, including cellphone location tracking, GPS tracking of government employees, Manhattan District Attorney's subpoena to Twitter for a user's customer information and a state effort to subpoena purchase records of Amazon customers. One of the things that I hope I will do in this short speech is to tie in how the discussion around the telephone records program revealed by Edward Snowden relates to themes that arise in other law enforcement efforts to collect information.

This extraordinary year that I referenced and that Jonathan has referenced all began with a shocking revelation in June of last year that the NSA has been for years conducting mass, suspicion-less collection of phone records of millions of Americans. This mass collection was, according to the NSA, authorized under Section 215 of the Patriot Act, which allows the government to obtain orders requiring the production of any tangible things relevant to an authorized foreign intelligence or terrorism investigations.

The privacy issues at stake in the NSA telephone records collection program are quite serious. While the program does not collect the content of any phone calls and is limited to so-called "metadata," the metadata it collects include information on an ongoing basis of all phone numbers called from a number, the time of a call, and duration of a call, for millions of calls across the United States. That kind of information implicates privacy concerns. Even at a level of a single individual or a single call, information about who a person calls and

when can reveal a lot about a person. For example, a person may call numbers dedicated to certain type of issues, like an abortion clinic, religious counseling, or mental health hotlines. Or a person may call our office hoping to speak to someone anonymously to report a problem or to seek advice. A person may call a certain person's phone number every weekend night at 1:00 a.m. All of these things reveal a lot of private information about a person, and when this type of information is aggregated over time and over a great number of people, the pieces of metadata paint an increasingly more detailed picture of people's associations, beliefs, and activities.

But despite the sweeping nature and the grave privacy implications of the phone records program, it operated for years without a challenge to its legality or its necessity. This was allowed to go on because of the structure of the Foreign Intelligence Surveillance Court, which meets is secret, generally only hears from the government, and does not ordinarily publish decisions. The most we heard as members of the public were oblique references from certain senators that the government had adopted stunning secret interpretations of Section 215. But they were unable to make more information available to the public. All of this changed when journalists informed by Edward Snowden's disclosure revealed the existence of this phone records program. And the reaction to that disclosure confirms the importance of transparency of government action because transparency has led to informed thinking and to accountability and to debate over the legality and the necessity of the program by all branches of the government.

First, at the very basic level of democracy, the disclosure has led to a more informed American public. The American public has the right to know in general terms at least the scope of its government's surveillance actions and the claimed legal basis for it. That is how our democracy is set up to work.

Second, the disclosure pushed the legislative and executive branches of government into action. Because of the disclosure of the program, Congress was able to debate the scope of the program and introduce a number of bills proposing reform. The debates also prompted President Obama to appoint two independent review panels to study how the White House should change or limit NSA's surveillance programs. Both of those groups came out with reports condemning the program from a legal and policy perspective, leading the President to propose reforms to the system.

Third, the disclosure allowed for cases to be filed in open federal court challenging the surveillance program—with the benefit of full briefing in an adversarial context, including from the people who are affected by the surveillance. This is how constitutional issues are meant

to be decided in our system. And indeed, the first federal court to decide the constitutionality of NSA's mass collection of telephone records found the program to be "almost Orwellian" and unconstitutional. Notably, the court rejected the government's primary defense of the constitutionality of this program: its argument that the Fourth Amendment did not apply because the program involved only metadata kept by third-party telecommunications providers. The government made this argument by relying on a case called *Smith v. Maryland*, which is a 1979 case in which the Supreme Court held that the police did not violate the Fourth Amendment by installing a pen register on a phone line of a target of an investigation over a few days. This pen register was a primitive device that recorded the phone numbers dialed from the phone line under surveillance. And the court had held that in those circumstances, the person had no reasonable expectation of privacy in the numbers dialed from that specific number that was under surveillance because he voluntarily transmitted them to his phone company and because it was generally known that the phone companies kept this information in their business records. The District of Columbia, which was the first federal district court to issue an opinion on the NSA's surveillance program, held that *Smith v. Maryland* is of little value in evaluating the constitutionality of the phone records program given the limited nature of the pen register involved in *Smith v. Maryland*. The court recognized the difference between collection with limited technology and the bulk collection of data over a long period of time that is happening now and that is possible.

Now although the second federal court to decide this issue—the Southern District of New York—decided the other way, the point is that this very important issue about the applicability of a thirty-plus year old Supreme Court case is now making its way up the appellate system. And this issue is very important not only with regards to the constitutionality of this telephone records program, but in relation to so many more police practices involving new technology. For example, cellphone location tracking is one area in which this issue comes up all the time. Basically whenever your cellphone is on, it is communicating with cell phone towers nearby. And the government is able to track your location or construct your historical location information and movements by seeking those location records from your cellphone location providers. In that context too, the government has been arguing that it can obtain those records from the cellphone location providers without a warrant because there is no Fourth Amendment protection under *Smith v. Maryland*—because that cellphone location records are business records maintained by third parties under *Smith v. Maryland*.

But just as the District of Columbia recognized in the NSA

surveillance case, the cellphone location technology available now is more invasive than the simple technology involved in *Smith v. Maryland*. *Smith v. Maryland* should not apply to cellphone location tracking either. Not only has the technology gotten more advanced, it has gotten cheaper, which means that there are few technological or cost barriers to conducting mass surveillance. This mass surveillance paints a highly detailed picture of a person's life. The debate about the phone records program is important because that debate is necessary for both the public and the courts to evaluate that program, but also because what we decide about it also informs a lot of other debates about the legality of other types of warrantless law enforcement requests for information.

There's still a lot of debate to come in the public and in all branches of the government about the legality of the phone records program as well as all sorts of other law enforcement surveillance issues. Undoubtedly, people will come out on both sides of that debate. But I think this is the value of transparency—it allows for this informed public debate. Thank you.

JONATHAN MANES: Next up, Chris Wolf.

CHRISTOPHER WOLF: Any discussion of this topic is not possible without a picture of Edward Snowden [displays picture of Edward Snowden on PowerPoint], so I'm satisfying that obligation. Sorry that the panelists can't see it.

Mariko Hirose, I think, did a good job describing the Section 215 issue and the metadata and *Smith v. Maryland* issue that it raises. The other aspect of the Snowden disclosure had to with the access by the NSA to data held by major Internet companies—Google, Microsoft, and others—and there was the suggestion in some of the disclosures that were made that the NSA had backdoor access to these companies and we heard almost immediately denials by those companies that such access exists. But even if that so-called "backdoor" doesn't exist, there is a general assumption that the government secretly is demanding troves of data from intermediaries through FISA warrants or through national security letters ("NSL") under the Patriot Act.

Last year the Information Technology and Innovation Foundation in D.C. said that U.S. cloud providers defined broadly could lose between $21.5 billion and $35 billion as a result of the Snowden revelations. "On the ground" studies since last June are really confirming those predictions. Earlier this year, NTT Communications commissioned a survey of 1,000 IT decision-makers from the United Kingdom, France, Germany, Hong Kong, and the United States about

their attitudes following the Snowden disclosures. That survey concluded that one in six businesses are either delaying or cancelling cloud computing contracts in light of revelations about the alleged signals, intelligence, and communications surveillance by the NSA and we've also seen data localization bills introduced in various countries to attempt to shield data from the NSA. Obviously the businesses using cloud computing are not the only ones that are concerned about the NSA's access. Consumers, civil liberty advocates, are also concerned about how exposed this data may be to government access.

So it's not surprising that reviews of government access have been conducted by a review group commissioned by the White House that issued its report in January as well as the Privacy and Civil Liberties Oversight Board ("PCLOB"), which is continuing to review the situation I testified before the PCLOB two weeks ago on their Section 702 analysis. And a number of bills have been offered in congress to try to both create more transparency and also try to restrain and pull back some of the NSA's authority.

Last August—as he was about to leave office—the general counsel of the Commerce Department, Cam Kerry, who is John Kerry's brother, spoke to the German Marshall Fund as this issue was still boiling. I think it was not so long after the revelations about the hacking into Angela Merkel's cellphone. Kerry was trying to provide some context and said that if you take into account the amount of data the NSA's touches—to use his term—and he was relying on the NSA for the statistics, then, the limited traffic actually reviewed amounts to only 0.00004 percent of all Internet traffic. Kerry then concluded by saying—I'm quoting now—very simply "the United States Government is not listening to or reading everything said by everything citizen of any country."

That wasn't completely comforting as you might expect to those who were still concerned about the degree and level of access, which is why Internet intermediaries have been very eager to release statistics on government access because they believe the statistics show that the actual legal requests for documents—assuming there is no back door, and they've denied the back door, and I'm not in any position to comment on whether that exists or not—that they actually get from the NSA—from the FBI—for national security-related investigations are relatively small. Google was the first to issue such a transparency report, and they've been doing it for a while. It wasn't prompted by the Snowden revelations. In the PowerPoint slide that you're seeing, Google shows how many user data requests it has received from the government since the end of 2009 and in that time—between 2009 and today—the requests have grown from 12,000 a year to over 27,000.

Google discloses the percentage of requests that were fully or partially complied with, and they explain that they have teams of lawyers—as many of the companies do—that look at the requests and determine whether or not they are actually logical and can be complied with, whether they are overly broad, and whether they really are not designed to get the information that they purportedly are seeking.

It is interesting that this data shows that Google only complies with—that means satisfies the original request as written—64 percent of the requests. These transparency reports are helpful, but they are not terribly detailed or granular. So the companies really would like the ability to provide more granularity and more detail, and they feel hamstrung by the limitations that are put them in terms of both the types of frequency of the requests that may be disclosed publicly.

The government, on the other hand, thinks that if there is more granularity, more detail, and more timely reports then the people being surveilled will be on to what the government is doing—as if they are not already. But the concern is that the methods and types of surveillance used will be revealed and will simply drive the terrorists and others to other modes of communication. Again, I'm not in a position to evaluate the merits of that argument, but those are the two tensions that exist.

Nevertheless, a number of companies sued the Justice Department after the Snowden revelations to get permission to be more granular and to be more timely when releasing information about the requests. In January of this year, a deal was reached between the Justice Department that allows Google, Yahoo!, Microsoft, and other Internet companies—I do not think Apple was a party of the case but benefits from it by filing an amicus brief or a letter in support of the case—allowing Internet providers to be more specific about the requests they are getting from the government and to divide the requests into categories but only to approximate the numbers in each category. The companies can announce the number of national security letters they receive in a given year, for example, and the number of users who will be affected by them. But they can only round those numbers to the nearest thousand. Providers will be able to report the number of demands for metadata or the actual content of messages made using FISA, but only with a lag time of six months. The report dated December 31, 2014, for example, could only report on the requests made between January 1 of 2013 and June 30 of 2013.

So the new rules also require—this is interesting—that Internet companies wait two full years before reporting the first FISA request for information on a new platform of communication; for example, a text or a video service that they might roll out other than something that they traditionally have offered. Like the count of national security letters

requests, FISA orders can only be counted to the nearest thousand. Reports can get more precise about the numbers—rounded to the nearest 250 rather than the nearest 1,000—but only if they combine all national security letters and FISA orders into one category rather than detailing them. So here's a sample report from Verizon that covers 2013 [shown on PowerPoint]. Again, transparency reports can only come out every six months under the deal reached with the government. And as you can see, reporting on FISA orders is delayed for the July to December time period. We won't hear about that for a while.

These reports by the way don't just report on what the United States is requesting but to the extent that these companies are permitted to do so under the laws of other countries, they report on the requests that they receive from other countries. So obviously a question—and this is a roiling question in international circles and it has—is whether or not data should flow from the European Union to the United States under the EU/U.S. Safe Harbor Agreement and whether the United States really is a data gobbler more than any other country, *i.e.*, whether the U.S. national security agencies are asking for data from companies on a per capita basis more than other countries do.

Last year Hogan Lovells, my law firm, published a white paper that looked specifically at these transparency reports and compared government requests across many different countries for what we described as law enforcement purposes because in many countries it's not—there's no distinction between law enforcement and national security access in reporting these numbers. So this chart [shown on PowerPoint] shows the aggregate number per capita and per Internet users' requests for Google, Microsoft, Skype, Twitter, and LinkedIn for the year 2012 for these investigatory purposes. And you'll see that U.S. government requests totaled approximately 96 per capita and 119 per Internet user for these five companies in 2012. The chart shows that Taiwan, the United Kingdom, Hong Kong, France, Australia, and Germany made significantly more requests of these companies on a per capita basis than the United States, Belgium, Portugal, and Singapore; which made fewer requests.

What this tells us, we think, is that the United States on a comparative basis is not unique in the world in its desire for personal information for purposes of either law enforcement or national security; and, in fact, on a per capita basis, appears to make fewer requests than some of our democratic counterparts. We are in the process of updating these numbers with more reports that have come out since we did our white paper last August and the preliminary results show that the government requests for data have increased across the board—not just in the United States but in other countries as well. I'll be happy to talk

about that during the question-and-answer session as well as the issue of due process and legal restrictions on government access, which has transparency as a major element of that in the United States versus other countries around the world. So why don't I stop there. Thanks.

JONATHAN MANES: Many thanks, Chris. Next up is Nate Cardozo from the EFF.

NATE CARDOZO: Hi. I'm Nate Cardozo. Thanks to Pamela and the rest of the Cardozo team. I am not related to Benjamin Nathan Cardozo although my father's name was Benjamin Cardozo and my name is Nathan Cardozo, I'm no relation whatsoever. I would like to start by asking a couple questions and then reading a few quotes.

So, first of all, how many people in this room believe that the United States government could access the content of your communications if they wanted to? It seems like most people. How many people in this room believe the United States government could access the content of just about anyone's communications who uses an American service provider, and by American I mean having major operations in this country? A similar number. How many people in this room believe that terrorists—I'm going to use that word, interpret it however you want—also believe that the United States government could access their communications? Around the same, maybe fewer.

Okay. Now I'm going to read a few quotes. I'm going to read these in chronological order. The first one is from 1971, *New York Times v. United States*. This is the Pentagon Papers case, Justice William O. Douglas concurring. "Secrecy in government is fundamentally anti-democratic. Open debate and discussion of public issues are vital to our national health. On public questions, there should be uninhibited, robust, and wide open debate." I totally agree with that 100 percent. The next quote is from President Barack Obama from a press conference in August of last year. "We can and must be more transparent so I've directed the intelligence community to make public as much information about these programs as possible." I totally agree with that as well. The next and final quote that I'm going to read is from the Department of Justice's opposition to the companies' lawsuit in the Foreign Intelligence Surveillance Court that Chris mentioned earlier. "Although the government is seeking to make public as much information about these activities as national security interests of the United States will permit, in the FISA context there is an unquestioned tradition of secrecy based on the vitally important need to protect national security."

One of these quotes is not like the others and I would submit that

it's the third. The "unquestioned tradition of secrecy" in my mind is no excuse for the continuance of that tradition of secrecy especially given what William O. Douglas and President Barack Obama both said—that open debate is vital to the democratic process.

Some of you might know the Electronic Frontier Foundation—the organization I work for that Lorrie is on the board of. We are an impact litigation civil liberties organization dedicated to protecting users' rights online, and much like Jonathan's programs, we have several active lawsuits in the area. That might be a hint as to where I'm coming from on these matters.

I'm going to talk about two things. First, I'm going to elaborate a little bit more about the deal reached between the five Internet companies and the Department of Justice ("DOJ") at the beginning of this year. And second I'm going to talk a little bit about the warrant canary. I'll talk about what that means later.

The deal that Chris introduced—I want to first point out that it is non-binding. It is not a settlement. It is not based in any sort of statutory language whatsoever and it is not enshrined in any sort of court order. It is simply guidance as Chris called it. What is it guidance for? It is guidance about when the DOJ will believe that there has been a violation of law. It's not guidance about when a court will necessarily believe that there was a violation of law—only the DOJ. And it does something interesting. It tries to create separate rules for reporting government demands in the name of law enforcement and those in the name of national security. I'm not sure that that's right, and I'm not sure that that's fundamentally consistent with the democratic process. And I'll talk a little bit about why I think that's true.

As Chris said, the guidance allows companies to disclose requests in two separate "buckets" if they want to disclose within ranges of 1,000 or in only one "bucket" if they want to disclose ranges of 250. The two buckets available to disclose in the ranges of 1,000 are national security letters and all other types of national security orders, primarily FISA orders. Notably absent from that as we could see from the Verizon transparency report that Chris pointed out, under the DOJ's guidance, no company is allowed to disclose the mass collection of Americans' telephone records that Edward Snowden keyed us into in June of last year—bulk collection under Section 215 of the Patriot Act. On one hand, the guidance allows companies to disclose the number national security letters that they receive, and the number of accounts affected. On the other hand, in the national security order context, companies are allowed to disclose the number or orders received and the number of *selectors* targeted.

All right. What does that mean? A selector is a name. If my

national security order says, "I want everything about Nate Cardozo," that is one selector targeted. But if my national security order says, "I want all of your company's metadata about everyone that you serve," that's zero selectors targeted. So in the Verizon report, we saw the company disclose the number of national security orders that they'd received, but no information about the number of customer accounts affected. We know of course from the Snowden disclosures that the percentage of accounts affected is 100, but the DOJ refuses to allow Verizon to disclose that fact.

The government argues—as Chris pointed out—that the justification for such restriction is that if terrorists knew which companies were cooperating with the United States government, they would move to other services. I would submit that that's not true. I would submit that everyone here in this room—and I don't think terrorists are any different—understands that all American companies cooperate with American government requests for user data. I don't think there's anything shocking about that. The government disagrees. In one of our cases—*Hepting v. NSA*, filed in 2009—we've been suing the NSA over the very program the Snowden disclosures confirmed in 2013. The government has consistently argued that the fact that AT&T cooperates with the NSA is a state secret, even though it was front-page news in *USA Today* in 2005, and has been front-page news ever since then. I think that in and of itself dismantles the government's justification for keeping these numbers secret.

But why else is it important that we know what authorities the government is using to get at our data? And what does this deal leave out besides the bulk collection? The deal leaves out the distinction between Section 215 orders as Mariko discussed, which are orders for metadata, and orders under Section 702 of FISA Amendments Act, which are orders for content.

While, the media will often talk about the PRISM program, which is probably a wildly inaccurate term, as we don't actually know what PRISM means. From what we can tell from context in the Snowden documents, we think PRISM probably means the interface that analysts use to search the data. PRISM doesn't mean the program itself, so I'll call it the 702 Program. Hypothetically, the 702 Program is limited to collection of what are called one-end foreign communications (a communication between an American and someone who is not based in the United States).

Companies are not allowed under the DOJ rules to disclose any information about the breakdown in requests they receive, between Section 702 orders for content and Section 215 orders for metadata. Why is that critically important? There are bills in Congress right now

that would reform Section 215 and not 702. There are other proposed bills that would reform Section 702 and not Section 215. And there are still other proposed bills right now that would reform both. When Congress is actively debating reforming authorities that the U.S. government uses to get our data, and the government attempts to prevent us from learning how often it uses those very authorities, we should understand that attempt as profoundly anti-democratic, just as William O. Douglas pointed out in the Pentagon Papers case in 1971.

We are at a crossroads in terms of reforming the American intelligence regime and the American surveillance state. We need to know how much the government is resorting to various authorities and about how often the government is not resorting to any statutory authority but instead uses its inherent authority, for instance under Executive Order 12333 ("E.O. 12333"), to collect certain information outside the United States. Even putting E.O. 12333 aside, when we're debating reforming surveillance authorities, we need to know how often they're being used. And the DOJ's guidance fails miserably to further that debate.

Who here is familiar with the term "warrant canary?" Very, very few. Okay. A warrant canary is a statement by a company or any kind of online service provider, that it has not received any of a particular kind of request. For instance, say I run a small server to provide email for my wife and myself. I could have a statement on my website that says I have received zero subpoenas, zero warrants, zero national security letters, and zero national security orders. That statement is true. I have not received any request for my data from the government. That's a warrant canary.

There are various interpretations of how this might work in context. R-Sync, which is a service provider that's been around a very long time, has one that updates daily. Apple introduced a warrant canary in its second to last transparency report and took it away in its most recent. Its second-to-last transparency report said that it had received no Section 215 orders. And as Mariko discussed, a Section 215 order is most likely used to require it to turn over metadata.

It's not particularly surprising that it hasn't received any Section 215 orders because it doesn't really have metadata. What it has is content. So when we saw Apple's name show up on the PRISM slide back in June of last year, those orders would be most likely Section 702 orders and so that wouldn't be triggered by its particular warrant canary.

The DOJ's guidance seems to prohibit the use of warrant canaries. Tumblr and Pinterest were in our view was extraordinarily brave and published strong warrant canaries. They said that they had received zero NSLs and zero FISA orders. They did not say that they received

between zero and 249 orders as the DOJ would have liked them to receive.

The DOJ's position is that Tumblr is in violation of the law although I don't know what law that would be. But the DOJ says that Tumblr should only be allowed to publish the fact that they've received between zero and 249 national security requests. I think that's ridiculous. I'm not sure that there's any possible justification for that.

And Tumblr did something, which I actually would consider a best practice for warrant canaries. They are not publishing a daily warrant canary. They are not publishing a monthly. They're not even publishing every six months. They're publishing the canary every six months with a six-month delay. Why is that important? Right now there's no case law on whether a warrant canary is legal. If Tumblr receives a national security letter in the next six months, it will have six months to litigate whether it can continue to publish its warrant canary, which in this case would mean simply removing the line item in its transparency report saying we've received zero. That means there will be some lag between the time that it receives its first NSL and the time that its next transparency report covering that time period would be due, so that it could file a suit for declaratory relief demanding permission to remove rather than lie about the number of national security letters that it has received.

If in the next Tumblr transparency report we see that they say that they received this many warrants, this many subpoenas and this many national security orders, but nothing about NSLs, we might start to think about what that means. And that's why it's called the canary—because it stops tweeting and the canary dies. It's a canary in the coalmine. So Tumblr did it right. Tumblr is only publishing these things once every six months and only—not for the prior six months but for the six months before that.

JONATHAN MANES: Many thanks, Nate. Next up is Ira Rubinstein.

IRA RUBINSTEIN: Thank you. I also want to thank Pamela and Jonathan for organizing a terrific conference and an excellent panel. I'm going to talk a little more about the effectiveness of transparency mechanisms with a particular focus on company transparency reports. And in doing so, I'm going to relate this national security-oriented panel to some of the discussion from the first panel on Disclosure and Notice Practices in Private Data Collection. I'm going to move pretty quickly through my first couple of slides because my co-panelists have addressed a lot of these issues. So just quickly to set some ground rules here. There are gag provisions in a lot of the relevant statutes under which the

government serves orders on private firms, and that's why we have debates over transparency. But I also want to mention that the treatment of the press and of whistleblowers is highly significant. I'm not going to say a lot about that, other than to point out that it's really the Snowden revelations that inspired the public debate over transparency much more than what's contained in the industry transparency reports and that has led to some of the significant litigation challenging government practices, as both Nate and Mariko pointed out earlier.

So one possible framework—although I'm not going to spend much time on it—is secrecy. And secrecy has a clear rationale in the national security context, but it's also often abused to hide or to avoid oversight and debate or to hide abuses or to resist change and censure. I list some obvious examples in slide 4 [classified documents, secret court orders, secret legal interpretations]. And I think what has driven a lot of the public debate are the secret legal interpretations—in particular around Section 215. Indeed, many Members of Congress were surprised to learn that Section 215 was the basis for the telephony metadata program. As Chris pointed out, there's a host of new transparency proposals coming out of the administration. Nate alluded to the Administration's recent announcement and that's something to take seriously.

There's been a significant effort by the FISA court to begin declassifying the FISA opinions, which used to be secret and unpublished. Another really interesting development is that the DNI—the Director of National Intelligence—now has a transparency website, which I recommend that you visit. It's quite interesting that this agency is now in the transparency business, and how they handle transparency is quite fascinating. There have been a number of industry proposals, too. Some are just reform proposals, and the industry has also sued the government to allow greater transparency and this case has now been settled. There's also a host of bills on the Hill and the two groups that Chris mentioned—the President's Review Group and the PCLOB have also heavily emphasized the need for greater transparency, both on the part of government institutions and the courts, and also on the part of industry. Then, of course, another lens through which to see this is that of transparency and democracy with its emphasis on accountability and trust. And again, both Nate and Mariko emphasized this.

Now I want to shift gears a little bit and discuss these transparency reports in the context of the first panel, which talked about what makes transparency effective. How is notice and choice effective? But this discussion requires that we step back a bit from the privacy context and discuss transparency more generally by looking at a book by Archon Fung, called *Full Disclosure: The Perils and Promise of Transparency*.

This book is the bible of disclosure. Fung and his co-authors looked at literally hundreds of transparency program ranging from our favorite—the nutrition labels—to corporate finance, to the restaurant hygiene program in New York (the ABC ratings), mortgage lending, and workplace hazards, to name just a few. And they found that these programs all shared a number of characteristics and that they tended to work in the same way and that the successful ones had common characteristics. So let's just look at these characteristics briefly.

Slide 8 identifies the five common characteristics they found (a specific policy purpose, specified disclosure targets, a defined scope of information, a defined information structure and vehicle, and an enforcement mechanism) and I'll talk about them just, again, very briefly with reference to nutrition labeling, which is one of the programs they studied and they considered moderately effective, as well as restaurant hygiene ratings, which they considered highly effective. So in terms of purpose, nutrition labeling seeks to reduce chronic disease; it targets consumers; in terms of scope, it's somewhat limited to grocery items rather than fast food and restaurants and delis across the board, although that's changing. In terms of structure, there are these easy-to-read labels with a standardized format. And in terms of enforcement, there's both a legal and a market mechanism.

On the restaurant hygiene side, you have something very similar. The goal is to reduce the risk of foodborne illness aimed at diners. The scope is reporting on hygiene inspection. The letter grades are the most telling difference between the two programs at least in terms of their common characteristics because there's nothing as straightforward and simple as a letter grade. How many of you have looked at these letters? If you see a C, you probably don't go in the restaurant. That's incredibly effective. It's easily understandable, it's extremely timely, and it has an immediate impact on your decision-making. But it's really more on this question of embeddedness in decision-making that distinguishes the programs. This is something that Lorrie Cranor would probably refer to more to in terms of usability or useable design. Fung et al. identified these three factors of value, compatibility and comprehensibility. And—as slide 9 indicates—they look at that both from the perspective of the users (for users, the three factors mean that information helps them achieve their goals, fits with their decision-making routines, and is readily comprehensible) and the regulated businesses that actually disclose the information—the disclosers (for disclosers, the three factors mean user responses (behavioral change) affect core organizational goals, such responses are compatible with how business managers receive, process, and act on new info, and managers readily comprehend them).

And here's where the nutrition labels are a little less effective than we might think. In terms of the consumers' goals, often consumers are driven by price considerations. And that overrides their ability or willingness to look at nutrition labels particularly if they're in a low-income bracket or if they're older and might also find it difficult to understand that information. In terms of company responses, I think there it's a little more mixed. Companies have maybe begun marketing to these nutrition labels. Maybe I'm in an older age bracket, so I look at this more seriously. But I certainly look at fiber content on cereal boxes. Maybe some of you make your own relevant comparisons there. But at the same time that's not the only response of companies. They've also continued intensive marketing of high fat, high sugar, and high salt products. So there's not just a single signal that these labels are sending in terms of consumer response and the company response to the consumer response.

On the restaurant side, again, you see some real key benefits of the program. The labels are highly relevant. The grade is directly related to a consumer decision whether to eat there or not. It's easy to understand. It's very timely. And I think the restaurants can very readily understand that a C means a drop off in revenue and if they want to they can easily measure that. So with this in mind let me conclude by asking, based on these characteristics, how well do the company transparency reports perform? Because as odd as it may sound, these transparency reports are still a type of company disclosure, which can be assessed in the same way as nutritional labels and restaurant ratings. Now during the discussion period, we can talk more about how the transparency reports inform the debate over secrecy as well as the debate over democracy. But at the end of the day if we're going to evaluate these reports, we need to think about them as a form of disclosure, which may or may not be effective. On the basis of Fung's work, I would have to say that they're pretty ineffective. The policy purpose they serve is not going to be evident to most consumers. If consumers are the target, it seems extremely unlikely that they're going to have any impact on a consumer's choice of whether or not to use a particular service, in part because they're not all that salient, in part because most consumers won't understand what they are being told and how it compares with what they might be told at another site. These disclosures report on the scope and the number of government requests, but what does that matter? Compared to what?

So there's this tension between who is disclosing and to whom and for what purpose, and whether the goal is to inform public debate, which is a very broad goal as compared to informing consumers about nutrition or hygiene. So I think if you want to see these in terms of

electoral politics, then perhaps they play a very significant role, but if you see them in terms of informing consumer or citizen behavior or as creating a feedback loop back to the companies they are much less successful. After all, what does a company do with the consumer response to these reports? I can imagine Ryan Harkins, a Microsoft attorney, having a conversation internal to Microsoft where he might say, "you know, P3P," which Lorrie talked about, "that was an interesting initiative. Let's rethink P3P. Let's really make it work." And he'll go to a product manager and they'll have a conversation about how that might serve public policy needs, how that might help Microsoft interact with government regulators, how it might even lead to some competitive marketing against other firms who aren't supporting it. But it's very hard to imagine a similar conversation taking place with Ryan going to the managers of Hotmail and saying, "These NSA transparency reports, let's really get behind those." That really just doesn't work. So my point is that there are some very significant limits to these transparency reports evaluated as such, as a manner of disclosing information to company's consumers. And much as in the first panel, I think the conclusion to be drawn from this is that it's really the substantive reforms around these extensive programs of surveillance that are going to be more significant than just transparency for its own sake. Thank you.

JONATHAN MANES: So thank you to our panelists. In the interest of time, I will just ask a couple quick questions and then open it up for everyone else to participate. I will also encourage our panelists to respond to one another's views.

There has been a lot of discussion already about the companies' transparency reports and I'm sensing perhaps some disagreement on the panel on two issues. One point of disagreement might be whether the information that companies are now permitted to disclose is sufficient. What are we not getting that we should be getting? Nate mentioned a number of gaps in the transparency reports. Another gap that wasn't mentioned was, for example, the inability of companies to disclose what is covered by a National Security Letter order—what counts as an "electronic communications transactional record" that can be swept up under such an order. This links up with the last panel where we were discussing how much personal information companies gather up about us. So perhaps we could have a discussion about whether the transparency reports are sufficient in these regards.

It also seems that companies may be an odd place to look for transparency about issues of public policy because they may not be particularly concerned with informing the public debate. You can

imagine that a company might not have incentives aligned with those of a citizen who wants to be well informed. Maybe we will start with Chris and then we can move down the line on those two issues.

CHRISTOPHER WOLF: Let me take your second question first, and I think that I disagree with your premise that companies are not in the business of trying to serve the public interest with their transparency reports. I think companies would like nothing better than to get out of the business of having to respond to national security requests. Believe me, it's a burden with no benefit. And they're now finding it to be a real commercial disadvantage and this gets to, I think Ira Rubinstein, your point. What purpose do the transparency reports serve? To individual consumers, probably not a whole lot. I'm not going to choose my email provider based on a transparency report. But it's a lot like privacy policies. I was sitting in the back of the room for the first panel. While the consumer may not be well-informed, advocates are, regulators are, people who care about these issues are, and decisions get made based on what is contained in those privacy policies.

Let me just start there by saying I defend a lot of companies that are under investigation because they've said one thing in a privacy policy and the allegation is that they've done something else. So the policies serve a really important public purpose there because it gives the regulators a hook to use. And I think, likewise, the transparency reports give the companies a hook to use first of all commercially when the rest of the world is saying it's terrible what's going on in America. Nothing like this would happen in our country and the companies say, "well, yeah, we get the same kind of requests from governments around the world," which is a subject of some of our other white papers that show how freely those requests are made even without what some people might think are loose due process protections here. But I also think the companies really do want to influence public policy by showing that they are subject to the request, and maybe even tamping down some of the heated rhetoric about the extent of the requests—assuming there are no backdoors—that the number of requests are relatively small. You wouldn't conclude that from *The Guardian* headlines, but if you look at the transparency reports, they really don't seem to be that many.

JONATHAN MANES: Does anyone else on the panel want to address that?

IRA RUBINSTEIN: Sure. I mean, if I were to ask myself my own question from the first panel about whether I'd eliminate transparency reports,

I'd answer "no." I agree with Chris that they serve an important function in terms of informing the public at large as well as advocates and regulators. So I would certainly want to maintain them. And yet at the same time I think that over time—once they're established—they'll not necessarily play a very significant role for some of the reasons that I've mentioned, and making them more salient will be difficult because saying that the number of actual orders is small—it's very hard to grapple with small as compared to what or small in what context.

And, most importantly, I would say that if you want a test of whether a particular focus on transparency is really significant—and I'm not necessarily suggesting that the companies or Chris or anyone is saying it should only be about transparency—but if you want to get the measure of how important it is, think about what a U.S. company would discuss with European regulators who, as Chris has pointed out, have begun to question and criticize the safe harbor agreement. And they are considering whether to impose some very restrictive regulatory ideas on what U.S. companies can do in response to data requests from U.S. government agencies. So ask yourself if a U.S. company would respond to these concerns by telling the European regulators: "We're more transparent now so don't worry." I don't think that works. I think what the European regulators are looking for, in order to be satisfied that European citizen data should and can be held by and shared with U.S. companies, are substantive changes in U.S. law, not just more extensive transparency reports.

MARIKO HIROSE: I definitely agree with Ira about the limits of the transparency report. I also agree with him that it is better to have such reports than not, but there are limitations. The problem is not only that the reports may not fully inform the public of what's going on, but it's hard to use these reports to get the legal issues before a court. In terms of other transparency measures, it's important to have courts publish decisions—the legal basis—for their decisions. And to have a system in place where it's not just the government making its case for requesting certain types of information.

On the question of whether a company is a good place to go to for transparency, companies have definitely been helpful allies on this issue in many instances. There are companies like Amazon and Twitter that have gone into court on behalf of their customers. That is important because without the companies standing up for their customers' rights, nobody may have known about these subpoenas. I think it's also important, though, to have other mechanisms to ensure that we don't have to rely solely on the companies' willingness to stand up for their customers' privacy rights. There has to be another mechanism to make

sure that these privacy issues are being raised on behalf of customers in courts.

JONATHAN MANES: I might just open it up to the floor for questions.

AUDIENCE MEMBER: [inaudible question]

MARIKO HIROSE: I think that illustrates one reason why we can't just rely on the companies to raise the privacy rights of consumers, because they also have other interests at stake and that's part of the reason that the ACLU intervened on behalf of customers in the Amazon case when Amazon went to court to oppose the subpoena served on them for customer records. So one way for customers to get involved is to intervene in a case where they can. I think also that judges are finding ways to address privacy concerns.

For a long time, for example, we didn't know that law enforcement was obtaining cell phone location information without a warrant because magistrate judges were granting the law enforcement applications and not publishing these decisions. There are many judges though who have started publishing the decisions—in particular, the legal basis for their decisions. And some judges have invited organizations like the ACLU and EFF and other organizations to submit amicus briefs so that the privacy arguments in opposition are fully briefed.

IRA RUBINSTEIN: I'd add two other things. One is that the FISA decisions are all based on a secret court that proceeds ex parte. Thus, in some of the reform proposals, there's an emphasis on creating some type of position for an attorney to argue the other side and represent citizens or consumers before the FISA court. A few different structures have been proposed, but this is a common theme in a lot of the new proposals. The second is whether the FISA court in particular ought to go in the same direction as the FTC has gone recently, which is to hire a chief technologist. And one of the criticisms of the President's Review Board, which was a four-person or . . .

JONATHAN MANES: Five.

IRA RUBINSTEIN: A five-person panel of primarily lawyers and policy people from the intelligence world, but no technologists. That was one of the criticisms of the Review Board. Both this group and the FISA court hear from the government about what's technologically feasible and consequential but the courts are in absolutely no position to assess

these statements from a technology perspective. And I think at least part of the deference and part of the naiveté you see in these opinions relates to a lack of technical sophistication that might be remedied by a court like the FISC having its own technological resources that can more critically respond to what it's being told in the underlying court documents.

CHRISTOPHER WOLF: I will say that I think judges are a lot more interested in becoming knowledgeable in this area than they used to be. The superior court in the District of Columbia, where I practice, had a full-day retreat to learn about privacy law last year, which is a good trend I think.

NATE CARDOZO: I would like to point a couple of things about the Twitter case specifically. The Twitter case was an extreme outlier in one particular way. The New York State judge in the Twitter case held that the user—the target of the subpoena—had no standing to challenge the subpoena. That's pretty much unique in American jurisprudence. Judges all around the country have agreed with—as far as I know only a couple of outliers, with the Twitter case you mentioned being one—that users have standing to challenge requests for their data. At EFF I run the "Who Has Your Back?" project where we give companies gold stars for best practices in defending your data from overbroad government requests and pushing back on inappropriate government requests. One of the gold stars we award is for companies that promise to give notice of a government demand for data to their customers in all circumstances unless they're prohibited from doing so. Because it's really the user that's in the best place—not the company—to push back on requests for data.

The Twitter case was unusual and we give them enormous credit for pushing back on that request. Then the judge, of course, used a financial lever to force Twitter's hand, a lever that is unavailable to be used against consumers. But giving notice to the user, I think, is something that takes the transparency-reporting concept and makes it personal and allows users to stand up for themselves.

Of course in some circumstances, users can't be tipped off if they're the target of the investigation, and there are methods to gag companies to prevent notice from going out, for instance Section 2705 of the Electronic Communications Privacy Act, some types of grand juries subpoenas, warrants, all allow the government to seek gags when it's really necessary to prohibit the user from getting notice. But unless the government uses one of those methods to prevent the company from notifying the customer, it is absolutely the best practice of the company

to give notice to the user so that they can bring the challenge in court.

JONATHAN MANES: I might take all the questions at once because we're running up against time. So over here . . .

AUDIENCE MEMBER: [inaudible question]

CHRISTOPHER WOLF: I think it's the latter.

NATE CARDOZO: It's absolutely the latter. The lawyers and policy people that I work with at the companies—I have relationships with over a dozen companies who do transparency reporting—love it. They view it as a way of saying, "look how little data we give to the government," not "oh my we're giving so much."

AUDIENCE MEMBER: [inaudible question]

CHRISTOPHER WOLF: That's a long discussion to have, but I would say stay tuned because I was out at the White House workshop on big data this week at Berkeley and John Podesta, who has been assigned by the president to write a report, came with a preview of that. So it's coming out on April 17th. Take a look at that and then I think we have another workshop we can have here.

NATE CARDOZO: I'm going to be on a panel next month in D.C. with Stu Baker, Alex Joel—and Joel Brenner—a difficult combination of names.

CHRISTOPHER WOLF: That will be lively. I can guarantee it.

NATE CARDOZO: Oh, yeah. The last panel I was on with Stu Baker he blamed EFF for 9/11. So . . .

JONATHAN MANES: We don't blame you.

IRA RUBINSTEIN: One final comment is that I think the calls for ending bulk data have been aimed more at the Section 215 program in part because there's an apparent lack of evidence that it's effective. But there are other programs that have been shown or may be shown to be effective, at least from the government's viewpoint, and this includes the Section 702 program. I think there will be a very serious public debate and very divided opinion with lots of people thinking that if these tools are effective then by all means the National Security Agency

should use them. So I don't think it's a foregone conclusion that ending bulk collection is the right outcome. It's just that that's obviously one of the policy levers to consider.

JONATHAN MANES: All right. Are there any other questions form the audience? Perhaps I can ask a few more. I wonder what Congress' role is in all of this. We have talked about how the courts might intervene in some of these cases to permit the companies to speak more freely. That's EFF's case currently pending in the Ninth Circuit. There are of course FOIA cases that seek to force more disclosure from the executive. But Congress writes the laws here. They wrote the laws that permit gag orders on companies. They wrote the Freedom of Information Act. They have enormous power to urge agencies to provide more transparency or members of Congress could even just release information themselves on the floor of Congress—although that would raise its own concerns. Does anyone have any thoughts about the role for that branch of government?

CHRISTOPHER WOLF: We've seen senators like Ron Wyden from Oregon and Mark Udall from Utah give us some hints of transparency on the floor of the senate. Right? They said, you know, Ron Wyden asked Director of National Intelligence Clapper a question that he knew that Clapper couldn't answer back in March of last year. So that was a nice little window. But in terms of what else Congress can do—all of the reform proposals currently in Congress have some sort of transparency aspect. Even Diane Feinstein's proposal, which would codify existing practices, which I would strongly urge you to call all of your Congress people and oppose—even Feinstein's bill includes a transparency portion.

JONATHAN MANES: Maybe one more question from me. There have recently been a number of remarkable speeches from the administration expressing a commitment to transparency. The President's speech on January 17th specifically addressed Section 215 orders as well as national security letters, which was something of a surprise to some because NSLs have not been a prominent part of the recent debate. And yet after these speeches, we have not seen all that much more affirmative movement. On the national security letter front, the companies did reach the settlement we have discussed, but it seems like that may be as far as the government is willing to go. Do you have a sense that this is as far as it goes in terms of proactive disclosure from the executive branch? Will there be a more durable transparency mechanism going forward once the current leaks run their course? Will

such transparency require the courts or Congress to step up and force continuing disclosures?

CHRISTOPHER WOLF: So the President's review group prepared a report, which the President is still reviewing, I think, and contemplating what to do with the recommendations. The PCLOB has not finalized its work, at least not on Section 702. So that report is still to come. So I think it's very much a hot topic in Washington. I think there will be a lot more discussion. I think that the Podesta big data report will probably touch on it in some respect.

JONATHAN MANES: All right. I think we should wrap it up. This has been a lot of fun. Many thanks to our panelists and to the organizers.

CLOSING REMARKS

MANAGING EDITOR*

On behalf of the *Cardozo Arts and Entertainment Law Journal*, a huge thank you to all of our distinguished panelists for your participation today. Thank you to Professor Frischmann and Professor Manes for their guidance in moderating these panels and to Professor Wu for helping us select and develop this data collection topic. Thanks to all of you for your active participation and gently letting me know when I was completely out of my element in discussing this topic with you. I have received more than my fair share of thanks today and I just want to pass on that gratitude to Francesca Montalvo, our fearless Editor-in-Chief. Thank you for leading Volume 32 to new heights and encouraging us to put together this panel today. And, of course, thanks to all of the Staff and Editors of AELJ for your dedication and hard work.

* Pamela Grutman, Managing Editor, CARDOZO ARTS & ENT. L.J. Vol. 32, J.D., Benjamin N. Cardozo School of Law, Class of 2014.