In the Matter of:

Informational Injury Workshop

December 12, 2017

Condensed Transcript with Word Index

For The Record, Inc.
(301) 870-8025 - www.ftrinc.net - (800) 921-5555

UNITED STATES FEDERAL TRADE COMMISSION
BUREAU OF CONSUMER PROTECTION
BUREAU OF ECONOMICS

INFORMATIONAL INJURY WORKSHOP

TUESDAY, DECEMBER 12, 2017
9:15 A.M.

CONSTITUTION CENTER CONFERENCE CENTER
400-7TH STREET, S.W.
WASHINGTON, D.C. 20024

Reported by: Susanne Bergling, RMR-CRR-CLR-CMRS

C O N T E N T S

P R O C E E D I N G S

- - - - -

(Proceeding called to order, 9:15 a.m.)

INTRODUCTORY REMARKS

MS. HAN: Good morning. Welcome. On behalf of my colleagues here at the Federal Trade Commission, I'm happy to welcome you here to our Informational Injury Workshop. We are happy to see so many of you here today.

My name is Cora Han, and I'm an attorney in the Division of Privacy and Identity Protection here at the Commission. My co-organizers for today's event are Jacqueline Connor, also from the Division of Privacy and Identity Protection, and two colleagues from the Bureau of Economics, Doug Smith and Dan Wood.

Before we get started with our program, I need to review a few administrative details. Please silence any mobile phones and devices. If you must use them during the workshop, please be respectful of the speakers and your fellow audience members.

Please be aware that if you leave the Constitution Center Building for any reason during the workshop, you will have to go back through security. Please bear this in mind and plan ahead, especially if you are participating in a panel, so we can do our best

to remain on schedule.

The restrooms are right outside the auditorium. The Plaza East Cafeteria is located inside the building, so you can use it without going through security again. The cafeteria is the place to go if you want coffee without having to leave the building. It will be open during the first part of our midmorning break until 11:00, and then after it will reopen at 11:30 and remain open until 3:00. Please remember, however, that no food or drink, other than water, is permitted in the auditorium.

Most of you received a lanyard with a plastic FTC event security badge. We reuse these for multiple events, so when you leave for the day, please return your badge to FTC event staff.

If an emergency occurs that requires you to leave this conference room but remain in the building, follow the instructions provided over the building PA system. If an emergency occurs that requires the evacuation of the building, an alarm will sound. Everyone should leave the building in an orderly manner through the main 7th Street exit. After leaving the building, turn left and proceed down 7th Street and across E Street to the FTC emergency assembly area. Remain there until instructed to return to the building. If you notice any

5

suspicious activity, please alert building security.

Please be advised that this event may be photographed, and it is being webcast and recorded. It is also being broadcast via Facebook Live. By participating in this event, you are agreeing that your image and anything you say or submit may be posted indefinitely at ftc.gov or on one of the Commission's publicly available social media sites.

We're happy to welcome those watching via the webcast. We will make the webcast and all of the workshop materials available online to create a lasting record for everyone interested in these issues.

There will be time for audience questions during each of the panels today. Question cards are available on the table in the hallway immediately outside of the auditorium and also from FTC staff inside the auditorium. If you have a question, please fill out a card, raise your hand, and someone will come and get it. Please understand that we may not be able to get to all audience questions, but we will do our best.

For those of you on Twitter, we will be tweeting today's workshop at #infoinjuryftc, and if you would like to ask a question via Twitter, please tweet your question using that hashtag.

Lastly, I want to say thank you to our panelists

6

for taking part today. We are grateful for your time and work in the privacy and security area. Aside from the folks that you will see on stage today, this program would not be possible without the great work done by Kristal Peters and Bruce Jennings, alongside our paralegal support today from Carrie Davis, Anne Blackman, Olivia Berry, Courtney Brown, Christine Barker, Amber Howe, and Aaron Kaufman.

Also providing invaluable support are Nathan Otlewski from our Division of Consumer and Business Education and Nicole Jones and Juliana Henderson from our Office of Public Affairs. Thank you, all.

Now, it is my honor to welcome our Acting Chairman, Maureen Ohlhausen, to give opening remarks.

(Applause.)

OPENING REMARKS

ACTING CHAIRMAN OHLHAUSEN: Well, good morning, everyone, and thank you for being here for today's workshop on information injury. I'd also like to thank all the staff who have worked so hard to make this workshop possible. Special thanks to Doug Smith and Dan Wood from the Bureau of Economics and Cora Han and Jacqueline Connor from the Bureau of Consumer Protection. The four of you exemplify the interbureau, cross-disciplinary collaboration that is a unique

7

strength of the FTC and that is particularly important for complex topics like informational injury.

So to start with the fundamentals, information injury is my term for the harms consumers may suffer from privacy and data security incidents. In a speech in September on painting the privacy landscape at the FTC, I described some of the types of injury to consumers that the FTC has encountered in its privacy and data security cases over the years. I also announced my three goals for this workshop.

First, better identify the qualitatively different types of injury to consumers and businesses from privacy and data security incidents. Second, explore frameworks for how we might approach quantitatively measuring such injuries and estimate the risk of their occurrence. And third, better understand how consumers and businesses weigh these injuries and risks when evaluating the trade-offs to sharing, collecting, storing, and using information.

My ultimate goal is to use what we learn today to guide FTC case selection and policy work going forward. Our discussions will explore the types of negative outcomes that arise from the unauthorized access or misuse of consumer data and consider factors in assessing consumers' informational injury. We will also

8

examine business and consumer perspectives on the benefits, costs, and risks of collecting and sharing consumer information. And, finally, we will grapple with how to measure informational injuries.

But before we get to these important discussions, I would like to touch briefly on the FTC's role in privacy and data security and talk about why this workshop is both timely and important. Now, as I never tire of saying, the FTC is the primary U.S. enforcer of commercial privacy and data security obligations. We take that charge very seriously. We have brought more than 500 privacy- and data security-related cases, both online and off, and under my leadership, this work has continued. Since this summer, we have announced six important privacy or data security cases, Uber, TaxSlayer, Lenovo, and three cases enforcing obligations under the EU-US Privacy Shield Agreement.

Our primary privacy and data security tool is case-by-case enforcement under Section 5 of the FTC Act to protect consumers from deceptive or unfair acts or practices. One significant benefit of this approach is that it limits the need for policy makers to predict future developments in the marketplace. This is especially important in complex, fast-changing technology industries and in areas such as privacy,

9

where consumers have a wide range of evolving expectations and preferences.

Case-by-case enforcement focuses on real world facts and specifically alleged behaviors and injuries. Each case integrates feedback on earlier cases from consumers, industry advocates, and, importantly, the courts. This ongoing process recognizes that markets, consumer expectations, and consumer benefits and risks evolve with new technologies, and it protects consumers while allowing innovation to occur.

In addition to Section 5, we enforce rules under specific statutes, such as the Gramm-Leach-Bliley Act and the Children's Online Privacy Protect Act, and we offer copious amounts of consumer and business education on these topics.

Now, given the FTC's past record, you may ask, why hold a workshop on informational injury now? I've chosen to focus on consumer informational injury for two key reasons. First, in making policy determination, injury matters. Although the free market is a powerful institution for improving human welfare, consumers can and do suffer injury from some business practices. Government does the most good with the fewest unintended side effects when it focuses on addressing actual or likely substantial consumer injury instead of expending

10

resources to prevent trivial or purely hypothetical injuries. We need to understand consumer injury to weigh effectively the benefits of intervention against its inevitable costs.

Tom Leonard, in his comment for the Technology Policy Institute, argued this point quite nicely, noting that privacy benefits us because it reduces harms from information misuse, but if there are no harms, then data use restrictions impose only costs and no benefits. Policymakers and enforcers, therefore, need to understand how and how much consumers are injured by various practices involving the collection, use, and disclosure of consumers' information. More precisely, we need a framework for principled and consistent analysis of consumer injury in the context of specific privacy and data security incidents.

The FTC's deception and unfairness statements provide such frameworks for thinking about consumer injury generally, but it's worth exploring more deeply how those frameworks apply in the specific setting of privacy and data security.

Now, speaking of unfairness, the second reason I've chosen to focus on consumer injury is because it is a key part of our Section 5 unfairness standard. The focus of our discussion today is on defining consumer

11

informational injury as a descriptive and economic matter, but I hope that what we learn today can help guide the future application of the unfairness standard's substantial injury prong.

As for why hold the workshop now, as I've mentioned, the FTC has been a very active privacy and data security enforcer, and many of our cases appropriately focus on the most egregious low-hanging fruits where the harms were obvious to the affected consumers, the FTC, and often the defendants.

For example, we have a series of cases, such as LeapLab, Sequoia One, and the recent Blue Global case, that involve data providers who sold sensitive credit and payment information when they knew or should have known that the buying party was a fraudster who was going to misuse that information.

There were other types of cases involving direct financial loss for consumers. For example, the Wyndham data breaches allegedly resulted in identity theft and fraudulent charges to consumers. In the recent TaxSlayer case, the breached tax return information allegedly resulted in fraudulent tax filings, delaying consumers' receipt of their tax refunds.

But as technology and business models continue to evolve, we have and are likely to continue to face more

12

challenging scenarios that involve harms other than financial loss. For example, we took action against Accusearch for selling illegally obtained personal telephone records of individuals, where we had evidence that stalkers and abusive former spouses used this information to surveil and harass individuals. We also brought a case against the operator of a revenge porn website whose posting of highly sensitive intimate photos and personal information generated threats to and harassment of victims. And consider also the news reports of at least one suicide associated with the data breach at infidelity-promoting website Ashley Madison.

A strong framework for assessing consumer injury in such cases will serve two purposes. First, it will help us think critically as we monitor new technologies and data uses for potential consumer injury. Second, it will help establish criteria by which we can judge if privacy and data security enforcement is the proper tool to address a practice or if other mechanisms, perhaps either other agencies, institutions, or laws, would be better equipped to address any particular negative outcome.

I believe our discussion today will help ensure we have such a framework. First, we need to examine more thoroughly the range of injuries that can occur

from privacy and data security incidents. We are generally familiar with the direct financial injuries from identity theft, for example. We have also seen examples of unwarranted health and safety risks and intrusion into seclusion, and our first panel today will talk about the different kinds of injury suffered by consumers because of privacy incidents and data security breaches.

Second, we need to understand the key factors that matter in assessing injury from privacy and data security violations. Some obvious ones are the type of data involved, the magnitude of harm, and the distribution of the injury. But what else? And are the same factors relevant in both the privacy and data security context? What is the relationship between risk and injury? Finally, when is FTC intervention appropriate? Our second panel will tackle these issues.

Third, we can benefit from learning about how companies weigh the potential benefits and costs of collecting and using information and how this affects the decisions they make about protecting or restricting such information. Similarly, how do consumers weigh the benefits and costs of sharing information? And our third panel will dig into those issues.

Finally, we seek a better understanding of how to

quantify consumer informational injury. Now, there's an old saying often attributed to management expert Peter Drucker. "What gets measured gets managed." If we want to manage privacy and data security injuries, we need to be able to measure them. Our fourth panel will discuss the challenges of quantifying informational injury and how we can tackle those challenges.

At the end of the day, Andrew Stivers, Deputy Director of our Bureau of Economics, who has already done valuable work on these issues, will provide some closing remarks. This workshop is the next step in an ongoing conversation about consumer informational injury and how we can address it effectively, both here at the FTC and in the marketplace. This is going to be a fascinating discussion, and I again thank all of you for joining us, and I look forward to starting the first panel. So thank you very much.

(Applause.)

PANEL 1: INJURIES 101

MR. WOOD: Okay. To start off today's workshop, we will be exploring the broad array of negative outcomes that result from unauthorized access or misuse of consumers' personal information. We are fortunate to have a panel of experts today who can speak to a wide and varied range of injuries. We hope that this first conversation will provide concrete examples that later panels can draw from while discussing various aspects of informational injury.

We're planning to devote -- so we have five wonderful panelists that we are very happy can join us, and we will spend the majority of the panel hearing from them. We are planning to devote the last ten minutes to audience questions. So if you would like to ask a question, you need to find a question card. They are outside the auditorium or you can also raise your hand and I believe paralegals have them. Then you pass the question card to a paralegal, and they will bring it up here.

So with that, let me introduce the panelists. So Pamela Dixon is the founder and executive director of the World Privacy Forum, a public interest research group known and respected for consumer and data privacy research.

David McCoy is an assistant professor in the Computer Science Department at NYU Tandon School of Engineering. His research interests are in the areas of security, privacy, and empirical measurement. Some of his current interests span from the socioeconomics of cyber crime to automotive computer systems.

Lauren Smith is policy counsel at the Future of Privacy Forum where she focused on big data, the internet of things as related to connected cars, data ethics, algorithmic decision-making, and drones.

Cindy Southworth is the executive vice president of the National Network to End Domestic Violence. She founded the Safety Net Project which focuses on the intersection of technology and intimate partner abuse.

Finally, Heather Wydra is a supervising attorney at Whitman-Walker Health's Legal Services Program. Her practice areas include discrimination in employment by places of public accommodation and in healthcare, as well as representing clients who have been denied access to health insurance coverage or disability benefits.

Now that we have those quick introductions done, let's get this panel and the workshop itself started by discussing the various types of informational injuries and consumer harm that our panelists have seen.

MS. CONNOR: Thanks, Dan.

17

My name is Jacqueline Connor, and I'm an attorney with the Division of Privacy and Identity Protection, and so today we are going to start the panel off by asking each panelist one question and kind of giving them some time to answer, and then we will jump into a more general group discussion when the panelists can jump in whenever they want to.

So, Pam, we are going to start with you, and you have the clicker for the slides. Unfortunately, today, identity theft as a term has been part of our lexicon, and the harm to consumers goes beyond what we consider traditional identity theft. Can you describe some of those other different types of identity theft?

MS. DIXON: First, thank you for the invitation to share my research and knowledge today, and I am really grateful that the FTC is holding this workshop. I think it's good timing and a good topic, so thank you.

So I think everyone is familiar with various financial forms of identity theft. We have probably all experienced an incident where we get a phone call from our financial services company, and they say, oh, by the way, we are going to issue you a new card because, you know, someone is using your card.

So that's a lot different, that annoyance is a lot different, than, for example, the woman I met and

18

worked with who was from Utah, who had her children taken away from her through the actions of a medical identity thief. In her particular situation, what happened is that an imposter had taken her identity information, just gleaned from a simple phone book call, and this woman went around to emergency rooms around Utah seeking painkiller drugs. And the police came, and when they arrested this person who was a problem, they came to the victim's house, arrested her, and took her kids away from her, because she was a bad mom for being a drug-seeking behavior person.

So it took her three months and a DNA test and working directly with the state attorney general to get cleared and to get her kids back. So medical identity theft poses extraordinary harms to its victims, and Acting Chairman Ohlhausen discussed quantifiable risks.

We released a report today called The Geography of Medical Identity Theft where we worked very hard to quantify the patterns and distribution of medical forms of identity theft. When medical identity theft happens, it usually happens by the actions of organized crime or various organized professionals who are working within medical and billing systems to create false billing situations. Sometimes it's the action of rogue individuals, like the woman who had a problem in Utah,

19

but no matter how it happens, people who have had their identity used by others to procure or bill for medical goods and services that they themselves did not seek, nor receive, have unique harms.

The first core harm that they have is that fictitious entries are entered into their medical file. Typically, it's a very expensive disease, for example, HIV/AIDS, sometimes cancer treatments. A popular thing to add to victims' files is hepatitis C treatment because it can run up to about a thousand dollars a pill, leading up to about $120,000 that people committing the crime can pocket for themselves. Meanwhile, the victim is left with hepatitis C on their medical file, not a great thing to have there for all sorts of reasons.

So what on earth do these victims do to get cured? Well, unfortunately, there's no legal recourse to correct or delete medical information from files. Our geographic research that we published today shows that medical identity theft is, in fact, growing, but it's also doing something unique. It's growing in very specific regions of the U.S., for example, Florida and other areas where there are elderly residents that are clustered. And what is happening is that this crime is victimizing certain states, certain genders, certain

20

ages, the very young, the very old, and what's happening is that you're seeing real pockets of quantifiable harm.

So let me move to the slides very, very quickly, and let's see if this works. Oh, it's going to work, okay. (Slide 6) So we did a very substantive statistical analysis, culling through almost -- well, just loads and loads of complaints to the Consumer Financial Protection Bureau.

Now, these particular complaints are just the simple count of reports. So, in other words, these are the pure counts. If you look at this diagram, you can see California, because they are a populuos state, they have got a huge roster of counts. The same with Texas and the same with Florida. You can see Georgia popping up there a little bit and New York. This is -- this is showing activity based on population.

But when you normalize the data and adjust it to count per 1 million, all of a sudden the population density issue goes away, and what you get is this. So now you have a different picture emerging. (Slide 7) What you see here -- and by the way, the report has 2017 data -- what you see here is an absolute crisis. You see a very deep distribution of medical identity theft across the U.S.

You can see, for example, Nevada, Colorado,

21

Florida remains very high, and actually when you're able to look at the data closely, the entire Southeastern chunk of the United States is a hot spot for medical identity theft crimes and, in particular, Florida.

(Slide 8) So here is a simple count of reports from 2016. This is, again, just a pure count, and I chose this to show you just so you could get a feel for, like, when people talk about medical identity theft, oh, California has a lot of cases. You see a lot of DOJ activity in California because there are so many cases, but then if you normalize the data, that's what you get, and you get these profound patterns where serious harm is happening.

(Slide 9) People are having their health records falsified. People who -- in our research for this report, we found that people are literally unable to remove false billing records from debt collector scenarios. So these are profound harms. They are being targeted geographically, and consumers and patients who live in these hot zones are going to experience greater risk than those, for example, who live in North Dakota. And it's not just based on population; it's based on targeting.

So I want to quickly pivot to biometric identity theft. We are behind the eight ball in the realm of

22

medical identity theft. It's in crisis proportions. In biometric forms of identity theft, we will certainly wish that we will have -- would have dealt with it early. We still can, and so that's why I'm very keen to present this to you.

In biometric expert circles, there is something that's being talked about a lot, and we are unfortunately already seeing forms of this crime occurring. So what's happening is that healthcare providers are installing -- and others, schools and financial institutions and even federal agencies and other areas -- are installing biometric measures to ensure identity.

So it could be an iris scan, it can be a fingerprint, can be facial recognition, all sorts of things, but the way that biometrics work, it's math, it's not magic, and it can be spoofed, and it can be foiled. This is one of the most difficult types of identity theft to prove your innocence of.

(Slide 10) So what you're seeing right now, so on -- if you see subject one, subject one is -- let's just say that his identity information is stored in a hospital. Well, subject two comes along, and he's like, oh, you know what, I would really like some free healthcare, or I would at least like to do some fake

23

billing so I can sell some, you know, things on the street. So let me -- you know, this hospital has an enrollment -- a biometric enrollment system. Let me go ahead and, you know, take his, you know, driver's license scan that the hospital so helpfully has stored for me, and let me use free technology to morph my photo with his, and we will create that middle image that you see, morph one plus two.

Well, unfortunately, it is very unambiguous research at this point that subject one, with the original photo, and subject two, with the -- his original photo, based on that center morphed photo, both of those subjects will be authenticated within that healthcare system. So what we are seeing on the street now is clinics and other kind of bad actors taking biometrically morphed authentications and stringing a dozen people off of that authentication and passing people through. So this is a risk that has not made it into the mainstream yet, but we are watching it very closely.

So based on these things, I have just a couple of things I want to say in terms of harms and what to do. It is very clear that victims of medical identity theft and advanced forms of identity theft like that have experienced very different harms than victims of lesser

24

forms of identity theft, such as credit card fraud. They are in a class by themselves, and I call on the FTC to study and treat this form of identity theft as a completely distinct and separate class, because there are completely distinct and separate classes of harms, and Congress definitely needs to pass specific legislation in terms of working to create remedies for these situations.

It will require dedicated work from the FTC and other federal agencies, including CMS, working directly with state-level AGs, because of the patterning of medical identity theft. For biometric harms, it's going to take a lot of money, because these presentation attacks through biometric morphing, the only way you can fix it really is to have a secure line where you have a secure photo, not a Kinko's photo or a photo taken at an insta-passport place or somewhere else, but you have a secure photo that goes securely to whatever vendor is creating the -- you know, the final product.

These systems are easy to spoof also in other ways. So if someone has taken over your biometric identity, it's very hard to prove your innocence. These are distinct harms. We've got to study them, benchmark them, and work very diligently on their specific and unique harms. We can't lump identity theft as a

25

category together anymore.
MS. CONNOR: Thank you, Pam. If you can pass the clicker on to Damon.
So Damon, as we just heard about medical and biometric identity theft, those might be new topics to some, but it's fair to say that most people know at least a little bit about identity theft; however, there are still abuses of personal information out there that people may not know about but that are incredibly damaging to people. Can you explain what some of those are?
MR. MCCOY: Sure. I am going to go through a little bit of the research that I've been doing on this, kind of starting with what's called doxing, and so doxing is basically the public release of people's information that they probably wish to keep private. (Slide 11) And so for those of you that haven't stared at lots of these doxings, these doxings are posted a surprising amount of times on the internet, and these doxes normally include the names of the people, their online aliases, their age, their date of birth, their addresses, phone numbers, sometimes medical kinds of information about them, ISP information, and also a lot of information on the family of these victims. And it also includes things like their online social networking

26

profiles and things like this.
(Slide 12) And so these doxings are oftentimes posted with the hope that they're trying to encourage harassment of the victims of these doxings. And so in one of my studies, we actually built a system based on some machine learning that was able to go out and find about 6000 of these doxings. So we ran the system for about 12 weeks, and we found 6000 of these doxings, and probably the ones that we found were only the most egregious of these doxings since our system was fairly conservative about what it calls a dox. And so based on this, we could do some analysis of these doxes.
What we find is that the victims often come from one of two communities. They either come from the gaming community, video gamers, or they come from the community of kinds of hackers and underground kinds of actors. The other thing we can see from this is that the victims oftentimes skew very young, and so the victims are oftentimes in their teens, perhaps in their early twenties, something like this as well. So this is very targeted and impacting, not the people likely in this room, but kind of the younger generation of people. As you can see from the, you know, 6000 doxes that we found -- and this was not a comprehensive study -- that this is happening very frequently and impacting a lot of

27

people's lives.
(Slide 13) And so one of the probably most egregious harms that can happen from this is in those doxes, again, they include people's social networking profiles, and, again, this is done to kind of encourage other people to pile on and harass these people. So what you're looking at here is -- we actually did our study in two parts. So we did our study in two six-week parts. We found about the same number of doxes in each part of our study, and with these social networking handles, we pulled these out of the doxes automatically and we monitored the privacy settings of these people's accounts.
What you can see from this graph here on that first one is that red part is essentially people closing their accounts down, so this represents people becoming socially isolated, likely being forced by harassment to close their Facebook accounts because of these doxes. And so what you can see in that first period is that -- and the thickness of this is the magnitude of the number of accounts that are closing it down. And so that first part is before Facebook started deploying filtering to try and filter out harassing comments from their platform, and that second part is after Facebook deployed filters to filter out harassing comments from

28

these. And you can see the huge benefit that was incurred by their population when Facebook did this. This was actually kind of a nice thing. Then we can show that this kind of harm can perhaps be mitigated by these filters that these online social networking sites are deploying.
(Slide 14) The other kinds of harm, probably the much more dangerous kinds of harms that can come from these is, again, the phone numbers, addresses, are also included in these doxes, and so this can lead from fairly innocuous things from someone ordering a pizza to someone's house to someone creating kind of a fake emergency situation, say, like, they make up a situation where it's a hostage situation or something like that, and they call up, you know, 911, and they say there's a hostage situation going on in this house.
This triggers a response by the Police Department to send, you know, armored swat teams to someone's house with guns blazing, on to someone's house, and this has resulted in many dangerous kinds of situations. Unfortunately, this kind of harm is more difficult to try and measure, but I think this is something that we should do a better job, from police databases and things like that, to measure this type of harm that's incurred by doxes.

29

And so the other question is, right, where are they getting this information from? And so one thing that we can notice from these doxes is that a lot of times they have ISP information, and a lot of times there's also online cookbooks as to how to gain information to dox people. A lot of times these include a method where they essentially socially engineer call support centers, especially ISP call support centers.

So they can start with someone's IP address that they somehow gained maybe by playing a video game with them, and they can spider out and start to get more information. So this is another thing where we should investigate how they are getting this information and perhaps understand how to make it more difficult for them to acquire information this easily.

(Slide 15) The other way that they acquire information is by sometimes socially engineering their victims into installing malware, and it's a particular kind of malware that they normally have installed on the victim's machine that's called a remote access trojan, or RAT for short.

So there's unfortunately underground ecosystems where the doxes are posted. So there's other hacker forums where they basically trade or talk about how to acquire victims, which they called slaves, to infect

30

with these remote access trojans. And when a victim is infected with one of these remote access trojans, this effectively gives the hacker complete access to that person's computer, including the camera, the microphone, all of the files on the system, the attacker can rummage through.

(Slide 16) So oftentimes what crops up from this is what we called sextortion, where the attacker somehow steals some sensitive information from the victims and tries to extort them for more images or perhaps for money and things like this. So a high-profile victim of this was Miss Teen USA. She was infected by one of these RATs, and she was caught up in one of these sextortion cases, but there are hundreds and hundreds of these victims.

So, in fact, we did a study where we enticed attackers to attack our own system, and we had over hundreds of attackers, unique attackers, that visited our systems, many of them from the U.S., but from all over the world, and the first thing that most of the attackers did is they tried to access our camera when they attacked our systems.

(Slide 17) And so the final thing that we have been looking at lately is spyware. I think actually spyware is kind of a misnomer for this. I oftentimes

31

called it stalkerware or something like this. So these are mobile apps that can be installed on people's phones. Some of these are in, say, the Play Store and Apple's App Store, and they have kind of dual-use features.

So they can be used, for instance, to find someone's stolen phone or they can be used to kind of -- without the victim's consent, to stalk them, and things like that. And then there's stuff that's off of these stores that's much more egregious. And so a lot of these developers, as you can see from this marketing material that I have posted up here, are essentially encouraging people to use it for this use case of stalking their partners, and so this can lead to a lot of different kinds of harms. Probably the worst of these harms is interpartner violence, so this can be used to facilitate intimate partner violence. It can also be used, again, to stalk people and things like that.

And so doing just quick internet searches, you can find, you know, tens to hundreds of these different apps that abusers are talking about installing on their victims' phones and things like that. And, again, we can see paid advertising where the developers of these kinds of stalkerware apps are encouraging and

32

facilitating this use of their apps.

And so throughout all of these studies, I can kind of show that oftentimes there are cyber harms that are hard to kind of boil down to dollars. They are, you know, physical threats and things like this, emotional harm, social isolation, that are very hard to kind of boil down into dollars, and they affect oftentimes young people and things like that, so the newer generation, and they probably have a lifetime of impact against these people that are harmed by these kinds of cyber attacks.

And so with that I would kind of encourage more of these kinds of studies to kind of understand the scale and other kinds of harms that are occurring with this and also the intermediaries that are kind of facilitating these kinds of harms. So with that I will hand it off.

MS. CONNOR: Thank you, Damon.

So, Lauren, so far we have discussed injuries that occur to particular consumers, but what kind of informational injuries are possible when many consumers' information is combined?

MS. SMITH: Great. Well, thank you for having me here to speak today. I'm excited to participate in the panel. So, yes, I am going to take a little bit of a

33

broader view, to take a step back, and think about the fact that as the volume of consumer data grows, the number of decisions that were previously made by humans are now increasingly being made by algorithms. As the number of data sources have multiplied, so, too, have the types of data involved in these decisions and the number of entities that may process and analyze this data across many sectors.

So I am going to be talking about potential harms that come not just from explicitly fraudulent activities or hostile approaches that we have heard from some of the other panelists today, but instead look at, you know, areas where the collection of the data initially may have been legitimate uses, may have been offered up by consumers in exchange for a service.

As you mentioned, this can create a number of sort of downstream effects. So information can be combined, so a consumer may provide some information in some context, but in what is sometimes called the mosaic effect, the information could be combined with other data sets and perhaps reveal information that they may not wish to have out there publicly, and also there can be inferences based on data that consumers have shared, you know, particularly with the advance of artificial intelligence, that there may be information that could

34

be inferred about a person from data about themselves that they may not have explicitly shared.

So in a lot of instances, analysis of sensitive data categories, such as race, gender, pregnancy status, can be used to improve services, promote inclusion, advance research, and actually be used to mitigate human bias, but in other contexts, it can raise the specter of disparate impacts on vulnerable communities.

As a number of folks have sat down to try to tackle this issue, I know my initial work on this started when the White House did a report on big data, looking at the benefits and potential risks arising from big data, and one of the biggest areas that was sort of a surprise to the folks working on this project and in the building was the intersection of big data and civil rights and the way in which this growth in automated decision-making could wind up impacting civil rights. The FTC has also done a lot of good work and hosted a workshop on that topic explicitly.

So as my organization sat down to try to tackle this issue, we found that conversations on the topic sometimes become mired in definitional challenges that can make progress toward solutions difficult. So these analyses sometimes fail to separate harms from the causes and from solutions. Sometimes we wind up

35

conflating human biases that have shaped society for years with digital causes. This can make it difficult for a consumer protection agency, like the FTC, to determine what kind of injury they should be protecting consumers from and how to define that.

So, you know, identifying financial injury around fraud, breach, or security failures may be fairly well defined, as we've seen in several FTC enforcement actions, but in other contexts, identifying the injury may be more complex. We can run into issues like this with things like differential pricing, which has given us things like senior and student discounts for years without, you know, creating what we might consider to be a financial injury, but can raise new concerns as consumer purchasing moves to the internet and pricing can shift based on a range of consumer attributes.

So we've found that there aren't many easy ways to navigate these issues, but we think that more can be done to promote fairness and encourage responsible data use. So to facilitate these conversations, we endeavored on a project that launched today, that I'll give you a brief preview of, and there are handouts also at the front. (Slide 18) We reviewed the leading books and articles on this topic and tried to distill the different harms that are identified in the literature

36

into a set of buckets and into a chart so that we could think about them at a very high level and try to understand what are the types of harms that could arise from automated decision-making based on consumer data, so that we could have them all in one place and determine whether some of them should be considered more of a risk, more important to address and focus on than others, and whether mitigation strategies might differ depending on which types of harms you're talking about.

So you can see a bit on this chart up on the screen, and so the first distinction that we made in groups of harms was the types of algorithmic decision-making that could either turbocharge or make more opaque harms to individuals, which you will see is the left two-thirds of the chart. And on the right is collective harms, which may be a little more difficult to pinpoint because there may not be, you know, specific harm to a particular person but may impact us at the group level, either as a society as a collective or specific groups.

And for the first, under individual harms, we drew a distinction between those that are unfair and those that are illegal, because we felt that there is sort of more clarity in especially civil rights law, as well as, you know, FCRA and some other areas where

37

specific harms to an individual could be -- have already been identified as illegal, as there is sort of a clear societal consensus that we do not want these harms to occur and can use technology to ferret out and to prevent some of those harms.

And then the ones that are not as explicitly illegal can raise questions of unfairness and of ethics but may require us to do a little more thinking to determine, you know, in what instances this would be considered an injury to consumers and how one might want to think about identifying and mitigating and being aware of some of those concerns that might arise with automated decision-making overall.

So when we got to the substantive grouping of the harms, we found that, by and large, they could be grouped into four broad buckets. So the buckets are loss of opportunity, economic loss, social detriment, and loss of liberty. We thought that these depicted the spheres of life where automated decision-making could potentially cause the most harm.

In any of these buckets, as I mentioned, the harm could be illegal, it could be not illegal; it could accrue wholly to the individual or to a group. So, you know, an example of an individual harm could be me not being able to rent an apartment because of a decision

38

made about me or something entirely to society, which could be something like filter bubbles that have become a topic of discussion, especially around the election, that many of us may believe have a harmful impact on society overall but have not been distilled as clearly, at least in existing law, and may not be something that can be tackled at that level at this point.

So for loss of opportunity harms, this group broadly describes harms that occur within domains of workplace, housing, social support systems, healthcare and education. Economic loss harms broadly includes harms that cause financial injury or discrimination in the marketplace for goods and services. Social detriment largely talks about harms to one's sense of self, self worth, or community standing relative to others. And loss of liberty applies to harms that constrain one's physical freedom and autonomy.

So these are up on the screen. In our packet, we give particular examples of each of these, and I'm happy to run through those individually. You know, we recognize that these categories that I just ran through aren't mutually exclusive, you know, economic loss can certainly lead to loss of opportunity, but in our attempts to sort of survey the literature and boil these down, we found these to be the most ready categories.

39

And then in an accompanying project, we also used these categories to then try to bucket the harms into different groups that can be approached with the same sets of mitigation strategies depending on the way they're characterized in this chart.

MS. CONNOR: Thank you, Lauren.

Cindy, given your experience with domestic violence victims, what types of informational injuries do you see arising and how?

MS. SOUTHWORTH: I just wanted to take a moment to go back to the "Hello Spy" slide that Damon provided, because it gives me great pleasure to trash them at every chance I get, because we proudly tweet at them. They are foul human beings. This is how they advertise an alleged family safety product. It's all about encouraging and facilitating stalking and domestic violence. So I last trashed them in 2014 at a Senate hearing, but I appreciate any opportunity to go after miscreants like this.

On a more serious note, just a little bit about domestic violence, stalking, and sexual violence. Many people think of them as these, like, tiny societal issues, this fringe thing, and when you actually look at CDC, FBI numbers, Department of Justice research, one in four women will be physically assaulted -- physically

40

assaulted -- by an intimate partner at some point in her lifetime. That doesn't count emotional abuse and control and lots of other unhealthy and controlling relationships. One in six women will be stalked by a partner or stalked, and one in 19 men will be stalked in their lifetime. And then one in three girls will be sexually abused before they're 18 and one in six boys.

Those are staggering numbers. That means your CEO, your boss, your neighbor, your friend, your postal worker, your, you know, barista at the Starbucks, your next politician is a survivor of one of these crimes. And so when you think about this large swath of society who's experienced significant victimization, then you think about how they're part of our daily lives, they're part of our data sets, they're part of our medical records, they're part of our lives.

And so the Safety Net Technology project that I founded 17 years ago on our blog, techsafety.org, we did a survey back in 2014, and we found that the local domestic violence organizations -- there's about 2000 of them nationwide that take those hotline calls from victims who are experiencing pretty horrific things -- and back in 2014, 97 percent of those front line victim service organizations said they were seeing cases where tech was being misused by offenders and abusers. Of

course, because tech is part of all of our lives. We all have our hands right next to a phone.

And three years ago, in that same survey, we found that 71 percent, almost three-quarters of those local programs, said that abusers and stalkers had started using technology to monitor internet and computer use, so essentially spyware. We don't have prevalence or incidence stats. This is just what the victim advocates are seeing on the ground, and the numbers are really quite concerning.

When you think about some of the specific harms, the obvious one is, you know, illustrated above. It's the physical harm. People running for their lives still happens, you know, domestic violence shelters end up housing or supporting 72,000 adults and victims in a given 24-hour period, but 41,000 adults and kids are kept safe in a shelter or a transitional housing program on a given day, 41,000 people. So they have significant privacy needs, and they are obviously quite concerned about data privacy and data security.

Some of the other things that can happen that you need to think about are credit issues. We find one of the most common tactics by abusers is to ruin a victim's credit, and so then if you have got commingled credit, if the abuser's criminal records get commingled with her

records in some way, or his records if he's the victim, those are really challenging things that can have devastating consequences.

You have job loss that you can have if you are identified as a victim. Unfortunately, there's still a huge stigma. Some people will say it's not safe for you to work here, you might be a risk to your colleagues, so you see all sorts of discrimination type things that Heather is going to be touching on more.

Other things to think about is just how data in combination can be problematic. And so years ago AOL released some of their search data for research purposes and didn't realize how identifying it was. A reporter actually took that data set, realized that somebody was searching for themselves by name and also searching for things like domestic violence, shelter, protection orders.

So the reporter called the victim up and said, hey, I got your information from the AOL search data set. Are you a victim of domestic violence? And, of course, she was. So how sort of crazy that would be to have a reporter call up and say, hey, I figured you out. Are you a victim? And given the stigma that is still out there, that is incredibly concerning.

Some of the other things to think about is not

just physical harm, but some of the data that victims are most concerned about is their home address. And so when medical records started becoming a real issue, one of the things people were sort of caught up on is the sensitivity of what's in a medical record, diagnoses, medical treatment, and, of course, that's true.

But for a victim of domestic violence or stalking who's literally on the run or living in hiding or trying to rebuild a life, just the home address being compromised, shared, you know, in any way, either through legitimate or illegitimate means, that could put a victim's life in danger, and then they may have to move.

And so when you talk about cost, it's, you know, physical cost of relocation, of plane tickets, bus tickets, movers, alarm systems, changing door locks, slashed tires. So those are the types of things that we see when there's privacy and security breaches.

MS. CONNOR: Thank you, Cindy.

And last but not least, Heather. Given your experiences, what kind of injuries result from the disclosure of medical information, sexual orientation, or gender identity that a consumer prefers to keep private?

And then to go beyond that, what are the effects

when people fear that this information, this medical information, sexual orientation, or gender identity, may be disclosed?

MS. WYDRA: Yeah. Thanks very much for having me here. I am happy to talk about these issues. I am an attorney at Whitman-Walker Health. We are a medical-legal partnership where we provide medical services to anyone living with HIV in the greater D.C. community and anybody identifying as LGBTQ, again, in the greater D.C. community.

So the reason that brings me here to this panel is because a large part of what I do in the legal services program is to counsel and represent people who have been affected by a disclosure of personal information, a disclosure of, for example, personal health information, private health information, or a disclosure of gender identity and sexual orientation that the person did not want to be disclosed.

And I see people harmed by these disclosures -- well, in countless contexts, but the three I'll talk about today are in the workplace, places of public accommodation, and then in an interference with personal relationships, and that can be in the home, in the community, in the workplace, and elsewhere.

So to begin to talk about the workplace, a large

45

part of what I do is employment discrimination, and a lot of times discrimination starts when somebody's health information -- in a lot of cases with the clients that I work with it's HIV, it could be something else -- is disclosed in the workplace. And sometimes the disclosure is necessary, sometimes it isn't, but regardless, the results are not always positive.

And the negative effects can be seen in every aspect, at every stage of the employment relationship. I have had clients who have come to me who have gotten through the initial portions of a job application, and then as you get a little bit further in the process, when you have what's called a conditional offer of employment, the employer is then allowed to ask for health information, can ask about HIV status, and can even make people undergo an HIV test as long as everyone who is applying for that position is required to do the same thing. An employer can't pick and choose.

And then, believe it or not, it still happens, after that testing occurs, after somebody's status is public, they find themselves suddenly without a job offer. That's illegal. Of course, it's illegal, but sometimes it still happens. Then it becomes my job to prove that the withdrawal of the offer was based on the HIV and not for some other purpose. That's a blatant

46

example.

Another one that can be more subtle but still happens is I worked with a client who was applying to be a bus driver in one of our nearby counties and, again, had to fill out a huge medical questionnaire. It didn't actually ask about HIV, although it could have, that would have been permissible, but she ended up disclosing it just because she had some blood work that was coming back oddly and thought, well, maybe if I just say this, this will end all of the inquisition.

But the doctor, instead of saying, okay, now I understand, she had letters from her personal physicians clearing her to do this job, said, well, I want you to go back and get these three tests, and I need all of these records. And eventually my client said, you know, they don't want me here, and I don't want this job. They don't want me here. So that's another example.

And then the example that I really see the most is someone who's going about their day, doing their job. I worked recently with somebody who's a maintenance worker in an apartment complex, and because he had to take so much time off for doctors' appointments, he finally said to his supervisor, you know, I'm positive. I've got to go to the doctor. I've got to get blood work. I have to have these medication checks.

47

Well, he didn't get fired, but all of a sudden, nobody wanted to work with him anymore, and he was put on these assignments by himself, and people said strange things to him, like, shouldn't you be wearing gloves? And day after day, there was just some little indignity that he had to endure. So, you know, in that kind of thing, I mean, yeah, there's a hostile work environment claim that we can bring, but, you know, it can be hard to find a remedy for that, and it's just something that he had to endure.

So the next context that I'll talk about where I see harm occurring is in what we call places of public accommodation. So that is any place that people go, you know, restaurants, hotels, gyms. The example I am going to talk about is the gym, and I am going to use the example of somebody who was transgender.

I worked recently with a client who was going to her gym, having absolutely no problem. She was a transgender female. So just to give you your Trans 101 training, that means she was born -- she was identified as male at birth but had been living outwardly as female for years, but she hadn't changed her driver's license. So she was going to the gym, using the women's locker room, everything was fine, but then one day she didn't have her ID card and had to show her driver's license,

48

which was old, had an old picture, had an old name, and after that, everything changed. The manager said she couldn't use the women's locker room anymore. So we're currently dealing with that case.

As some of you may know, in D.C., the law is very clear. It doesn't matter whether identity documents are formally changed or not. If somebody identifies as male or female, that is the gender where they need to be treated under the law.

And so, finally, I'll talk about the harm that can come in personal relationships when personal information is disclosed. There was an incident that happened with Aetna. It was in the news. It was public. Aetna sent a mailing to dozens of its HIV-positive clientele who were -- they were making some change to how they were covering HIV medications, and the mailing that went out had a very large window that shows the address, you know, where you can see what the address is. The window was so big that it actually showed the first couple sentences of the letter, which talked about HIV and medication.

Well, these letters went to people at their homes. They went to people at their apartment complexes, where, you know, mail is just thrown everywhere by the front door. A lot of people ended up

49

with their personal health information disclosed to family members, neighbors, friends, when they really didn't want it to.

Sometimes that was okay, but we heard horrible stories, people who came to us wondering what they could do, because -- I mean, again, it ran the gamut from, you know, people just look at me weird now to somebody wrote, you know, bad words across my apartment door and vandalized my garden and left me notes saying we don't want your kind here.

So those are the types of things that can happen to people when their personal information is disclosed without their permission.

MR. WOOD: Okay. Well, thank you very much, Heather, and thank you to all the panelists for describing to us a very wide range of injuries. I hesitate to ask, but are there other types of informational injuries that we haven't touched on yet?

MS. DIXON: Well, if no one is going to respond, I will. There are certainly more types of informational injuries, and they are hiding in every corner of the digital ecosystem. I think, though, that the important thing is to focus on what causes substantive harm and harm that has meaningful impacts on a person's life.

You know, we can all spend time working on

50

issues, but there are big issues. I do think that medical forms of identity theft, I think biometric harms, I think these are big issues. I think the issues we've heard around this table are big issues. If someone is going to have their life threatened or their livelihood threatened, these are profound harms.

So in some ways, I mean, it's -- I would really rather look at quality and say, look, here are very meaningful harms that we have quantified, we've studied, we know about, so let's roll up our sleeves and let's do something about them. We've had plenty of time to identify these harms. Why not have the FTC write a new report, for example, about domestic violence, about medical identity theft, about these other harms? Hold a separate workshop. Let's find solutions. Let's work collaboratively. I'm all for it. Instead of, you know, breadth, let's go depth, and let's solve the problems.

MS. SMITH: I would say one area of harms that we haven't provided examples on yet up here are the loss of liberty harms that can have a very significant impact on folks' lives. So there has been a lot of good research on predictive policing and the fact that relying on data to determine how policing resources are targeted can reproduce -- if it's not done with an understanding of the risks, can reproduce historical bias and have very

51

significant impacts on communities at large.

I think those are, you know, things we're just beginning to understand and figuring out how to mitigate, and things like use of recidivism scores, making sure we're understanding the type of data that is going in, assessing whether that is the correct data, assessing whether that data has its own biases built into it when we're making decisions about folks' literal freedoms.

MS. SOUTHWORTH: Just on that note around sort of what police data can do, police data can do a lot of mitigation of harms in terms of helping identify where you have got either overpolicing or underpolicing or institutional bias dilemmas. One of the downsides of police data is, inadvertently, it can actually be identifying.

So there was a 12-year-old rape victim that was -- her identity was basically compromised by a very well-intentioned police department publishing their police data, and they used the block to try to anonymize the data, but there was only one adolescent girl on that block. So even, you know, the attempts to an anonymize the data just -- it wasn't anonymized, and so in that case a pretty traumatic life experience then became amplified by the whole world finding out about it.

52

MS. CONNOR: So we have been talking a bit about the actual injuries that come after the fact, but what about -- and we have heard a little bit about the risks of injuries, but I guess for Heather and Cindy, I'm wondering if the people that you work with, if you see them acting differently or, you know, are they not -- are they afraid to give and use their information for positive purposes to help them because they are afraid it's going to be misused in some way?

MS. WYDRA: Yeah, I can start with that, and, Cindy, I'm sure you will have plenty to add. Absolutely. Where I see that a lot is in the area of healthcare, especially with people who have experienced discrimination someplace else, will be afraid to go and see a doctor because they're worried that either a doctor won't want to treat them because they are HIV-positive -- and, yes, that still does happen -- or someone who is transgender feels, you know, like a doctor won't know what to do with me, they will be biased against me.

There are plenty of instances of doctors being very culturally insensitive to someone who is transgender and not understanding at all what their particular medical needs are or even why someone would be transgender. So, yes, that is one area where I see

53

people then afraid to go for healthcare outside of Whitman-Walker.

MS. SOUTHWORTH: There's a real conundrum. I know some of the attempts for data security have been to mail things to people's homes. That came up during some of the pretexting of phone records dilemmas, and from a domestic violence victim standpoint, that's not necessarily helpful.

The same sort of scenario of somebody being afraid to go get healthcare after a victim is relocated to a different town because the health insurance will send notification of the care that was provided and the location of where it was provided to the original home address, and it's really hard as it is to safety plan with victims of domestic violence and stalking who are trying to relocate and sort of talk about make sure you grab the birth certificates, you know, get the favorite teddy bear, make sure you have got the bank records and all this, and oh, by the way, you need to notify the insurance company, this data broker, all these other places, and change your address.

People are a little busy running for their lives to do all of those steps, so what on the one hand is an attempt to be providing transparent notification to the sort of address on record, it gets very challenging in

54

terms of other places where victims sort of opt out of society, sort of what Damon was showing in terms of after people are targeted, do they shut down their Facebook accounts, do they opt out of social media, do they choose to sort of exclude themselves, or are they told explicitly?

We have had many survivors told by law enforcement, if you don't want to be threatened by your ex, get off Twitter or get off Facebook, which is isolating, especially if your entire family -- one of the ways that my now family communicates is I know who's expecting a baby because I'm on Facebook with all of my cousins and relatives.

Telling me the only way I can be safe and be female in the world is to get off the internet because, well, this misogyny just happens, so just get over it and get offline, that's not possible and it's isolating. It also impacts job possibilities. If I can't safely use technology, I am far less likely to go get a well-paying job.

MS. DIXON: I would like to add on to that very briefly. The other thing I see a lot in my work is that people don't understand what it is that will actually protect them meaningfully, as opposed to things that they just feel like will protect them. A really good

55

example of this are people who have had some kind of exposure that's problematic, and in order to reduce the information flows, I often recommend to them to actually opt out of certain websites. They won't do it because they don't want to give any additional information.

So I've actually had people who refused to opt out of, for example, financial information sharing because they don't want to give their Social Security number. People are so afraid of giving their Social Security number, when in that instance it would actually help them.

So I do think that there needs to be a lot better understanding of, okay, what piece of information can I give up that will actually help me? What pieces of information will genuinely be harmful? I'm not persuaded that, you know, getting off of all technology is the right answer here.

I would really like to take the position that, look, let's be able to use technology and have it not harm us. That's really the situation I want to see people get to.

MS. CONNOR: Thank you.

So moving on a little bit, what are the costs of fixing these informational injuries that you've all been talking about, financial or otherwise? And I know

56

Heather was speaking to this a bit, about legal remedies, but what legal remedies are available to consumers to address these harms? And if there's not any, what gaps are there and what are their consequences? I know that's a loaded question, but...

MS. WYDRA: Well, I can start talking about, you know, the clients that I work with and the cases that I see. The remedies, if they exist, they come under the state, local, and federal antidiscrimination laws. So that could be Title VII, that could be the Americans with Disabilities Act, and then the state counterparts, the D.C. Human Rights Act, the Maryland Civil Rights Act, and so forth. So that's one way for people to get a remedy.

For the disclosure of medical information by a healthcare provider, there is HIPAA, but for those of you who know anything about HIPAA, it doesn't provide a remedy to the individual who has been harmed by the disclosure. It can provide a penalty and sanctions and other types of punishments to the medical provider or the insurance company, but the person who has been harmed, there is no remedy for them.

And then, of course, you always -- there are some things that there is -- just no amount of money compensates for. I mean, you try, but loss of

57

reputation, emotional harm and suffering -- I mean, certainly, yes, we try and fix those things with financial remedies, but there are some things that can just never go back to the way they were before.

MS. SMITH: So I think, you know, part of why we undertook this exercise is to understand that, depending on the type of harm, it may have a different set of remedies. So for some of the more well defined harms, such as employment discrimination, so even if there is existing law saying that a certain group cannot be discriminated against in this context, technology can enable that to happen anyway, where data could be used as proxies.

So you may not be able to not hire someone based on their race, but if you know what type of shampoo they use based on their online purchasing history or browsing history, that shampoo could be used as a proxy for race. So I think with some of these we need to think about different methods to ensure that we are making sure that data are not being used as proxies for protected classes. I think that is a very clear top-level approach to preventing some of these harms, but ensuring that -- you know, that is sort of built into how we're thinking about these issues, writ large, and that there are technological solutions to some of these, you know,

58

algorithmic designs to consider whether protected status should be used as an input in certain contexts.

And sometimes it may be helpful to include protected class status so that it could be checked to ensure that decisions are not being made based on that, but it also may be important to consider how it's being used throughout the process.

And then when you get down to sort of less crisply defined harms, you know, it's important to think about -- you know, in some of the conversations recently, you know, are we calling on tech platforms to define societal norms? And is that sort of too much to ask as we're thinking about these issues?

And considering that, you know, for things like network bubbles and narrowing of choice, we haven't created a clear set of societal norms yet, but if there are business processes that are taking these into account, that are considering the ways in which these products could have impacts like this, sort of creating ethical frameworks, creating best practices to monitor and check for ways in which data in a data set could be used to have some negative impacts, that that is really something that should be baked into how we're thinking about new technologies going forward.

MR. MCCOY: So I think one big area that we have

59

seen in society that's lacking is this online harassment stuff, right? In the doxes I showed you, it's hard to point at a law that currently exists that would outlaw that type of behavior. And as the FTC showed, they had to kind of creatively go after these revenge porn sites, and then finally the states are kind of stepping up, piecemeal, and enacting laws, but there's a huge gap in terms of understanding, right, the laws defining nonconsensual posting of people's information online.

MS. DIXON: I would like to focus on one of the harms that is very, very tied to medical forms of identity theft, which is very aggressive, unethical, and often illegal medical debt collection. So individuals who have written in complaints both to the FTC and the CFPB have routinely stated that they are working to get debt collections off of their credit bureau files for one, two, and three years, and there's narratives actually in the report we just put out of victims of medical identity theft who literally cannot buy homes and cannot move on with their life because they're being held hostage by debt collectors.

And some of the debt collectors will even tell them, you know, if you just pay us $200 on this debt, we will let you go. It's just a horrifying mechanism. And I do think it would not be a profoundly expensive item

60

to fix, to really take a closer look at medical collections and really clean up that niche of a sector. I think it's an area that causes a lot of harms to a lot of folks over extended periods of time.

MR. WOOD: Okay. Well, this has been incredibly interesting, and I think we could go on all day, but we promised that we would leave some time for audience questions. So we have two questions from the audience, and the first one I guess is addressed generally to the panel. They are both addressed generally to the panel.

So it is, can the panel speak to how to characterize the lower level of informational harms that exists and affect day-to-day life but don't rise to the level that is being discussed by the panel?

Anyone?

MS. DIXON: I wish we could ask a followup question more about what they meant. You know, information annoyances beleaguer all of us, and it really does take some time to figure out, okay, is this annoying or is this a problem? Some of the doxing and the spyware, I mean, these are large problems, but they may start as smaller ones. I think that this is a difficult question without knowing more from the asker. I apologize.

MR. MCCOY: And you can clearly point at the

61

target advertisements, things like that, ads that are showing up on people's pages that are clearly revealing other things they are searching, and if people can see their screens, they can quickly infer other things that people are searching that are, you know, sometimes or oftentimes very private.

MS. DIXON: You know what, I did think of something that actually I think is a low -- well, not a low level, but it's an everyday informational issue. Most people who call our office are interested in stepping up their privacy game, and one of the first things I ask them is how they're using their financial tools and services, and really, if you want to really see a lot of privacy informational issues that are on a lower level, not necessarily hurting you but just circulating about you, really you have to learn how to, you know, shut down some of the financial data flows.

It's easy to do, there's a lot of opt-outs, but I do think that's a general annoyance that kind of filters through all of our lives, you know, unless you, of course, want to pay cash for absolutely everything, and I don't know anyone really who's able to do that anymore, it's just not a very easy way to live.

MR. WOOD: Anybody else, or should we move on to the second question?

62

Okay, so the second question is, to what extent are the harms experienced in these events purely due to the actions of a miscreant versus the negligence of a data steward?

MS. SOUTHWORTH: I can jump in on this one. I think that is the million dollar question. You know, one of the -- I would say the easy bad actors for me are the ones like Remote Spy that were actually advertising to allow you to spy on anyone from anywhere. That's just an easy lift of, yeah, they're pretty much malintended.

I think where it's more challenging is if my information is compromised through a fairly innocuous, minor data breach, does that information identify something about me that then has later long-term impact? You know, do I not get hired because of my history of domestic violence or sexual assault, or whatever stigma it might be, because something that was accidentally leaked and then makes its way through the stratosphere?

I mean, I think the blatant bad actors are easy targets for us to point to, but I do think there are unintended consequences of breaches that at least my constituency sort of runs into.

MS. SMITH: And I think the reality is that personalization has become a service and something that

63

we are very used to seeing on online platforms and technology devices that we use, and there's a lot of benefit there to consumers, but it also sort of leads to downstream impacts.

And so, you know, just as there is more data built on an individual, if that information is used in a way that an individual does not know about and would not prefer, then that could be perceived as a harm to them, and it could have, you know, depending on the type of harm, could have either slight or significant impacts on their lives.

Again, I think some of these questions are not settled questions, you know, are -- is having a -- the set of information that is presented to you when you log into an online platform personalized to you based on what your friends like? based on what you've shopped for? That may provide you with a great service that you really enjoy, but it also, you know, could be used in other ways that you may not prefer.

And so I think some of these harms we are still just starting to understand, and as much, you know, transparency and control that consumers are able to have, I think the more we're able to sort of safeguard in the long run how they play out.

MR. WOOD: Anyone?

64

Well, I guess we have about a minute each for final thoughts from the panelists. Does anybody -- should we -- can we just start with Pamela?

MS. SOUTHWORTH: I want to just sort of -- not being a legal expert but instead talking about very pragmatic things, I loved the example that Damon used about how technology filtering can make such a difference to somebody's online experience and help mitigate harassment.

We've seen very pragmatic approaches to information harms in my 27 years of doing this work. One was a domestic violence shelter years ago that their address was accidentally published by the local phone company, and the phone company paid off their mortgage and helped them move, not a light lift, but very practical and very pragmatic.

I think the granular attempts to allow people to control their own experiences, such as filtering comments within Instagram, Facebook, Google is working on it, will help people be able to stay online and have less assault and harassment coming at them.

MR. WOOD: Okay. Well, thank you.

Other final thoughts?

MS. DIXON: I have a comment about solutions. Very often I hear people talking about solutions and,

65

oh, well, we are going to solve our medical identity theft problem by installing a biometric system to make sure all these patients are who they say they are. Meanwhile, they have got a huge security problem in their biometric system.

There is not anymore -- it is just not possible to find a single silver bullet, gorgeous, perfectly formed solution that will solve 100 percent of an informational problem. It doesn't exist. We have to look at layers. We have to look at different facets. We have got to have a multifactorial approach to any kind of problem that we're trying to solve.

And I really like the idea of collaboration across states and federal agencies. I like the idea of a collaboration among different expertises and against different types of victims and a variety of stakeholders. I think it's an important approach, it gets a lot of ideas on the table, and it avoids the dangers of the quick, easy solution.

I always like to say that the absolute longest shortcut or the absolute longest distance in the world is a shortcut. You have got to take the time to find the nuances in the solution or it's just going to be a disaster.

MR. MCCOY: So I will echo what Pamela said, is

66

that I think we really do have to put a lot of effort into trying to understand and measure the kinds of harms that are going on and get a good grasp on these measurements, because when we start deploying solutions, it's going to be impossible for us to say whether the solution actually helped or potentially harmed people unless we have these good, deep measurements of the kinds of harms and the scope of the harms that we're dealing with.

MS. WYDRA: And I will make a final remark. You know, I have been actually impressed by the discussion here and its mix of technology, but also of compassion, and it's not necessarily what I would have expected from a panel talking about or at the FTC talking about these types of things, especially that a lot are computer-based and technology-based, but I think that's been a part of everything that we have all talked about.

Of course, we want these types of data breaches and personal information to be protected, but when it isn't, it has to be dealt with with compassion, and it's important to understand, as you just said, the types of harms that happen, because it makes people more motivated to make sure that they don't happen.

MS. SMITH: And I would say as we think about these injuries, writ large, to ensure that we are sort

67

of approaching them methodically and really identifying the harms that we're most concerned about and separating that from sort of what the causes might be and using that as a step to get to what the solutions might be, and then understanding that technology in its many forms today can create some of these harms but could also be used as a tool to prevent some of these harms that may have been perpetuated by sort of humans in a less technological form before but now could potentially be mitigated with these technologies.

MR. WOOD: Okay. Well, that's nice to end on a hopeful note, but we do have to end. So there will be a 30-minute break after this panel, and so the next panel begins at 11:15. The cafeteria is open until 11:00, so if you --

MS. CONNOR: Not it's -- oh, open until 11:00, I'm sorry.

MR. WOOD: Open until 11:00, and then it closes. So if you want coffee, that is your best bet. We will see you back at 11:15.

(Applause.)

(A brief recess was taken.)

68

PANEL 2: POTENTIAL FACTORS IN ASSESSING INJURY

MS. MITHAL: My name is Maneesha Mithal, and I'm the Associate Director of the Division of Privacy and Identity Protection, and with me is my co-moderator, Neil Chilson, who is the Acting Chief Technologist of the FTC.

I want to introduce the panelists quickly. Their bios are in your packet, so I won't go into detail. We have Alessandro Acquisti from Carnegie Mellon; James Cooper from George Mason; Michelle De Mooy from the Center for Democracy and Technology; Geoffrey Manne, from the International Center for Law and Economics; and Paul Ohm from Georgetown University.

So before we get started on this panel, we just want to set the stage a little bit. Now, on the first panel, you heard a lot about the bad outcomes, the really bad outcomes that can come when bad actors, in particular, get your data, and in this panel we are going to be talking a little bit more about the responsibilities of commercial entities that collect and store your data.

And so what we are going to be doing is we are going to present a privacy hypothetical and a security hypothetical, and we're going to ask the panelists to raise their hands in the hypotheticals when they hear

69

that there has been injury taking place. And the goal is not to come to any legal conclusions but to really have a policy discussion and a policy back-and-forth about why people raised their hands when they did.

We also want to ask the panelists if you could raise your name tents when you have something to say so we know who to call on. We do hope that there is some really interesting back-and-forth. And Neil and I will be switching off moderating duties, so with that, let me just turn it over to Neil.

MR. CHILSON: Thank you very much, Maneesha. Thanks to our panelists for being here. And thanks to all of you.

So, yes, so we are going to do a hypothetical here, and when -- the panelists, as I read this along -- and there will be accompanying bullets on the screen for the audience. Once you raise your hand, unless you hear something that changes your mind about whether there's been consumer injury, leave your hand up. And then, like Maneesha said, we will be discussing why you identified injury at that particular point.

So with that, on to our privacy hypo. So in this hypothetical --

MS. MITHAL: I'm sorry, while we are getting the technology cued up, I just want to give one disclaimer,

70

which is that we're really not here to talk about the law and the kind of legal ramifications of Section 5. We are really here to talk about injury as a policy matter. So, again, when people raise their hands, this is not kind of what qualifies as injury under Section 5, but this is kind of when do you think injury has occurred.

MR. CHILSON: And part of this is to explore less the line that the participants have drawn in raising their hand and more why they decided at that point to raise their hand. So on to the privacy hypo. (Slide 21)

A pharmacy uses retail tracking in its stores to determine the most effective way to display greeting cards. (Acquisti, De Mooy and Ohm raise hands.)

The pharmacy then begins to track aggregate consumer interest in over-the-counter HIV tests.

The pharmacy begins selling this aggregate information to interested market analysts.

One marketing company uses its own algorithm to associate this aggregate information with other data to estimate the probability that a specific consumer has purchased either a greeting card or an HIV test. (Cooper raises hand.)

The marketing company then uses the data to target advertising to identified consumers, including

71

Carl Consumer. (Cooper lowers hand.)

Now, continuing with the HIV test example, the marketing company advertises HIV tests to friends and associates in Carl Consumer's social network.

MS. MITHAL: James, is your hand up or down?

MR. COOPER: It was up and now it's down -- it's resting. I am very weak.

MR. CHILSON: The ads mention that Carl Consumer recently purchased this product. (Cooper and Manne raise hands.)

A local insurance company gets this information and raises rates for Carl Consumer. (Cooper lowers hand.)

Carl Consumer's employer sees one of the ads and fires Carl.

So those are the eight sort of framing sentences of the hypothetical. We got an early jump there, and I am very curious about what that is. I think --

MR. COOPER: Before you even said anything.

MR. CHILSON: Yeah, and I'm very curious about why that is and what it was in that first sentence -- I have some suspicion that "retail tracking," as a term, has some baggage, and so let's just kind of run down the line here and have each of you explain why you raised your hand when you did, starting with Alessandro.

72

MR. ACQUISTI: Well, this is my thinking. Clearly, if you are defining injury or harm specifically as a realized and a quantified economic harm, I suppose that most of us here and the ones who raised their hands in the first scenario there would agree that there was no realized quantifiable economic harm.

However, there would be a very reductionist definition of injury which would ignore scholarly research on privacy, not coming from the legal profession that I know you want to avoid for this panel, but coming from social sciences.

Think about the work by Irwin Altman, for instance. Privacy is not the protection of data. Privacy is a dialectic process of boundary management, which includes both the opening of the self to others and the closing of the self to others. These boundaries are affected by social norms, expectations, individual preferences.

So in the context you are bringing up in the very first -- with the very first scenario, some of the key questions for me would be whether Carl was, indeed, aware that, as he was walking through the store, his behaviors would be tracked. Did he consent to this information being used for other purposes? If not, then there is a possibility that that boundary has been

73

broken, and when the boundary has been broken, well, then, that can be considered an injury.

In addition, I can easily jump from scenario one to scenario nine, which is the -- perhaps the most of the actual realized harm by creating a slightly different hypothetical, which is the pharmacy is using tracking via video. The video gets leaked. Carl's employer sees the video with Carl and fires Carl.

So we jumped entirely the other eight steps -- seven steps, and we went directly to the harm. So the point being here that when there is a breakage of the boundary, we increase the likelihood of a potential downstream cost, what economists refer to as expected costs, which are very important to consider, because agents, economic agents, both consumers and companies, make decisions based on expected benefits and expected costs. So we have to consider that in analyzing privacy harm.

Finally, I tend to steer away from a purely narrowly economic definition of injury and harm because the harm itself, the economic harm, even when it's there, it's incredibly hard to quantify, and for a number of technical reasons, which I hope I can -- we can get into later. I probably can pause here, let others talk, but I would like to go back to the issue of

74

why quantifying economic harm is so hard.

MR. CHILSON: Great.

James?

MR. COOPER: Yeah, thanks, and thanks for inviting me. It's great to be here.

So I get -- you know, I raised my hand -- I kind of went up and down a lot, and so one -- one, two, and three, you know, we still have that aggregate as the -- as the qualifier there. So you think that I have something private that I want to control the information, so if this is aggregated and there's no -- it is not really known, I'm not -- nothing has been revealed about me.

So, you know, in -- and you can even -- and I'm willing to even entertain the notion that you may want to keep your interest in greeting cards private. I mean, no -- I'm not here to dispute -- this isn't about, you know, well, that's obviously innocuous, who cares? I mean, someone could genuinely have utility loss from having people see the greeting cards they're looking at or something like that. But at this point, no individual person knows, certainly no algorithm knows about you.

When you get to number four, I think it gets -- and I put my hand up here -- and I think it gets to be a

75

closer call because at that point you're taking this aggregated information and somebody's saying, okay, well, now I want to find out more about Carl. I want to pull out -- I've got this giant lump of data, of people who have been at this drugstore, but now I want to see, what's Carl into? What is he buying?

And at that point you're starting to reveal something about Carl, and so I think there you start to get into -- if we're talking about privacy harms or informational injury, I mean, if we're thinking about the kind of harms that can involve privacy -- and I think I can talk about a distinction there in a minute -- that at this point you could have that because something is being revealed specifically about Carl.

So to me, one of the big differences between -- you know, just to sum up -- you know, between one through three and then four is you're going from aggregate to individualized, and then you think that's where you can get into the dignitary harms, the things that we think about with privacy.

Now, when we get into number five, at that point, the -- you're -- and I guess it was maybe unclear what the hypo -- this is the kind of thing where my student would come to me saying, I don't understand, it wasn't clear, and that's why I missed it.

76

MR. CHILSON: There's no right or wrong answers here.

MR. COOPER: Yeah. Well, there are on my finals.

MR. CHILSON: There might be later on.

MR. COOPER: Yes. So, anyway, the -- here, you know, there are two potential harms. So the -- you know, you're targeting ads to customers for this -- and it didn't really say are you targeting ads for greeting cards or HIV tests? And here I think that -- now, with the greeting cards, maybe there's not a harm with the HIV test. I still think it's a targeted ad. You're not really -- you do have the potential cascade where someone's looking over your shoulder or sees that you get an ad for an HIV test.

Well, what does that mean? Kind of like the Target baby ad thing that people talk a lot about, but I have a harder time with, you know, targeted ads in general. I mean, it's kind of an intrusion into seclusion maybe, but I have a harder time there.

And the same thing with, you know, number six. Here it's like, okay, I've got this data, and it seems that Carl's interested in these HIV tests, and so Carl's got a network of friends, and I'll advertise it to those. And, again, as long as they're not necessarily linking that -- you know, maybe somebody's light bulb

77

goes off and says, oh, I'm getting this because I'm friends with Carl.

I think that's a tenuous -- maybe too tenuous of a link, but my hand then goes up, of course, when we get to number seven. I think there that that's clearly a privacy harm at that point, because you're revealing something about Carl. And, again, part of the hypo here is we don't even know if it's true or not, but regardless, you know, you have made a prediction that Carl is purchasing these HIV tests, and all the -- and it's certainly something very sensitive, and you're telling all of his friends in his network, you know say, "Hey, buy this, because Carl has bought that."

And I think at that point you're revealing something sensitive to people, and I think that there is some -- you know, there's some empirical lit out there to suggest that people care more about revelation to other people than they necessarily care about an algorithm or a server somewhere, you know, knowing something about you.

Now, finally, and, you know, when we get to eight and nine, this is where I kind of make a distinction between -- I think there are two things you have to think about. There's the direct disutility harm that comes from someone knowing something private about

78

you. I have written about this a little bit, I call it intrinsic privacy harm, but essentially that dignitary harms, loss of autonomy, all of the sort of truly -- the loss that you feel when someone knows something private about you or you lose control of your information.

Now, the out -- then there's the -- the outcome that when a third party knows that about you, they act on that information. Now, if it's -- and I know this is maybe something we will get into a little later. I mean, maybe it matters if it's true or not, but assume for now it's true. If the insurance company now knows something true about Carl and makes a decision based on that, now Carl would like to keep that secret, I mean, that's strategic on his part, like I would like to pay a lower insurance rate, and the insurance company now knows that.

Now, that's -- is that -- it's a consumer harm to Carl, yes, he's paying higher insurance rates, but just -- you know, it's a reduction in adverse selection. So in some ways the surplus for society as a whole is getting larger. Now, I know this sounds cold-hearted, I'm not saying, "Oh, you know, poor Carl," but it's not to say -- and then, you know, going further, I'll lump the boss in there, too. So the boss fires him.

Now, there are a million reasons the boss may do

79

that, legitimate, illegitimate reasons, and that -- but, again, if they're acting on truthful information, then I -- but I wouldn't categorize that as a privacy harm. I mean, there are very legitimate reasons why we as a society prevent classification on certain -- based on certain attributes, but those are third parties acting on truthful information, and I think that's not the same as a dignitary-type harm that comes from privacy.

And I think you should make a -- there's a distinction in that -- again, we can use discrimination law to get at these, we can use other means, and this isn't -- this panel is not about law, but I think that there are -- to me, in my mind, there's a distinction between those two.

And I've probably gone on too long, but...

MR. CHILSON: No, that's fine.

Michelle?

MS. DE MOOY: Everyone, thank you so much for having me.

So where I'm starting is with the idea of privacy as a core principle to democracies, in particular. And so when you start with that, the reason that I raised my hand at the very first part of the hypothetical, which I knew would be kind of funny in a sense, it's not that I think that this person has been injured in a physical

80

way, per se, but I do believe that the violation of privacy has occurred, and the reason is because, first of all, their expectations matter.

So when you walk into a pharmacy, I think most of us -- or any kind of store -- don't have the expectation that our phones will be pinged repeatedly by a tracking system. Also, the idea of whether or not Carl was asked for consent, did he -- was he asked for permission to ping his phone? And, of course, when that happens, it's typically not just one small piece of data that's getting extracted, but many. So had he given his permission for that to happen? Did he have any control over the level of tracking that occurred? In other words, did it every single time he went into the store happen or just this one time for 15 minutes or when he was near the greeting card area?

Also, what was the benefit to Carl in this scenario? I think this is something I want to bring up later, because I think it's hugely important here. You know, a lot of the discussion around privacy, and particularly I think in FTC cases, sort of assumes that there's a benefit to consumers, to individuals, through whether it's behavioral advertising or tracking of any kind. And I think the question to ask is, where was the benefit here? Did he receive any kind of

81

benefit for this transaction?

Also, did he have access to or understanding awareness of what was occurring? This goes along with expectations and consents. And then also the idea of risk I think should come into play. So we know that the way that privacy harm happens is through small privacy violations perhaps, right? And I think this was discussed a little bit on the last panel. It begins small.

And so, therefore, the very first part of collection and tracking, that is where the risk is raised. So the fact that this information, what was taken without permission, et cetera, which is my assumption here, that means that his risk for identification, his risk for all of the other harms that come later has been elevated, and so that triggers obligations of the tracking company in terms of whether or not they are providing benefit, whether they're providing control and access.

And then I would -- you know, my hand would be raised for the rest of them, of course, but for different reasons; similar sort of principles but different reasons. For example, the aggregate to me is meaningless. The fact that it's aggregated is -- you know, it's one method, but it's very meaningless when it

82

comes to protecting information and protecting it from identification.

For example, how much data was there in this collection? Maybe there was three people who bought an HIV test, and, therefore, you know, Carl is pretty exposed in that aggregate data set. So, you know, it's not clear there, but I would say that aggregation, in and of itself, is not a panacea to any of these other issues.

And then the idea that this is also related to health information, you know, I don't think that sensitivity of data is everything by any means, but I do think when you're talking about health information, and particularly highly sensitive health information, like HIV status or a concern about HIV status, that elevates it, because this is immutable information, right?

Our health is not information we can replace easily. It's not information that can go somewhere else. It is immutable and intrinsic and inherent to us. And so, therefore, I think it raises more risk in terms of harm.

I think, finally, the idea of whether or not a person has recourse or -- you know, this ties to awareness and consent and expectation, but do you have any means to change this or to say I don't want this

83

information to be marketed or I -- you know, I don't feel like this is in my best interest and, therefore, I would like to reduce my risk of some of the harms that I see occurring by not allowing this collection to happen in the first place?

MR. CHILSON: Great.

Geoff?

MR. MANNE: Thanks, Neil, and thank you for having me here and thanks for listening to us pontificate, as if we know something.

That is a big part of what I want to say here, is that there's a lot less that we know than that we don't know in this area, and one of the really crucial things that I have been sort of thinking as I have been listening to people talk is that people are identifying something as injury, but they are not the sorts of things that we would all clearly understand as injury, in ways that are -- it's just not clearly the case that those things are, in fact, injuries, that they harm utility, that they are a painful or otherwise objectionable thing to, let's say, most people. Even that is hard to know what the right categorization is.

And so, you know, so one of the things here is that all of the things that we've been talking about and all of the things on the hypotheticals are all

84

describing aspects of information relationships. They are talking about how various entities interact with consumers around information, but that isn't the same thing as an injury.

The fact that information may be involved in something that's happening and has probably happened in some form or another since the beginning of time doesn't convert it into an injury. It helps to describe it, and it may help to understand how it could lead to injury. It may help us in certain contexts to understand things that are, in fact, injuries, and this goes back to my first point. We don't know that yet, but with enough data and enough analysis, maybe we can figure that out.

And so I -- all the way up until at least number seven, my sense here is that anyone who says that there's an injury here is either generalizing from their own experience, which is really all we can do, but still, we need to be very cautious about that, or intentionally or not converting an information relationship into an information injury, and I want to caution very strongly against that.

I think that risk is, of course, a really important part of this, but a risk of an injury is not actually an injury, and that's another really important piece here. For example, with number six, the marketing

85

company advertises HIV tests to friends and associates in Carl Consumer's social network, let's say -- I don't actually know for certain that this is true -- but let's say that it's fairly clearly the case that Carl would be injured if the information about the HIV test were revealed to people who could identify him, right, to people he knows or something.

The fact that the -- the marketing the test to friends and associates in Carl Consumer's social network -- and I'm assuming someone named Carl Consumer is going to have lots of friends, so it is not going to be clear that it's him -- but that may indeed increase the risk that someone will be able to figure out that he had purchased an HIV test, and that may, indeed, impose harm, but that fact itself does not strike me as anything that we should recognize as itself being an injury.

If risk of injury were enough to constitute injury, literally everything, literally the existence of these businesses would increase the risk of injury and, therefore, be actionable. And I think we would all understand that that can't possibly be the case.

I think it's also difficult in the context of how the panel was set up, but it's difficult to talk about injury without also talking about countervailing

86

benefits and talking, therefore, about sort of net injury, right? James started to talk about this a little bit, and I don't know -- you know, maybe we should be more clear about this as we go ahead. It may be that, for example, with the revelation that Carl bought an HIV test, that Carl himself was injured. It may also be that all of Carl's sexual partners now know that they should go out and buy an HIV test themselves, and the net benefit may be quite positive. That's also the case, as James pointed out, with respect to the insurance company, but I think particularly acute in that instance of Carl's social network knowing that Carl bought an HIV test.

Now, again, obviously -- I think, again, with the caveat that none of us really knows -- but let's say obviously it's a harm to Carl, but is it an injury that we really want to stop? Is the conveyance of information that could lead to net social benefit something that we should be calling a harm?

I think we have to be a little -- you know, sort of careful about that, but, again, just -- it's important in this context to caveat that. I understand why it's a harm to Carl, but the relevant question here is, is it a harm that the FTC, for example, should take account of? And that's a little bit harder.

87

MR. CHILSON: Great.

Paul?

MR. OHM: So all four of the co- panelists did such a great job kind of dealing with this, so I am going to spend most of time now responding to things that have been put on the table.

Nobody I think, if I recall, tried to kind of define what they mean by "harm" just to begin, so one working definition that I think philosophers and legal scholars have used is, are you worse off than if the conduct had not occurred, right?

And I think that, frankly, the liberating conceit of that we're not supposed to think about the law, we are just supposed to think about the word "harm" and "injury," makes these hypotheticals really easy, in ways that I think Alessandro and Michelle pointed out, that in almost every -- in every single one of these, we can point to something that is an injury.

Now, I think a lot of Geoff's comments portrayed the idea that they may not be injuries that we want the legal system and the legal apparatus in an enforcement agency to be able to remedy, but that's a different question than the question of has an injury occurred. So let me just address some of the things that were said.

88

One is risk of injury is not an injury. That makes absolutely no sense to me, right? We have many examples economically, but also, if we take a broader lens on things, where if something is in state one and then, because of the action of another, it becomes a much riskier state two, you have been injured, right?

We do this in medical malpractice contexts. We do this when it comes to the value of our consumer goods. If you didn't face a risk, and because of the action or negligence of another actor, you now face that risk, that's an injury. I don't even understand how it is not, and that's kind of in broad economic terms.

Layer on top of that, in the way that Alessandro urged us to, a hundred years of writing about emotional distress and anxiety, the kind of things that befall every one of us given the information insecurity we all live in, and I know this is the privacy question, but it goes for privacy as well, right?

I mean, if any of you were you in the room for the first panel, I bet your pulse started to quicken about midway through and probably hasn't come down yet, right? Knowing about the harms that we are all subjected to -- and, again, there might be countervailing benefits that justify these harms -- but knowing that there's an increased risk of certain harm

89

is itself an injury, and I say that kind of both from this abstracted, philosophical conversation, but I'm happy to say that as we move into thinking about the law later.

Let me say two more quick things. James put on the table and I think Michelle capably rebutted the idea that the word "aggregation" is some sort of kind of holy shield that can protect you from the idea that you're putting at risk the people whose information you are handling.

Now, the one thing that I've said in a lot of my writing and I think is pretty intuitively understood now is that utility of information of the kind Geoff was talking about is the other side of the coin of privacy invasion, that if you aggregate the data so much that you're reducing the risk of privacy harm to almost zero, you've also rendered that data totally unuseful for any commercial purpose.

On the other hand, if what you mean by "aggregate" is, yeah, it's aggregate. We don't know your name, but there is so much rich information about your transaction or the transactions of a few people that we are going to be able to, like, sell it to advertisers and insurers and employers, then that also means that the risk of privacy is still latent within

90

there, and so you can't have one without the other, right?

I wish there were a magical wand that we could wave that would suck out all the privacy risk from a pool of data and yet retain the utility. That doesn't exist. It's the exact same attribute of data that provides both of those things.

Last, but not least, when I read the first hypothetical, I thought, what is a retail store? And then I remembered these past memories from my childhood where you would walk into these buildings and buy things. One thing I think we should think about when we think about privacy and harm is what kind of population is affected by the harm, right?

And I think it's fair to say that for certain retail establishments today, we're talking about an older population, a less digitally connected population, maybe a less educated population. I think that's fair game to bring into our harm analysis as well, that if there's a harm that maybe isn't visited on most of us because we all are Amazon Prime customers at this time, but it is targeted at, you know, older people who go to a particular pharmacy, I think that should factor into the way we assess the harm.

Thanks.

91

MS. MITHAL: So there is a lot to unpack here, and, Geoff, I will give you a chance to respond, but let me just tee up the question I think everybody wants to talk about, which is, okay, you have all raised your hands and identified where you think there's a harm. Now, at what point do you think there should be government intervention?

And so maybe we can start with Geoff and go back down -- or actually, why don't we start on the other side. We will start with Paul and come back down this way, and maybe just tell us the number of where you think that there should be government intervention.

MR. MANNE: Can I just respond to one thing Paul said first before we do that that I think will help, because it's not about that?

MS. MITHAL: Sure.

MR. MANNE: I do want to just sort of reiterate what I said a minute ago, Paul, and try to hear more about how a risk of harm can itself be harm when, again, literally every bit of activity increases the risk of injury, right?

I mean, that's like saying a drugstore that serves one additional customer has created a cognizable injury, because by increasing the amount of activity, it's also increased the risk that somebody would be

92

injured by whatever might happen -- befall that drugstore. So I don't think it can be that.

I also think that it's essential in trying -- I like that Paul actually tried to identify a little bit the kind of disutility that one could experience from even a risk of harm, this sort of idea that the knowledge of risk can create an anxiety. And I'll admit that that could very plausibly be an actual injury, but then I get back to my very initial point, which is, you know, we can speculate about that all we want, but I don't think that makes it so.

And I think it behooves the FTC and many others to try to figure out whether there is actually something cognizable there, you know, obviously within the context of the legal regime, but, I mean, even just independently in terms of defining what injury is.

And then, finally, I don't see how any of those things could be injury if the information is already known or is already out there or the risk of it being exposed is already there, so the anxiety is already there.

So one of the really important things here, it seems to me, from the way I hear a lot of the panelists talking about it, requires some awareness of the preexisting risk, in the case of risk, or the

93

preexisting exposure of information, if that's itself going to create a harm.

I look, for example, at number two, the pharmacy begins to track aggregate consumer interest in HIV tests. Well, Carl certainly knows that the drugstore already knows if he bought a test. They have that information about him. Whether they're tracking other people's information or not doesn't actually affect the anxiety he might feel about somebody knowing this, because he already knows perfectly well that the person at issue here knows it. So it's hard for me to see how that could increase the risk, at least to Carl, although I understand you may think the aggregation of information creates a second risk, okay?

MR. OHM: So I am happy to jump into this second question, and my overly lawyerly answer is, depending on what some of the words mean in the hypotheticals, I think every one of them could justify government intervention.

And as part of the backdrop, when I think of government intervention, I think more broadly about legal recourse. Is there a court, under any theory of law, with any plaintiff, that can get recourse for significant injury, right? I want to make sure we're only talking about significant injury.

94

The courthouse doors are being closed to kind of tort plaintiffs left and right, mostly because judges fear that if they allow too many class actions to proceed, it's going to get out of hand and there's going to be a lot of vexatious litigation.

So in that climate where none of these things are going to be easily redressable in tort -- or maybe most of them won't be -- I think that kind of raises the urgency for an agency like the FTC to step in, especially when they think there is a serious harm befalling a lot of consumers based on information imbalances. So I think it behooves the FTC to kind of step up and fill the gap of the closed courthouse doors that I'm referring to.

And so let me just, you know, go through two really quickly. HIV, right? HIV is not only a kind of significant medical condition that still today, sadly, has a lot of unfounded stigma attached to it, a devastating effect on reputation as we heard in the first panel, but it's also tied intimately to sexual behavior, right?

And so to kind of address Geoff's direct point, the hypothetical says interest in HIV, not purchase of HIV, and I take this to mean, you know, perhaps one of these new in-store retail scanners that will tell that

95

you lingered by the HIV test shelf for a while, or maybe if you tie in RFID, that you picked up two of the boxes and then put them back, right?

And so to me that's where we're starting to kind of tread into significant sensitive information, and Michelle kind of said that's not the be-all, end-all. I think it's a really useful rubric. I think it's a widely accepted thing outside of the context of the FTC in the law, that we should identify as kind of shorthand aspects of information that are sensitive, and if that's the kind of thing we're talking about, collection in an unexpected or new way or using a technologically new ability, that's where we should be much more worried about the kind of risk of injury that is, in my mind, harm. So that's one line. If you want a line, one line is if information is sensitive, then maybe the government ought to intervene, okay? So that's number one.

Let me give you one more. We think a lot about kind of the violation of some other positive law, right, some expression by Congress or by a legislature that some act is not only unexpected in the way Michelle described but also violates some law, arguably violates some law.

I think it's become common for retailers and

96

stores to kind of look at MAC addresses emanating from smart phones. I know there has been some activity about this in the FTC. That is an easy violation of the Pen Register and Trap and Trace Act, right? Congress, in its infinite wisdom, has said that this is a misdemeanor crime.

Now, there is no plaintiff's action attached to that, so you never see this enforced by anyone, but Congress said, in the same way they did with wiretap law, that there's something about this collection of this kind of information that is illegal, right?

So that's a second heuristic, rubric, call it what you want, that would say that the FTC, or some other mythical government agency, should exercise its ability to kind of vindicate the rights of people who are injured in hypothetical one, hypothetical two, and then most of the other hypotheticals flow from one and two.

MS. MITHAL: Geoff?

MR. MANNE: Okay. So just very quickly, I think one of the things that Paul said and that is sort of at issue in this hypothetical with the retail tracking sort of idea is that when you have a new technology or a new form of data collection, that's where we should be the most vigilant. I think exactly the opposite is true, of

97

course.
I think, or at least I think it's imperative to point out, that it is in the cases of newer technologies and innovations that we want to be the most careful about overenforcing the law and overdeterring investment in innovation.
And so, you know, one of the big problems I see with overenforcement is overdeterrence of experimentation in the sorts of areas that we actually really want, and I don't just mean experimentation with new technologies, but I also mean experimentation with new forms of information relationships that people may or may not actually care about.
They may -- that they may actually prefer, that they may be willing to pay for, any sort of number of relationships that you can describe. And if every single effort at trying one of those out leads to potential liability, none of them will ever be tried out. And so that strikes me as being, again, exactly backward.
I see also sort of a related reason, a really serious problem, where we're making illegal the collection of data. There should be a really, I think, significant distinction between the collection of data and the use of data; the -- even the increased risk of

98

some actual cognizable harm via a data security problem arising from the collection of data. I think that's really problematic, but at least that has a logical coherence to it.
But, again, the idea that by collecting data, we're going to over -- to dramatically deter that, we are never going to find out all of the things that we could do if that becomes the sort of thing that no one actually wants to engage in.
And so in terms of trying to identify where the government should get involved, I do think that we should err on the side of, you know, where we actually can really identify that there are viable harms here.
Again, in some cases, we are going to know where that is. In some cases, we don't, and that means that before the FTC should start intervening, it should start collecting data. It should start with things like this workshop -- which is great, this is a great start -- but I think there's, you know, years of work to do beyond this before we're going to start -- before we should start identifying that the government should be deterring these examples of data collection.
MS. MITHAL: Okay.
Michelle?
MS. DE MOOY: I just want to push back on one

99

thing that you said earlier, that there's a relationship here, and I think that that's debatable, and in most cases, I think it's debatable whether -- usually a relationship involves at least two parties, and I'm not sure that Carl is aware that he is in a relationship here, right? Maybe it's a stalking relationship, I don't know.
So I think that that's an important point to make, you know, that the -- his expectations, his understanding of the situation is probably different from the tracking at the pharmacy and the continued other interests involved here.
And part of the reason I bring that up is because, again, the question of whether or not he benefits from this exchange I think should be a part of any kind of legal rubric to determine where -- you know, the level of risk, and I think also, of course, ties to consent and the person's expectation.
I think the idea -- you know, the government already intervenes when it comes to sensitive information, so I agree with Paul that sensitive information should trigger obligations. You know, this data, in particular, is not, of course, covered by legal frameworks, but in my opinion should be, not because it should be illegal, but because it should be a part of

100

the assessment for -- whether it's the government saying you're not allowed to do this, and the levels that reach up to that.
And then there's the other threshold of maybe harm where there's remedy for the individual, right? I think those are maybe better ways to think about how government intervention would make sense, and this is something that you can see in other frameworks, where it does make sense.
I think I fall on the side of collection increasing risk because there is, of course, always the risk of surveillance. This is a fact in our data-driven world, and it is a part of almost every product and service that an individual interacts with in the digital age. So the idea that this information can somehow get out and get loose is not a fantasy.
This has, in fact, happened over and over again, and sometimes the impact is worse on some populations and not on others, and, therefore, I think the other part of this assessment should include what Paul said, that it depends on the population, and their particular place in the ecosystem does make a difference in terms of the effect of the harm and the impact of the harm. And so perhaps that would inform whatever remedy was offered by the government to the person who was harmed.

101

And then just sort of generally speaking, I fall, I think, in the category of -- the idea that -- of sort of Professor Calo's rubric on this, which is that there is subjective harm and objective harm, you know? So subjective harm is the perception of loss of control, and it's the sort of -- that results in kind of a fear or discomfort. And then, of course, there's the objective, which is where there's an actual adverse consequence. And, again, I think those should be divided by the idea of what is permissible, what raises risk, and what should involve consumer remedy.

MR. COOPER: Thanks.

So I think, you know, we're moving here from talking about -- we went down the first line saying what is harm, and if we're talking about an individual, we're talking about, you know, Carl, well, you know, again, as I mentioned, he may be harmed if people know about his greeting card habits, but when we're -- and, you know, I mean, legitimately.

I mean, there could be -- I know it's hard to say that with a straight face, but by the same token, you know, everyone -- there is no accounting for taste, and everyone has their utility function, and economists are agnostic about that.

MS. DE MOOY: Sure, and maybe he is buying it for

102

his mistress.

MR. COOPER: Good. That will show up on next year's final.

But when we think about regulation and government, we can't make -- at least not in this role yet -- we can't make individualized rules, right? So we have to look at distributions of -- we have to look at, you know, kind of the distribution, where we can draw the line.

So I think that, you know, in this case, when we talk about the greeting cards, we think, well, we can -- I think it would be pretty easy here to say that, well, there could be some people who are especially sensitive about their greeting card -- greeting card habits, but it would be -- it would be hard to see for me, you know, government intervention, especially some sort of aggregate or even individualized, like I want to have an algorithm to predict what kind of greeting cards people like and send out ads from my Hallmark store to say, well, try this greeting -- you seem to like these greeting cards, try this, because, again, you know, there may be some people way out on the tail, but we can't individualize our rules.

But this -- and what I'll echo is something that both Paul and Michelle and I think Geoff -- maybe, maybe

103

not -- but, you know, I do think that the type of -- the type of data do inform that. So if we're talking now about the HIV status, I mean, for all the reasons that Paul and Michelle gave, I mean, there's a lot of -- that's very, very sensitive -- that's very sensitive information, and I think here, when you think about what's the right enforcement posture, I mean, where would this -- where would we have government intervention, you know, I think you have to balance a lot of things.

I mean, first, there's just a direct utility harm from Carl, okay? So some -- this data about my interest or actual purchase -- again, not clear yet -- but in HIV testing is out there. That's a legitimate loss. There's also dynamic losses, and I think this is maybe the flip side of what Geoff -- I like what, you know, Geoff had said, okay, this information -- this information is out there, and perhaps there -- perhaps there are benefits to that to Carl's partners, but you also have to think about incentives to acquire the information in the first place.

So if I'm concerned about my HIV status being out there and it's something I want to keep private, well, maybe I will engage in privacy protective behaviors that keep me from learning that valuable information. So

104

there's kind of that dynamic part, and it's related to autonomy benefits from privacy. You know, how do you act under observation versus not observation? So, I mean, these dynamic things you have to consider.

But, you know, there are also -- there are also beneficial -- the data being out there is -- as Geoff said, there's benefits to that. There are benefits potentially to the insurance company. So I think that it's a different balancing.

I mean, we kind of look at where -- I think speaking as an economist, I mean, you want to unite information and control the person who's the highest valued user of that information, and maybe Carl's the highest valued user of his interest in HIV testing or his concern about that because of direct utility loss and because of the dynamic benefits that come from his -- you know, actually acquiring the information about his HIV status, which can be, you know, quite -- quite beneficial.

I would also agree here -- and I want to echo something Geoff said -- because I do think it's better to focus on uses rather than collection of data if we're thinking about a regulatory posture. You know, so I think that, you know, if we're concerned, say, about numbers eight and nine, like what the insurance company

105

is going to do with this data or what the boss does, I mean, you know, we -- you know, we could go all the way back to one and say we're just not going to allow you to collect any data in a drugstore at all, because drugstores are -- you are doing a lot of sensitive things in drugstores, so no retail tracking in drugstores. That could be one -- one because there could be this cascade of risk.

But we could also just say -- if we're concerned about HIV status, we could say, well, look, you can't -- you know, as an insurance company, you can't base insurance rates on HIV status, you know, if we're concerned about that outcome. We could say that bosses -- that employers are not allowed to fire based on HIV status if we were concerned about that.

We do that in -- we, as a society, make cuts like that all the time, that there is certain information you can't act on. I mean, I think rather than suppressing, you know, truthful information, that it's better to just prevent the uses of that information where we can, and I think that, you know, that has -- that has the two benefits in that.

So one is that if we suppress -- prevent wasteful -- like, if we don't want insurance companies to use this information, but we prohibit -- but we don't

106

bar them from using it, we just prohibit the collection of the data, well, nothing's going to prevent people from investing in signaling behavior.

So if we're worried about this -- I mean, so you are going to get signaling, you are going to get people trying to come with -- hey, I don't have HIV. And the other thing is -- and we have seen this with some recent studies with the ban the box initiatives -- is people engage in statistical discrimination. So if there is something that is actually useful in making your decision and can't ferret -- you're not allowed to have that information, then they'll find proxies for that.

And so there have been, you know, a couple of ban the box -- which is, you know, don't look at -- you can't -- there are about 25 states have them in different forms, but it basically says you can't look at whether you have a criminal record, at some -- you know, on a job application, and there's a good study by Agan and Starr, then another one by Doleac and Hansen, that basically -- basically finding the idea of ban the box is incredibly well meaning. You know, let's break the cycle of, you know, going to -- I go to prison, I get out, and there is no meaningful -- no one will hire me because I've been in prison, kind of this catch-22, so I just go back to criminal activity. So let's break the

107

cycle. Let's not allow -- and so we do this by suppressing truthful information.

An employer may actually have a legitimate reason to say I don't want to do this. So what happens when you bar -- when you engage in ban the box? Again, the two really good studies that are out there, they show that, well, discrimination against African-Americans goes way up, I mean, like ten times lower call-back rates in New Jersey and New York after ban the box than before for African-American males between, like, 18 and 25.

So I say that by -- you know, their cost is suppressing truthful information. People want to make decisions based on info, but -- so the way to go is not to say you can't collect it. Just say you can't use it. So anyway...

MS. MITHAL: Alessandro?

MR. ACQUISTI: Well, I feel that, although coming from different directions, both Geoff and Paul made a point I agree with, which is not all injuries necessitate government intervention, and there may be countervailing benefits arising from those injuries.

So the way I try to think about this problem, it is to go back, as I often do when I work in this area, to the seminal work on the economics of privacy coming

108

from Chicago School scholars in the seventies, such as Posner and Stigler. They pointed out that privacy protection is inherently really sticky, okay? It creates economic losers and economic winners. It affects the distribution of wealth.

And I believe they were correct in pointing that out, but I believe also they stopped short of recognizing that also the absence of intervention, so the absence of protection, is creating winners and losers. There is no way out.

If you intervene, you are going to affect the distribution of wealth. If you do not, you are still affecting by not intervening. So the dilemma for the regulator is how to choose whether to intervene or not. Some of my colleagues in the economics discipline believe that, well, when things are so complex, actually, take a step back. Regulate only when there is some dramatic, quantifiable, provable harm, and let the market do its magic.

Well, as an economist, although I do believe in markets, I also have reasons not to believe that, in the case of privacy, they work that optimally. First, we have ample evidence, which we have described in the Journal of Economic Literature Review, we published last year, we have -- Curtis and Liad -- we have ample

109

evidence, theoretical and empirical, that without government intervention, it is not a given that the markets will end up with an optimal amount of information sharing, information protection, so from the aggregate perspective. So we already have that.

Second, we also know that there are enormous information asymmetries when it comes to personal data -- how much information about myself is being collected, how it is being used, what the consequences will be -- which renders the individual responsibility argument, which is essential for good market outcomes, essentially untenable.

So what do we do when we face a scenario where we have stakeholders' interests in contrast, in tension with each other, as it comes to how much data should be collected and analyzed? And these interests are not just economic interests; they also relate to things such as autonomy, freedom, and dignity.

Well, I feel that the way to tackle this is not just to have a sound economic analysis, which we should have -- and this workshop is very useful in that direction -- but also listen to the will of the citizens through their elected representatives. That could be a good metric for government intervention. If a majority of others think that privacy is important, perhaps we

110

should listen to them.

MS. MITHAL: Okay. So there's a lot to discuss, but I think we have to move on to the data security hypothetical, and we might get some time to come back to tie the two together, but -- so let me -- we will do the same exercise. We will read out a sentence from the hypothetical, and just raise your hand when you think that there is injury, okay? (Slide 22)

So Company A stores consumer SSNs. A security researcher discovers that Company A has a security vulnerability that exposes its entire computer network, but no unauthorized access has occurred. (Acquisti and Ohm raise hands.)

Okay. Two, unauthorized access occurred, but confirmation that no consumer data has been exfiltrated, okay? (De Mooy raises hand.)

Unauthorized access has occurred, and it is possible that consumer data has been exfiltrated.

Okay. Unauthorized access and consumer data from Company A has been found on the dark web, but there is no evidence that it has been used for a fraudulent purpose. (Cooper raises hand.)

And then, finally, unauthorized access and consumer data from Company A has been used for fraudulent purposes. (Manne raises hand.)

111

Okay. So let's see, so why don't we -- we'll switch around the order this time with the -- why you raised your hands when you did. So why don't we start with James.

And actually, we have about 34 minutes left. We have the data security hypo, hoping to wrap up, so if you could keep your interventions short, and we can probably get in a few more questions or issues to drill down.

MR. COOPER: Okay. Yeah, so I went with number four just because at that point that's where I think that the risk of bad things happening is sufficient enough. I mean, you're on the dark web. We have evidence that these data are with bad actors for potentially bad purposes, essentially, and certainly number five, you're there.

But one through three, at that point, it's too speculative to me to say. Again, it's a -- there's probably some increased risk of harm, but it isn't sufficient, and I don't know if at this point we're just talking about, you know, legal intervention or where we are in the hypo, but I would say one through three doesn't -- isn't sufficiently cognizable in my view.

MS. MITHAL: Why don't we go this way -- oh, okay. Go ahead, Michelle.

112

MS. DE MOOY: Okay. So where I land on number two is a privacy violation, perhaps lesser on the assessment of harm, but still, nonetheless, a violation, and that is because -- I would go back to actually a former FTC Commissioner, Thomas Leary, who framed unfairness authority as "a tool best deployed in circumstances where third parties with whom consumers have no relationship do unfair conduct, practices prey on vulnerable consumers, involves coercive conduct, or creates significant information deficits."

So my assumption on this -- and this could be incorrect -- but my assumption here is that the consumer is not aware of this unauthorized access and in this case I think should be made aware of it. So as we all know, the limits to what we understand versus the limits as to what hackers and others understand, there's a great information asymmetry there, too.

In other words, if there was unauthorized access, I don't think that it's fair to assume that it's fine. In fact, I think it's fair to assume that it's probably out on some level. And so, you know, I think it just depends on which way you lean, which way you decide to assume, and I think, you know, from -- if you come from privacy as a core principle, you lean towards the protectionist idea.

113

And so, therefore, number two would not necessarily trigger government intervention or laws but that it might trigger some kind of awareness, some kind of notice and control for the consumer to be made aware of the unauthorized access and perhaps be able to take their data away, out of the company who, in this moment, has failed.

MS. MITHAL: Geoff?

MR. MANNE: I want to just -- let's see, I raised my hand with the last one, and even it is actually somewhat questionable, not that there's injury. There obviously is sort of -- it's defined in terms of injury, but in part to respond to what James said, I just wanted to ask -- I hate to introduce my own hypothetical -- but let me just ask, the Equifax breach, everyone agrees that that was injury? Anyone not think that was injury?

So -- I don't know, no one raised their hands either way. So no one thinks it's injury? The reason I ask is because I -- you know, I don't know if this is true or not, but I know that according to the IRS, of the 150 million records that were exposed, some 100 million, they estimated, were already on the dark web, that -- and the language they used in one place was we actually think it won't make any significant or noticeable difference.

114

Again, let's just take that as true for the moment. It may very well not be, but I think -- so, James, for example, pointed to number four. I think that even number four -- and even, for that matter, number five -- can't really be injuries so clearly.

Again, if the information was already out there and if it was already being used for fraudulent purposes -- now, there could be additional fraud and there could be additional costs, so we can see how number five could be an injury, but number four -- and number four, too, it doesn't really matter if -- anyway, it's whether the information is already out there, which goes to this point about making risk of harm into a harm itself.

I don't want to keep harping on this, but I will, because I think it's really -- I think it's really problematic here, especially in the data security context where -- again, so we don't know anything about any of the conduct here, all -- you know, we know that unauthorized access occurred, for example.

Well, I can tell you that there is a nonzero chance that even the most secure systems could be subject to unauthorized access, and we know that because the NSA was subject to unauthorized access, right? And so -- and, again, in the narrow confines of this initial

115

question, that doesn't actually tell you anything about whether there's been an injury or not, but I think it's -- in other words, whether someone was harmed -- but I do think it's essential, when we're starting -- that the things -- as you've seen, as we've been talking about this, they all blend together quite a bit, and it's -- if we're going to be talking about risk, for example, as being -- that risk of injury as being an injury itself, I don't think we can talk about -- we can talk about that without talking about the things that create the risk.

And I find it especially problematic if we are defining as injury something that can result from firms taking the utmost, absolute, you know, blockbuster care, well beyond what we would ever actually want them to pay to take, and if that can still increase the risk of some cognizable harm and, therefore, constitute injury, again, I think we're -- things are really problematic then.

So just to clarify, so maybe five, probably five, but nowhere before five.

MS. MITHAL: Okay. So Paul, then Alessandro, then James.

MR. OHM: So I warn the moderators, I am going to fight the hypo, but I'll fight it very quickly, because

116

I know we are short on time, and then I'll answer it not fighting the hypo.

Let me put another hat on. When I was a network systems administrator before I went to law school, I defended networks. If I ever met someone who said no unauthorized access has occurred, we can confirm that no consumer data has been exfiltrated, then you know that serious violations have occurred, because those are naive statements. Those are impossibilities, right? We can never be sure of things like that.

And it's the companies that are sure that their data has not been exfiltrated because their Sys Admin said we have an IDS, those are companies that I guarantee you are like, you know, just crawling with hackers at the time. So fighting the hypothetical, I want the government to investigate number one because that company is naive or lying, okay, but let me not fight the hypothetical, right?

So another way to kind of respond to Geoff's comments is do we find a world in which Equifax has occurred to be an acceptable state of the world, right? Is this a problem that we, as a collective and as individuals, should try and solve? And I think that is tied to the injury question, right?

Before Equifax, there was another -- you know,

117

there was a company a week who would demonstrate to us time and time again that, to speak like an economist, there are externalities that need to be internalized and they are not, under whatever mechanisms we have. The state of data security in the world is horrid, is horrible, and is causing concrete repercussions for everybody in this room.

I, like Elizabeth Warren -- I like saying I like Elizabeth Warren -- spent three hours on the phone and on the website trying to, like, go through the Byzantine data monitoring protocols that they have, and I wasn't able to do that after three hours. That's time wasted. There has been anxiety, people face financial ruin, right, people self-chill, they self -- they change their behavior.

They don't apply for jobs they might apply before because they are worried about their credit report. They don't try and buy a house because they know they are not going to get a home loan. And there are documented cases time and time again -- and Geoff has impressed upon us several times -- that we should look for documented cases. They are not hard to find.

And so for all of those reasons, I think all five of those can be defined as injury, particularly if we're not asking, is it actionable and legally redressable.

118

But let me end these comments with one point of agreement with Geoff, which is, yeah, I absolutely agree that we should account for how reasonable your data security was, right?

If the world's greatest hacker broke into the world's greatest security, then yes, the law probably should not offer redress against the world's greatest security purveyors, because they're behaving responsibly, and they have internalized the externality that I'm talking about.

It doesn't make it less harmful, but it does mean that there might be some notion of causation or countervailing benefit or something else that we should take into account. And so, yes, I think I agree with you there that, that in making these hard choices, we should think about how good was your security at the time.

MR. ACQUISTI: I raised my hand for number one for reasons similar to the ones Paul mentioned and very similar to the ones I brought out under the previous scenario, where there are boundaries, and whether the consumer even knew that the company, Company A, stores his or her SSN, whether Company A had a right to, in fact, have this information, how did it acquire this information, why and how it is using this, et cetera, et

119

cetera, et cetera.

Now, again, I will agree that not necessarily quantifiable, realizable economic damage had occurred, but the possibility of it -- in fact, there's increased risk of the downstream cost, which allows me to go back to a promise I made earlier in the more technical discussion of the economic harm, right?

Even if you want to narrow down the definition of injury to economic harm, then we have to face the enormous challenges of quantifying the term even when we know that it does exist. The challenges are enormous, and I will give you some examples.

One is that the harm is incredibly context-dependent. The very same piece of information could be harmless or even beneficial in one context and extremely damaging in another context. The harm can take very different economic typologies. There is the direct harm, such as Carl being fired or this consumer data being used for fraudulent purposes.

There is the opportunity cost. If my data is used by others, my ability to use it strategically decreases. There is the loss of earning. Someone else may be benefiting from my data, and perhaps I'm not getting a fair share from that -- those benefits.

And then there are also other differences, other

120

nuances, which, again, make it incredibly hard to pinpoint, quantify the harm. There are costs which are exceedingly small but happen continuously. The time I have to spend deleting the Spam message which my Spammer filter didn't catch, the increasing time, perhaps few fractions of a second, but across many consumers and across a long period of time amount to a huge waste of time due to the fact that when I log a page on the internet, the page is loading in its loader because of the tracking going on behind the browser.

On the other hand, on the opposite end, there are the costs which are catastrophic but are very low likelihood, such as catastrophic medical identity theft. And then the differences between the harms which occur immediately after some privacy invasion has occurred, and in your example, scenario one, Carl's employer firing Carl immediately after this information has arisen, and the harm which may take -- which may happen months or years after the fact, such as maybe someone suffering from the Equifax breach one year from now, making it incredibly hard for us, as economists, to prove causality, even though there is a high suspicion that there is a direct link.

So the point being that these costs are so diverse and so nuanced, the idea that we can create just

121

a simple metric to capture them all and their -- and a simple formula that the regulators can use to decide, oh, yes, I should intervene, or no, I shouldn't intervene, is really hard, seems almost to have a premise that is untenable to me.

MR. CHILSON: Well, great. Thanks to all the panelists for walking through the hypothetical and identifying why you raised your hand at the point, even if you didn't all embrace the hypothetical as written.

But, James, you wanted to respond?

MR. COOPER: Yeah. Well, I just wanted -- I guess since I put my tent up, a lot has been said. First, I am going to agree with both Paul and Alessandro. I mean, this is, as an economist -- I mean, there are -- as far as internalizing the externality, I mean, it's hard to think of how this can happen, but for all the link between, you know, I committed -- you think of the normal tort. I drove carelessly, I hit somebody, they got hurt, and that's easy, causation, you know, and then you -- so you said you calibrate toward law to internalize that. That can't be done here.

I mean, it is really -- as far as linking, I think it's very difficult, and I think that one thing we have to think about as we go forward and think about how to deal with -- and I think this tees off of -- goes off

122

on something Michelle had said and a theme throughout, is, you know, there's data everywhere. There's -- we don't know, if there's a breach, where that data came from.

Ultimately, as a system, I mean, we think -- if this is something we just live with, I mean, do we think of this as a first-party insurance world, where we just kind of all either self-insure or buy insurance policies against cyber risk and just let -- and then we reduce some of the -- some of the incentives that tort can potentially bring, but it's hard -- for the reasons that Paul actually talked about in the previous case, you know, that many or if not most data breach cases get thrown out either for lack of standing or lack of meeting -- pleading harm sufficiently. So, you know, the tort system, it's unclear.

So, you know, it would -- does it make sense to have first-party -- you know, to -- rather than being insured through the court system, to add first-party insurance, and then maybe backed up by some kind of FTC intervention when applicable? I'm not sure.

I do think it's super-complicated, but when I first put my tent up, I did want to kind of amplify something that Geoff had talked about is risk, and I do think that this is the big question. You know, as Geoff

123

said, maybe -- you know, I raised my hand at number four, but, you know, you raise a valid point, that, okay, number four, even number five, so there's harm. I mean, you were kind of fighting the hypo --

MR. CHILSON: If it was unclear, next year we will write the questions more clearly.

MR. COOPER: Yeah, yeah, well, that's how I -- but the idea that we can't -- it is going to be really hard to link up. So it's on the dark web. Maybe it was already on the dark web, and maybe it's on the dark web that has nothing to do with this data breach, you know, and so I think about this, and, you know, I was -- like probably everyone in this audience, I think about this as like a Bayesian updating problem -- that's a joke, by the way, or maybe it's not, maybe it's not, maybe everyone is thinking exactly like that -- but I think you kind of start off with some kind of view -- you know, prior view of the world for the odds of my data being misused as part of a breach, and then, you know, ultimately what we're trying to figure out is what are the odds that this breach is going to lead to some kind of demonstrable -- this conduct, whatever -- not the breach, I'm sorry, the conduct -- the vulnerability that this firm is engaged in, likely to lead to harm.

And so what you think about, you update your

124

priors by thinking how often, when I see a breach, is it associated with this kind of conduct? I mean, that's what's called a likelihood ratio in updating, and so the thing is, is how much does what we know change our -- about the -- the likelihood that this conduct is related to harm change our -- change our priors? And it could be that this really -- the delta, the change in the odds of harm are really, really high. It could be like a factor of two, but it also could be that the baseline of harm conditional on breach is so small that the posterior, my final -- you know, so it moves from the odds of harm from this breach being 1 percent to 3 percent.

So at that point, do we look at the delta, which could be really large -- and these are like epidemiological studies, where you start with a really low baseline of a -- some kind of condition, and then there's a drug, and the drug reduces that condition by four times, but it just goes from like, you know, 3 percent to 2.5 percent overall. So it's the same kind of thing. Do we look at the change in the risk or do we look at the overall risk?

And I think that -- to me, I don't know exactly where I come out on that. I know that there's a case whose name can't be spoken up here that that's one

125

theory, is to look more at the delta in harm, look at how this conduct is likely to change the likelihood of harm as opposed to the overall incident of it.

And also, I know that this is -- we haven't really gotten into this, but, you know, the extent to which we actually -- the conduct has been out there and it hasn't happened for a while, I think that does inform, but -- but anyway, I see other people with their tents up, so I will not go on.

MS. MITHAL: Geoff?

MR. MANNE: I will be much quicker than James was. In fact, I was going to put what James said into English.

MS. MITHAL: Thank you.

MR. MANNE: I think there's a big problem --

MS. MITHAL: Epidemiological?

MR. MANNE: -- in -- no, actually.

No, I think that there's a big problem, and it was reflected in what Alessandro said and what Paul said. No question this is really complicated, and James makes great points about assessing risk as a -- I won't repeat everything he said, but all of that and all of that discussion and all of the discussion, the acknowledgment that there are problems out there and that certain conduct actually can cause harm or cause a

126

risk of harm, I mean, all of that says nothing about the optimal level of injury, the optimal level of data security breaches, and the optimal level of care that's supposed to be taken.

And I sympathize or I agree, in fact, with everything everyone has said on this score, except nothing they've said is really operationalizable or even really particularly relevant until they've established that we're deviating from the -- from optimal or from some identifiable baseline, because while we can point to lots of injury, as long as we're all going to acknowledge that the cost of making the injuries zero or the risk of injury zero is far higher than we are willing to pay, you haven't yet established that we are actually at a point where we should be intervening more or identifying more, looking at changes in risk, you know, from 1 percent to 3 percent as actionable, and all of those things are totally possible, except none of them can be determined unless we have some better sense of what the optimal baseline is, and I don't think we have anywhere close to that. Thank you.

MS. DE MOOY: Okay. So I reject that, and I will get to that in a second, just a second.

MR. MANNE: It's tautological.

MS. DE MOOY: One thing that I want to say, just

127

to push back on something, James, that you said about tort, which is that the tort of assault requires imminence, and I think you were saying that wasn't part of that assessment when it is.

MR. COOPER: No, the imminence is -- like the tort has -- it has been -- these cases have been thrown out, most of them have been thrown out either on standing grounds, or even if they make it past standing, they -- so I think we were in agreement.

MS. DE MOOY: So I just want to finish my thoughts. So I think -- you know, fine, but I think the point that I want to make is that the FTC can look at these issues in a much broader, richer way than the court system, right? I mean, we can look at social harms in a way that the court system cannot. I think that's an important part of determining how to govern here.

And so I think that the imminence of -- is akin to the idea of risk, and I think that's important. I think -- I just want to also mention, while we still have time, that I think the way that the FTC can approach this to respond to your sort of fatalistic feeling that we can't actually --

MR. MANNE: No, optimistic. I'm saying we have to do it, that we should do it.

128

MS. DE MOOY: We do have to do it, and I do think that there are, baselines, like no breach. There's a baseline. Now, the idea of how you penalize breach or practices, of course, is up for grabs, and I think there are ways to do that also. There are precedents for what is permissible in data security, and, of course, those might change over time. And so this has to be a fluid framework that can do that.

I think unfairness has that potential. I think unfairness has a much broader reach than deception, and I think that is where the FTC can begin to explore how to assess a risk, how to assess harm in that framework. For example, you have -- you know, it -- under the FTC Act, substantial injury cannot be reasonably avoidable, is not offset by benefits, right?

So all of the sort of areas that I mentioned, the idea that it can't be readily avoidable is a huge issue. This is absolutely impossible most of the time for people to avoid being in this database in the first place. It is not necessarily possible. Many, many people I spoke to had no idea that Equifax existed or had data on them, right?

So, I mean, the information asymmetries, the lack of a level playing field I think is absolutely crucial, and you cannot just sort of go past that and say that

129

that's not a part of the risk assessment. It has to be a huge part of the risk assessment, and I think the way to do that for the FTC is through the unfairness doctrine.

MR. CHILSON: Great. Do you guys want to put your cards down?

MS. DE MOOY: Oh, sorry.

MR. CHILSON: I don't want to keep calling on you.

MR. MANNE: I have more to say.

MR. CHILSON: So one thing that I -- tying together the responses to the two sets of the hypotheticals, while Paul openly admitted that he was pushing back against the hypothetical, I think pretty much all of you pushed back, which is the point of hypotheticals, and I was particularly interested in, both Michelle and Alessandro, you both said -- not in exactly the same terms but essentially -- like, this might not be harm -- and I think, Michelle, you actually did say this -- might not be harm, but it is a violation.

And so I am interested in teasing out why the difference there, and I think, Alessandro, you sort of laid out a sort of -- a boundary framework, that when you cross the boundary, that's a type of harm, and I

130

think it sounded to me a little bit like, Paul, you pushed back against that idea a little bit in the sense that injury is a sort of collective thing that we've developed over time and that law has a role to play in that.

So I just wanted to throw that out there for the panel, whoever is interested in talking about that, but especially Alessandro and Michelle, about why would we have violations where there aren't harms?

MR. ACQUISTI: Well, my two points on this would be that there are violations which may not arise to economic harm, but other forms of harm, autonomy or dignity or freedom.

Secondly, there is the increased likelihood of downstream harm, and then we can debate, as we did on this panel, whether the increased risk we felt materialized damage is enough for intervention or not. That's a fair point to debate.

And third, there are all these categories of economic harm which we know are there, but we find it very hard to quantify. So this is in a sense my argument.

And the last point I will make, if there is -- if there are a few more seconds, is that I would suggest that as important as this workshop and type of panel is,

131

I would suggest also a different workshop, a different type of panel, where the burden of proof is not put on a consumer, demonstrate that you have economic damage, otherwise, we should not -- we should not intervene to protect, but the burden is twisted around and put the data holders into the position of demonstrate that you cannot do these transactions you are doing now in a more privacy-protective manner. And if you do, and if you claim that there are costs of doing so, demonstrate -- tell us, show us -- where the costs go. Only to you? To consumers? To society?

So for the moment I am going back to the essential problem we are facing here, which is enormous information asymmetry at the individual level, because individuals don't know how information about them is being collected or used, and then societal level, because as much as we like to believe in data analytics, much of the internal data economy right now is a black box where we do not exactly know what is happening. We know that value is being generated. We don't know exactly how it is being allocated. That is to me a pretty crucial question that, as economists, we should address.

MR. CHILSON: Great.

Michelle?

132

MS. DE MOOY: I think that was well said. I think the only thing I would add is that the person -- you know, privacy is contextual, as we know, and, therefore, an individual's perception of the situation matters, and I think the way that the government can play a role there in leveling the playing field here is by assessing what are reasonable expectations, what types of user controls are available to this person, what sort of access rights do they have.

And then when we move down the spectrum of risk to economic or quantifiable harm, that's when you can assess whether remedy makes sense, whether there is a justifiable remedy. And I think that is logical and exists in so many of our laws, but for some reason, as Alessandro kind of pointed out, this is skewed in this environment as if the benefits of data collection are so great to consumers that it's ridiculous to think that there could be sort of violations to harm, but I think that's absolutely what occurs and I think has been borne out, you know, in example after example.

MR. CHILSON: Paul, we have some questions from the audience, so --

MR. OHM: All right, go ahead. I'll find a way to say my answer in response to whatever you ask me.

MR. CHILSON: Great. I trust that that is true.

133

So we have a couple questions here. Some of them have been somewhat addressed since I got them, so I am going to focus on one that has not been addressed yet, and I am open to you guys taking it in any direction, but a focus on harm would be particularly interesting.

The question asks -- considering the panel -- do you all accept the notion that privacy is a critical component of democracy, as Michelle stated? Kind of a big picture question, but if you can tie it back to harm --

MR. OHM: I totally have a segue to the point I was going to make.

MR. CHILSON: Let's start with Paul, then.

MR. OHM: Yes. So let me start small and I will end up at the question. So I wanted to respond to the Bayesian brothers, the idea that we are going to examine the delta --

MR. COOPER: I didn't know I had a brother.

MR. OHM: Yes. So what I find problematic about using that as the sole way of defining harm in a data breach case is it means if you are in a space where there really isn't much harm and there is a lot of responsible practice and then there is one really bad actor who is below the standard of care, then the FTC has jurisdiction, but then as the world goes to hell in

134

a handbasket and we end up in a cesspool, where all corporate actors, for whatever reason, you know, malevolent or benign, are not protecting us, are causing anxiety, causing the kind of fear I'm talking about, then suddenly you've stripped the agency of jurisdiction, that seems completely backwards and a little warped to me.

It feeds to a point that Alessandro has made a couple times, but I think because he's an economist and because he's a polite Italian, he hasn't made quite forcefully enough, which then leads to your question, which is -- Alessandro has repeatedly said that the economic tool kit can be very helpful when it talks about harm, but it should not be considered complete, right?

And don't mishear me, we still should be empirical and we still should be rigorous, but I think in many ways the economic tool kit is deficient when it comes to this. And I know I'm talking to an agency that happens to have a Bureau of Economics, that has people who helped put this workshop together. I think we need to look at other social sciences, we need to look at legal scholarship, and we have to understand, as you get into the very next panel, that sometimes it's going to be hard to kind of measure results that come from those

135

other fields with what the economists say. And if you're only looking to the economists, you're thinking of this to narrowly, which goes to democracy, right?

So the idea here is there are absolutely ways -- whether or not democracy falls within the FTC's core mission, I don't know if I'm ready to say it -- but there are ways of saying that when we talk about privacy harm, we are talking about broader societal problems, and Congress in its infinite wisdom said, look, the courthouse doors are going to be open or not to traditional tort law principles, but we are going to write a capacious, broad statute because we can't read the future, and we want to create an agency that can stand by the consumer today and tomorrow and the day after, and I think that's how they wrote their unfairness provision, and I think it's kind -- a responsible agency would take advantage of that and try and protect consumers in the way that Congress had in mind.

So thanks.

MR. CHILSON: Yes.

Geoff?

MR. MANNE: Well, I think absolutely it's the case that privacy from the government is essential to democracy. I think we have to always bear in mind that

136

we're talking about private entities here and our fellow citizens, and I think it's -- I think often -- not always, certainly, but often the two are allied and then they are extremely, extremely different in my mind. And we do a real disservice when we say something like, you know, the drugstore in the first example knowing something about me is -- not that anyone said this, but one could say -- is the -- you know, it's just as bad as the government knowing this about me, right? The next hypothetical after the insurer and the employer could be, you know, the government. And I think it's crucial that we keep those things separate.

With respect to the ability to keep information private as being - sort of from other people as being crucial to democracy, I don't even really know where to begin to answer that, and that -- and therein lies the problem with what Paul just said. No one knows where to begin to answer that, and Paul is right, that one could read an immense amount of discretion into Section 5, and we could have an FTC that supercedes every legislature in the country and every other statute in the country, and indeed, you know, you could say that's what Congress intended. I mean, you'd be wrong, but you could say it.

And the idea that sort of trying to implement some idiosyncratic principle like democracy at the level

137

of enforcement against real companies engaging in real commerce with real consumers who are, for the vast, vast, vast, most part of the time enormously benefited by that -- something else to remember when we talk about all this, we highlight all of the problems -- but as a practical matter, to me, they are kind of few and far between, really.

That doesn't mean we shouldn't care about them, doesn't mean we shouldn't do something about them, but let's not forget that they are the exception, not the rule, and -- anyway, in sort of authorizing an agency to say, well, we're protecting democracy, and, therefore, we should be able to do basically anything we want, without having a need or an ability to quantify it strikes me as so dangerous as to undermine democracy.

MR. CHILSON: James?

MR. COOPER: All right. Thank you, Geoff, for going on long enough that I might get the last word.

MS. DE MOOY: I will just reaffirm what I said before and say that I don't think that anyone was suggesting what you just said --

MR. CHILSON: Sorry, James raised his card. That's the only reason I moved past --

MS. DE MOOY: Oh, I see. Okay.

MR. COOPER: No, no, go ahead, and then I will

138

say the last thing.

MS. DE MOOY: You get the last word. I would just say, yeah, I don't think anyone was suggesting this lawless world of, you know, immeasurably damaging our democracy through the FTC's unfettered -- you know, I don't think anyone was saying that at all. In fact, I think the point of the FTC's involvement in privacy is -- first of all, it's an agency of the United States Government which is charged with protecting the Constitution, which, of course, is embodied by democratic values, and some of those include the space for political thought, the space for choice, the space for control, the space for deciding who can see information and who cannot.

All of -- a lot of the data ecosystem violates these principles in different ways, and so, therefore, we can look at -- deeply into those to figure out what exactly what makes sense for the FTC's role here. And, again, I think it has a lot to do with leveling the playing field, and that is a democratic principle, to not have information asymmetries dictate all of these practices and policies, but to have the level playing field where the American consumer can make a choice.

And just incidentally, I would also say that the distinction between government and commercial entities

139

is much blurrier than I think you were painting it to be. You know, the government is acquiring commercial data all the time. They are working with private contractors all the time. So I think, you know, the -- it's not possible to make that distinction, per se. I mean, in law we can, maybe, but in this discussion, I think, you know, especially when you talk about health data, which is more my area of expertise, it is absolutely -- the government is constantly selling and buying commercially generated information about people.

MR. CHILSON: James. Two seconds, James.

MR. COOPER: Yeah, okay, and Paul's wrong. No, I'm just kidding.

The only thing I would say directly to Paul is you had said that we -- the FTC needs to think -- you know, incorporate a lot of other things other than the economics. I think actually that's one of the issues, is you say legal scholarship needs to be incorporated. I think there has been very little, if any, economic incorporation into a lot of the privacy, if you look at the two privacy reports. So I think that moving away from the legal scholarship, more into empirical work, or at least balancing them more, I think that would be -- I think the balance is certainly more on the other side.

But the last thing I'll say, I agree with, you

140

know, the question about democracy. I would agree with Geoff, I think, when it comes to privacy. It's vis-à-vis the government, not really vis-à-vis private corporations, and I think the big -- you know, I'll leave it at this. I think one of the big picture questions here is, you know, I completely agree with the -- that there are information asymmetries here, and there are -- Alessandro's great body of work has shown, you know, a lot of, you know, this contextual dependence, a lot of biases and Dahlman effect exist in this, but there -- asymmetric information and behavioral biases exist across a lot of markets.

The question -- I think the big picture question here -- and I will just end on this -- is, you know, we think about what we want to do. What's better at mediating consumer preferences in this case, the market or the government? And I think that that -- the more it's informed with empirical literature, I think the better. So I will just leave it at that.

MR. CHILSON: Yep. Well, thank you very much to our panelists, and thanks to all of you. I believe up next we have lunch.

(Applause.)

(Whereupon, at 12:46 p.m., a lunch recess was taken.)

141

AFTERNOON SESSION
(1:46 p.m.)
PANEL 3: BUSINESS AND CONSUMER PERSPECTIVES

MS. HAN: All right, if folks could take their seats, we will get started with this afternoon.

Okay, everyone, welcome back. Dan and I will be moderating this next panel, which will build on this morning's discussion and explore how businesses and consumers perceive and evaluate the benefits, costs, and risks of collecting and sharing information in light of potential injury.

The panel will examine the considerations businesses take into account when choosing privacy and data security practices and also how consumers make decisions about sharing their information. So we are lucky to have a great group of panelists here with us for this discussion, and we will take questions at the end. So just as a reminder, there are comment cards available in the hallway and also with FTC staff inside the auditorium. If you just fill it out, raise your hand, someone will come and get it from you. For those of you viewing the webcast, you can submit questions via Twitter.

Now I would like to introduce our panelists. So closest to me is Omri Ben-Shahar, who is the Leo and

142

Eileen Herzel Professor of Law and Kearny Director of the Coase-Sander Institute for Law and Economics at the University of Chicago Law School. He teaches contracts, sales, trademark law, insurance law, consumer law, e-commerce, food law, law and economics, and game theory in the law, and writes primarily in the fields of contract law and consumer protection.

Next to Omri is Leigh Freund. Leigh is the president and CEO of the Network Advertising Initiative, where she leads the organization's growth and helps with the agenda and strategic priorities. Leigh joined NAI in 2015, after an 11-year career at AOL, where she served as vice president and chief counsel for global public policy.

Next we have Jennifer Glasgow, who has served as a global privacy and policy executive for over 40 years, originally with Acxiom and most recently with First Orion Corp. She is very active in numerous international efforts to develop effective public policy with maximum harmonization across the world.

Then we have Bob Gourley. Bob is cofounder and partner of the cybersecurity consultancy Cognitio, which helps companies fight cyber crime and corporate espionage. He is the author of the book The Cyber Threat. His first career was as a naval intelligence

143

officer, and he was the first director of intelligence at the Department of Defense's Cyber Defense Organization.

And last, but not least, we have Katie McInnis. Katie is a policy counsel in Consumers Union's Washington, D.C., office. Her work focuses on technology and the consumer's right to privacy, security, control, and transparency. Before joining Consumers Union in 2016, Katie served as a privacy and technology fellow at the Center for Democracy & Technology, and in the Enforcement Bureau of the Federal Communications Commission.

So thank you again to our panelists for joining us today.

MR. WOOD: Okay. So the format of this panel is basically going to be a loosely organized group discussion, and we are going to start off with some business-oriented questions. The first one is pretty broad.

So what are the risks and benefits businesses consider when deciding whether and how to collect and share consumer information?

MS. FREUND: Good afternoon, everybody. Thanks for having us on the panel this afternoon. We are thrilled to be here. I represent the digital

144

advertising industry here, and we represent third parties, first parties, and other companies involved in the digital advertising ecosystem. So when we think about the benefits and the risks and how companies are weighing the benefits and risks of the use of data that's collected online, you know, we think through kind of the entire internet ecosphere.

So, you know, revenues from online advertising support the free internet. I know that's a very broad statement, but we support and facilitate e-commerce. We subsidize the cost of content and services that consumers really value and expect, and this is a really, really valuable kind of benefit.

We have to weigh that with the risk of what are we doing with consumer data and what are consumer expectations, which is something I think we will get into in a little bit, but, you know, we have got just some statistics to share with you.

Eighty-five percent of consumers in a survey that was conducted by the DAA, which is one of our fellow trade associations, say that they prefer kind of an ad-supporting internet model so that they don't have to worry about paying for costs and services, and so, you know, online marketing has really big, direct, significant benefits for consumers. It helps them

145

connect, create, publish.

You know, we are talking not just about Google and Facebook when we talk about the internet, but we're talking about my favorite hypothetical example, joesknitting.com. We can support advertising that supports joesknitting.com to be able to reach their consumers because the advertisements are reaching individual users as they travel across the internet and not individual websites. And so we support kind of the long tail of publishing online and contents and services that consumers might not otherwise be able to access.

So when we think about risks, however, part of our genre is self-regulation in this internet advertising ecosystem, and so we think about what are the specific, concrete injuries that cause or are likely to cause kind of substantial injury or harm to consumers. You know, we think through the Section 5 framework about how things are not reasonably avoided by consumers and what are the benefits to consumers.

So we think through data practices that are reasonable. We try to think about the types of injuries that might be quantifiable and economically harm a person. We think through some things like eligibility for employment. We think about eligibility for insurance requirements, and now we collect and use a lot

146

of data in advertising, and the serving of a targeted ad to you or to me because of our interests online is different and so weighed differently than the actual use of that data to make a decision around eligibility for employment or other things.

And so we've got guardrails set up in our own environment, in our self-regulatory code, that kind of thinks through these potential concrete injuries and puts restrictions on our company's use of that data. So when we think about benefits and risks, we really try to spend some time thinking specifically about how to avoid the risks that might be inherent.

We heard about some of them already earlier, in the earlier panels, and we try to put guardrails around them so that we can have responsible but vibrant data collection and use throughout the internet.

MR. WOOD: Great.

So, actually, when you -- to make things easier, let's raise our name cards when we -- if you want to say stuff.

Jennifer, please go ahead if you were going to. Sorry.

MS. GLASGOW: You want me to answer the same question, right?

MR. WOOD: Sure. Yes, yes.

147

MS. GLASGOW: Okay. Well, my experience is in both marketing, so I'm very familiar with what Leigh and NAI have done, but also in risk management and government use of data. So I feel like, you know, when businesses are trying to assess the risks of information collecting and sharing, they tend to start by saying what are the benefits, what are the opportunities that information creates?

And the reality is -- and you probably got a sense of this from the morning panels for those of you who were here -- the reality is that every piece of information that you collect has the potential for some benefit, and it may be different. It may be many, many benefits. Very rarely does a piece of data stand alone and only create value in one instance.

However, there's two sides to that or that sword is sharpened on both edges, because there are risks, as Leigh has mentioned, that need to be taken into consideration. We look at laws first to decide, you know, is whatever activity we're doing compliant? Sometimes that's an easy call, sometimes it's a hard call, because many of the laws -- I'm most recently working in the telecommunications industry, and we have laws going back to 1934. So none of those laws even contemplated the environments we're living in today, so

148

you have to kind of read between the lines and figure out what you should be doing and what you shouldn't or what the intent of the law was.

And then we have a lot of different self-regulatory groups, like the NAI and others, that focus on very specific, typically either activities like marketing or industry sectors. And I think it's a good balance, because the industry groups can move quickly when new issues surface or when new technology or new business practices evolve, and, if necessary, you can back them up with law on something that's really egregious or that industry is not able to get wide adoption for.

So the risks, you know, most businesses today, I think certainly those that have any kind of international presence, with maybe the exception of some small ones, go through privacy impact assessments on what their business is involved in, and this -- and any assessment, you read any of the standard guidelines for a PIA, say you have to look at the risks and you have to mitigate those to the degree that it's possible or to the degree that it makes it acceptable, an acceptable level of risk, versus the benefit.

MR. WOOD: Okay, great.

Katie?

149

MS. MCINNIS: We know that businesses evaluate the risks and benefits of data collection, but we also know that they exhibit the normal human tendency to overestimate the benefits of data collection in comparison to the risks, and we have seen that through many data breaches recently, that companies aren't sufficiently internalizing the risk of all this data collection when they evaluate the benefits of such collection.

MR. WOOD: Great.

Bob, do you have -- as a cyber security expert, do you have a different take?

MR. GOURLEY: Well, I would say that answers to questions like this I think vary from industry sector to industry sector depending on where you are at. For example, in the financial sector, those companies are all built around trust, especially if they are consumer-facing, so they take an approach where absolutely everything is evaluated from risk. Compliance is critically important to them. Any heavily regulated industry, compliance is extremely important. We all know compliance does not equal security, and compliance does not automatically reduce all risk of data loss, but compliance is extremely important, because if you fail, it can be a company-crushing event.

150

So I would say that my approach to questions like this is to look industry by industry.

There are other industries where you're collecting information that is already publicly available. You're just pulling it together and using it slightly different. So if that information is lost, is that any risk? Does it hurt your brand at all if it was all publicly available? And so the kind of question like this, I think it's very important to figure out what's the business model in the particular industry, what are best practices for that industry, and then what's the right approach.

MR. WOOD: Okay. So is reputational injury important in advertising and privacy considerations as well?

MS. FREUND: Oh, absolutely. That's why we're here. I think, you know, what we represent is, you know, people like Jennifer's former company and others who came together recognizing that there was -- with all of this data comes great power and great responsibility, to use a Superman analogy, but I think that thinking through -- it wasn't Superman?

MR. GOURLEY: Spider-Man.

MS. FREUND: Spider-Man, sorry. I have my superheroes wrong. You can tell my kids are out of the

151

house already.

So I think that we have got a group of companies that are dedicated to responsible data collection and use practices because there's a need for that. Consumers, you know, have said that there's a need for that. Industry and regulators have recognized that there's a need for that. There has been some talk of legislation with no actual legislation enacted here in the United States, at least, but we think about Europe and the data protection legislation enacted there.

And so, yeah, I think in -- especially in advertising, when you've got relationships with consumers by first parties, but also the very backbone of that advertising community is made up of companies that are unknown to consumers, so it carries a bigger responsibility. So because you're not necessarily known to a consumer and you're collecting and using their information, you've got to put certain safeguards in place.

So we have lots of safeguards in place through our self-regulatory code. Some of them I mentioned. You know, as the data gets more sensitive, the restrictions get more onerous. And so, you know, we have, for example, a definition of sensitive data within our code, and that covers a lot of some of the things

152

that were discussed in the earlier panels with respect to health data. You know, some of those, as the sensitivity increases, the requirements for what you may and may not do with that data and the types of consumer notice and choice that you put around that is really important.

And so we base our code on the very basic principles of privacy, on the FIPPS, and it centers around notice, choice, accountability, and control. And so in advertising, we think this is really important. I think, you know, when it comes to, like I said, the use of data to serve me a Nordstrom shoe ad -- just hypothetically -- is less injurious to me than the concept of somebody using, you know, something around HIV status or gender identity.

So we've put guardrails around those things in our code to try to make sure that folks understand that they can trust -- I think trust in advertising -- and I know a lot of what folks have mentioned on earlier panels is the complexity of this environment, the data environment. I'll use an analogy that is not my own, but I'll use it anyway in terms of developing trust.

So you fly on airplanes every day, but you don't necessarily know every word of what those mechanical manuals say about how that airplane flies. You just

153

have to trust that the people that do know what they're doing with it. And so part of our purpose and mission is to hope that consumers will understand that people that know this very complex digital advertising ecosystem and the data that's around it will do responsible things with it, and by adhering to our code and the code of the DAA and other self-regulatory orgs out there, we think that we, you know, are striking that right balance between innovation and protection.

MR. WOOD: Okay. So do businesses -- I guess you've addressed this, Leigh.

Other panelists, in your experience, do businesses consider directly the injury to consumers?

MS. GLASGOW: I'm sorry, do they consider what? There's an echo.

MR. WOOD: Sorry. Do businesses directly consider the risk of injury to consumers?

MS. GLASGOW: I think they do, but I think it's really difficult for them to understand what those risks might be. Certainly the ones that I've been involved in, whether it's Acxiom where I was an employee or whether it was clients that we were serving, would ask the question, what are the risks? And I think from this morning's panel, if I hadn't said, either I as a privacy officer for the company or the company themselves

154

probably would have not recognized over half of the risks that were actually -- or the consequences that were actually discussed this morning.

So I think there's an opportunity to continue to educate the business community about how serious some of these risks really are, because most of them, at least in -- and I totally agree with Bob that it does vary to sector to sector and how regulated or how -- what kind of compliance obligations fall on the company, but they -- it does -- it is one of those situations where we need -- we need to understand all the potential risks, and they're evolving.

We have new risks that surface, you know, with new technology, and maybe the FTC, in conjunction with industry, could do some more with that. IAPP, the International Association of Privacy Professionals, would be a great audience to carry those messages to the businesses.

MR. WOOD: Okay.

Katie?

MS. MCINNIS: I just want to jump off of what Leigh was talking about with the differences between different kinds of information and the kind of sensitivity we ascribe to them. I think that what's getting complicated in this space is that with increased

155

data collection, nonpersonal identifying information is now quickly becoming personally identifiable information. So although you may not know my gender identity, you may know what kind of products I'm looking for online, and therefore can make an inference that I have this gender identity. Or maybe you don't know my HIV status but you know that maybe I'd be in the market for an at-home test. So it does get very concerning, this huge collection of data and the kind of inferences you can draw from it, and I think we need to be aware of the kind of data that you're collecting can really be highly sensitive even if the individual data points don't seem sensitive on their face.

MR. WOOD: And do you think businesses are inadequately aware?

MS. MCINNIS: I think it's hard to tell. I'm not a business. I haven't worked for a business. So that's a little hard for me to tell personally, but I do know that with the number of data breaches that have occurred, especially in industries that do not have a direct relationship with consumers, Equifax is a great example, we know that there's a huge amount of data about individuals that they never really opted in or gave permission to be collected about them, but it can have a really injurious effect on their livelihood and

156

also some clear informational injuries beyond just financial injury.

MR. WOOD: Okay, great.

MS. FREUND: Could I quickly respond? Oh, sorry, Jennifer had her card up.

MS. GLASGOW: Yeah, I want to make a point that I think we need to think about data breaches in one light and then we need to think about inappropriate use of information by the company in another light, because both of them can be harmful, and quite often we conflate the two, and I think resolution and mitigation for those kinds of bad practices may differ dramatically between the two categories.

MS. FREUND: Yeah, so Jennifer took my answer.

MS. GLASGOW: Sorry.

MS. FREUND: But I will add to that that I think the concept that Katie rightly brings up is, you know, what data do you need to do the action that you are doing? You know, the less data you have, obviously, the -- the less risk of data breach, and I agree that we need to keep them kind of separate and distinct, but, you know, in our case, advertising data is only valuable in very limited circumstances and for very limited time periods, so one of the things we counsel our companies or our members to do is, you know, practice very good

157

data minimization. So if you don't need the data and you don't need it any longer because the person's bought the car or bought -- in my case bought the shoes -- you know, get rid of it.

So that is something that I think, you know, good data responsibility and good data collection practices and use will -- will engender that concept of data minimization, and I think it's important.

MR. WOOD: How -- how -- do businesses in general seem like they're cognizant of data minimization and how important it is?

MS. GLASGOW: I think they're aware of it. Again, it is going to vary with the industry and it may also vary with the size of the company. It's something that they talk about. If there's any kind of international footprint, because those are more regulated kinds of activities in other geographies, it becomes more important to them, because they know it's important in Europe, and it's easier to just deal with it worldwide than it is to do things differently in one country from another.

So -- but I want to caution against data minimization being too much of a panacea. It's -- knowing how long you need the data and how long it's useful is important, but some data has a lifetime of

158

value, and so, you know, the idea that all data has an expiration date that's fairly short I think may lead you into a situation where you're not maximizing the benefit you could bring to it.

For instance, if you want to contrast marketing, which as Leigh has said, has a much shorter lifetime of value, to identity theft, where if I have a pattern of where you've lived for the last 20 years, I can do a better job of predicting that you are the person you claim to be than if I have two years of that data. So it can be very variable.

MR. WOOD: Okay.

MR. GOURLEY: I would like to push on that just a little bit, too, and say I have not seen a correlation between how much data you have and the likelihood that it will be breached. Your data needs to be secure even if it's a very small pile of data, and a lot of times a small pile of data in a small company is actually of greater likelihood of breach because that small company can't protect it as well. I would say what would GAFA do? We have got to watch what GAFA is doing, because they are leading us all forward. Who is GAFA? What's GAFA? GAFA is the big industry players when it comes to data. Google, Apple, Facebook, and Amazon. GAFA is leading us all into how we are going to do privacy in

159

the future, and they're moving a lot faster than I believe the government or even consortiums can, and they're adopting things like very smart encryption, ways to stripe and secure data and spread it around places where it's harder to penetrate, not perfect, but harder. They're also advancing cryptologic concepts called, like, differential privacy. Have you guys heard of differential privacy? It's a cryptographic construct that has been around for several years. It was now brought into Apple's iOS 10, so we are all using differential privacy now, didn't even know it. It's in Apple. Facebook is using it. Google is using it. What it does is it allows you to extract data from data sets in ways that do not reveal the underlying data but give you statistically meaningful information. A simple way to think of it is when you use the traffic on your Google Maps and it says that it's red on this road, it knows that because they're watching the traffic, and they do that in a way that does not take away the privacy of the individuals and how slow they're moving. So my bottom line point to all this, watch GAFA. Watch where the big guys are going with privacy and protecting data. I think it's extremely important. So the question to ask yourself, what would GAFA do?

MR. WOOD: Okay. So we've heard that various

160

factors businesses take into account in thinking about the risks and benefits, so maybe GAFA, the trust that the consumers are placing in them, consumer expectations, compliance, and regulatory issues, and sometimes direct the sort of concrete injuries possible to consumers.

How well do these factors businesses consider correlate with potential injury to consumers?

MS. FREUND: Could you --

MR. WOOD: So how well do the various factors businesses think about when they're weighing whether and how much data to collect and share correlate with the potential injury to consumers?

MS. GLASGOW: I'll start it off. The -- I think this is where I want to answer it -- answer the question first for -- from a security breach perspective and then from an appropriate use of data perspective.

From a security perspective, I think they correlate pretty well. We're now looking at all kinds of security related to data at rest, data in transit, and so on. So -- and to pick up on the GAFA analogy, the fact that we're using cloud computing, where many of the security precautions that a small company would never be able to put in place are made available, is actually a really positive step forward on the security

161

side of the house.
Now, that doesn't help at all with, you know, internal use of information and whether it's appropriate, and I think that we're beginning to realize that as data becomes part of every business practice, regardless of what it is, and much of it is personal data, whether it's anonymized or not, it -- it has both opportunities and risks, and those need to be evaluated as carefully as possible.
I used to say that, you know, we could put some guidelines out there, whether they were self-regulation or legal guidelines, that governed or they gave companies guidance about what they should and shouldn't be doing, but as information has proliferated, everything we do and the opportunity, the analytics that we have to turn it into valuable insights, it gets very easy to do, not just -- not just available to large companies with sophisticated data scientists.
We have surfaced a different issue, and it is every single use of data in a company needs to be looked at through the lens of a privacy impact assessment, and that should include the risk to the individual as best they can assess them. I -- in work -- if you look at, say, an industry such as the risk side of the house where you're doing identity verification to try to

162

understand if you're dealing with legitimate people or may be required by law, as it does with Gramm-Leach-Bliley, where there's a know your customer rule, to actually verify that you're dealing with who you think you are, those are industries that really pay a lot of attention to these kind of issues.
I find that some of the tech industries and some of the startup industries are not as consistent, if you will, in their evaluations of the internal uses of data outside of security.
MR. WOOD: Katie?
MS. MCINNIS: Data security should be in line with what a consumer expects and the thhe business is really putting a high priority on; however, the number of data breaches we've seen even just in the past five years show that that isn't significantly guarded, right? We have Aetna, Uber, Yahoo, Equifax, so many companies have not sufficiently protected the data that they were entrusted with from consumers, and that's highly concerning, and that should be their priority, and that's what consumers expect, yet we're not seeing that being fulfilled.
And as far as misuse of data or privacy of consumer data, we're also seeing that that's not necessarily a priority. One good example is Uber's

163

recent breach, right? They had a breach. They worked to cover it up. And then they also misused some internal data in ways that consumers felt was really, really creepy, and led to some changes in their policies.
We see that change in their policies because they are consumer-facing and they depend on the consumer trust. So, yes, both of these things should be important to companies, especially data security, but we're not seeing sufficient protection of either.
MR. WOOD: Okay.
Leigh?
MS. FREUND: Yeah. So just a couple quick things. I mean, one is when we think about our industry in general, I think, you know, it's important to kind of keep in mind that without the consumer clicking on an ad, you know, none of this is here and none of this works. So obviously keeping the consumer kind of top of mind as we go through these processes is important, whether or not we do it enough is a matter for debate, but -- and I do think it's important to think about, you know, Jennifer mentioned, you know, it's harder for some of the smaller companies, the startups and the smaller tech companies to kind of put privacy top of mind due to either resource constraints or other things, and I will

164

say, you know, just a little plug for us that the compliance process that we undertake with our member companies, big or small, is the same process for every company, every year, on an annualized basis, and it helps make them put a priority on privacy that wouldn't otherwise be there, and I think, you know, we have -- Anthony, my compliance person -- VP is sitting right here, and he can tell you the rigorous and onerous process that folks go through, but I think it's important to think, when you think about consumers and putting them at the corner stone of things, priority -- the priority of privacy becomes really important, and so we build things around it that help us, you know, make sure that even the smaller tech companies and the startups put priority -- put a priority on privacy.
MR. WOOD: Okay, great.
MS. HAN: Great. Thanks. So, Leigh, I want to build on something that you just said, putting consumers sort of at the front of mind, and my question for you all is do businesses think consumers are informed about the benefits and risks to consumers from collecting and using their personal information?
MS. FREUND: Yeah, I think that is something that folks will probably disagree with on -- even among our panelists here, but I think, you know, one of the things

165

that we -- we do, at least for our industry, is make sure that our members are aware that one of the cornerstones and pillars of ourself-regulatory program is consumer outreach and education. That's part of what they're required to do as part of our compliance process, et cetera.

So we try to make -- you know, through either our own organization's efforts or through our member companies, try to make consumers aware of the fact that they do have choices about what happens with their data and try to bring folks to our website where they can learn more about what this digital advertising is, why am I being targeted, why is this ad following me around. We have some educational, you know, components of our website that try to help people understand that a little bit more and help understand what the options are, if they choose to exercise them.

I think, you know, in general, we have found, through some research, that consumers are pretty -- are becoming more aware. I think this is where self-regulation has an advantage over some of the other -- the other types of things that we could do, like regulation or legislation, because we can respond to things a little bit more flexibly and quickly as new technology is developed, but I think consumers'

166

expectations are also evolving in this area quite a bit. You know, I think about, you know, what we know and understand in the industry, but also about, you know, my parents living in Michigan, who didn't grow up with this technology -- I think they're listening, hi, mom and dad -- and what their expectation is might be completely different from the folks living in Silicon Valley or the folks living here, working in this industry.

So it's really incumbent upon the companies to try to use those efforts to make sure that they are understanding -- that they're trying to educate the consumers about what their choices are and what type of data they use. I mean, privacy policies, we can argue that those aren't necessarily as informative as perhaps they should be, because they're buried, but trying to think of proactive ways to make consumers informed is really important, at least to our companies.

MS. HAN: Thank you.

Omri?

MR. BEN-SHAHAR: Yeah, I would like to take an issue with that. I think I'm deeply skeptical about education, transparency, privacy policies. It's a nice kind of thing for the companies to brandish and for the law to require, but it has no effect whatsoever. We know that for a fact from numerous studies, and there is

167

no way to make it work either, in part because we have to realize we're sitting here talking about privacy, but there are panels around the country and around this city and around -- where people are talking about other aspects of the consumer transaction that consumers care and some law makers and some advocates care not less than data privacy, for example, the choice of forum, the arbitration clauses, and the ALI cares a lot about the warranty disclaimers and the way laws are being completely disclaimed by the -- so it's hard to educate consumers about everything and especially about, I think, data policy, because it's a moving target. It advances so fast that by the time people catch up, they are already two or three years behind.

That said, you know, I would think that if -- that doesn't mean consumers don't largely understand what's going on. I think maybe they -- there is some indication that they do. I don't want to -- I can't say that as a fact, but that they -- people say in surveys that they are concerned about what's going on. That means that they think that the data collection and data usage and sharing is done in ways that they can't pinpoint the details, and they will not be able to understand, but it causes them to -- them concern. So I think that there is a sense of the risks that comes

168

maybe from experience and from some other contexts, not from privacy policies.

The other thing that I want to say about that is that I think consumers view -- return to that basic point, that consumers view data policy as one aspect of the transaction. The data collection and practices of firms is -- in part, consumers view it as a good thing. Privacy or information is the new money -- I've said that before -- and people are delighted to pay with information rather than with money. I mean, they probably would prefer to do neither, but I think that there is a general understanding as well that there is a grand bargain here where information becomes the currency.

Also, personalized services are largely good. People enjoy them. There's both a private -- a great private benefit. I mean, this is a symposium about informational injury, but I think from a social point of view, we really also want to think about informational benefits. There are -- Catherine Tucker, who sits here and will talk later, did some work about the value of personalized service, about digital medical records and the life-saving effects that they have.

I've looked at some work that is basically -- you know, digital records and data collection in the auto

169

insurance industry and these pay-as-you-drive arrangements that are, by the way, prohibited in places like California because of privacy concerns, and the value that they can bring, under some estimates it could be as much as 20 percent reduction in accidents. When the insurance company knows how you drive and when you drive and how much and where and how abruptly you stop and things like that, and it changes the premiums accordingly, that may be a privacy issue, but it also leads people to drive in a safer way. One estimate that I saw is a reduction of up to 20 percent in auto accidents. That's, like, 3000 lives a year, I don't know, maybe more. That's -- there must be a very large privacy concern to override that benefit, and I'm not even talking about the fact that people whose pay-as-you-drive habits are measured, drive less, and another estimate is about 8 percent less, and that is a reduction equivalent -- that's an enormous reduction equivalent to about $1 of carbon tax, so there is an environmental benefit.

There is also a benefit here to low-income drivers. They usually drive less, so they will get lower premiums. All of that stuff comes -- are benefits that come from data collection that the consumer, if they enjoy them, are enrolling into. They like these

170

things. That's the context of what the people realize.

Of course, there are the bad things that people are not -- either not aware of or when these things hurt them, and I can't say that this is considered by consumers as worse than other "fine print things," like termination fees in cell phone contracts or warranty disclaimers or things like this.

So consumers are -- to wrap up, consumers are aware that this thing is going on. They are not aware of the details and how advanced some of the collections are, but not concerned enough to buy privacy shields that are not that expensive and to protect themselves from these things.

MS. HAN: Thank you.

So this is a great transition for us into sort of the consumer perspective and how consumers weigh the risks and benefits in determining how to share their information. So, Jennifer and then Katie.

MS. GLASGOW: Yeah, I'd like to pose that we have -- consumers have things they expect to happen, and I think security certainly falls into that category. We don't give them a choice about security. That's a must-do. Whether we do it well or not is, you know, something we can debate.

But then there are things that we should give

171

consumers choice about, a la, online advertising, as Leigh's organization has put forth, but I tend to be a bit skeptical about their understanding, and I think we need to separate notices and policies from understanding.

I've been knocking around this information world -- I started out on the technical side of the house many, many years ago, and I find that even as a professional in it, it's sometimes very tedious for me to understand what-all is happening, and I think as we move into more big data applications and more analytics, where decisions are being made by the analytics engine, it's going to get even harder.

So I feel like we're going to have to get industry to rally around setting guidelines, like the marketing and digital advertising industry has, that industry has to follow and that we don't have to ask the consumer a choice about, because there's a lot of research that shows that when consumers are given a choice, if they aren't sure about what they're really being asked, they take no action. So that means the default becomes extremely important in terms of is this going to be something that's allowed or is this going to be something that happens and they have an opportunity to stop?

172

MS. HAN: Thanks.

Katie?

MS. MCINNIS: So I do think that there is an understanding that consumers have. To a certain extent, they are trading information for a free service. I think people do generally understand if they sign up for Facebook, they're trading a certain amount of information away in order to use this free social platform.

However, I don't think that that's always the equation here, especially when you have a number of unknown data brokers and other companies online that are collecting my data at a rapid rate. I have no understanding that I'm getting anything from them by their collection of my data, nor is there any benefit to me.

And I also want to challenge the assertion that people want personalized services. Yes, to some extent, I would benefit from personalized services from some of the organizations and companies that I interact with, but on the whole, I do not want targeted ads across my web service, which is why I'm one of the many users of ad block services.

And the increased use of ad block services is showing that people do not really want this supposed

173

personalized service that we so think they want. The personalized service also is clearly not about me but about making revenue on the other side, and I think consumers do understand that.

And as far as how much consumers are understanding the kind of disclosures they agree to, I agree it's incredibly hard, and there is a lack of complete understanding, in part because a lot of it's happening on the back end, which is why you have even like Facebook that does a lot of work to make sure that you know what kind of people are seeing your data.

People are still asking if Facebook is listening to them. So there is this -- this gap of understanding of, oh, I have a lot of contacts, and, therefore, that's why they're recommending a client that I saw last week on Facebook.

So there is this gap in understanding, and I don't think we should really always put that on the onus of consumers, right, and that's one reason why other intermediaries like Consumers Union is here to try to help evaluate the kinds of policies and disclosures that are out there and help consumers decide which companies and products to use.

MR. WOOD: Any other comment?

MS. HAN: Leigh, then Bob?

174

MS. FREUND: Bob, please go ahead.

MR. GOURLEY: First of all, just like so many other people, I love Consumers Union. I'm so glad that you guys are doing what you are doing because we all need that. And I also want to say that I feel ignorant after studying these things for decades, and I realize I don't have a full understanding of what it means when my data is breached and lost.

Sometimes I get angry and I know it's wrong, like the OPM breach, stealing my data, that was bad, horrible, but it had no real cost to me, or the Equifax data breach, I got angry like 190 million other people because the data is out there, but as I thought about it, I realized that had zero impact on me. I locked down my all credit records anyway. It had a big impact on that company, because they are in the business of selling my data, but that's not my business. It didn't hurt me at all. My data had already been stolen, and my Social Security number is out there. Why? Because I was in the Navy, and every year that I got promoted, it was in the Congressional Record, here's Bob Gourley, this Social Security number, achieves this rank. So with that data being stolen, it has zero impact on me. So how am I supposed to quantify that? And for me, I don't want any of my data stolen, but I have no way of

175

understanding, if there's a big breach, what the cost to me is. I do want my privacy protected like everybody else. I mean, what I do is my business, but I just -- I am not quite sure how to quantify or put a number on the impact of a breach.

MS. HAN: Thanks.

Leigh?

MS. FREUND: Yeah, I just wanted to touch quickly on the ad blocking issue. You know, folks that don't want targeted advertising, you know, many of them download ad blockers which actually block the advertisements that are on your site in general. You know, I would say the answer to that is go to the NAI website and exercise your opt-out, but I will say that we've done a lot of work through some coalition building in our industry on the concept of ad blocking, and we're realizing that privacy is not the top reason that people are downloading ad blockers.

It's not that they don't want targeted ads. It's that they don't want their data to be used. It's that they don't want a terrible user experience. And I'll suggest that, you know, when we use targeted or behavioral advertising and try to use data minimization, to not use as much data as we need to do that, but that it actually creates a better online experience for

176

users, you know, to get the same economic value out of a targeted ad. When you're not using data, you would have to have 20 popovers, 20 popunders, and some flashing jiggly belly ads on the page in order to try to make up that revenue for the publisher.

So ad blocking is harmful to that kind of continuation of free content and services on the internet, and the reason people are using it is not necessarily privacy-related. Although it is one of the reasons, it is kind of down the list.

MS. HAN: Katie?

MS. MCINNIS: I just wanted to touch on something that Bob highlighted, which is these injuries to consumers are highly contextual, right? You have an incidence in which OPM, an organization you knew about had your information, was breached, and you already knew about this and, therefore, acted proactively and froze your credit reports with the three major credit bureaus, which is great.

However, for a lot of people -- first of all, they didn't have a relationship with Equifax, so they were surprised not only that their data was breached but also weren't really sure how to handle it; and secondly, now they had to go take the step you did a while ago, and, therefore, they don't know what kind of future

177

injuries could happen to them, in addition to the fact that a lot of this information is immutable, like a birth date, et cetera.

So the injury maybe hasn't happened now and it's hard to quantify, but it could happen in the future, which is one reason why it's so hard, especially for regulators and other organizations, to really analyze and understand what kind of injuries these cause to consumers because it is so highly contextual.

MS. HAN: Omri?

MR. BEN-SHAHAR: Yeah, thank you. I'd like to touch on some of these issues that Katie raised. I think that, you know, it's -- the question to ask all of us, what I said and others, what do people care about? You know, I think that that's the key issue, not what do people say they care about, but what do they actually know when they are -- when there are stakes on the line, what do they care about?

It really shouldn't -- I don't think the FTC should listen to privacy advocates or privacy skeptics, to alarmists or deniers on either side about saying, hey, I download ad blocker. People -- you know, that's -- you know, we are not representative of anything.

Also, I don't think that the FTC should listen

178

very much to what people say in surveys they care about, because they will say they care about anything, you know, and there's a lot of things that just don't add up. The question is, what do their behavior shows when there are stakes -- real stakes on the line, when they have to make hard choices, what to spend money on or what to sacrifice or what burden to take, what kind of ads -- popping ads and all sorts of these things, flashing, nonpersonalized ads, to suffer through?

So as to not be personalized, then maybe they will say, you know what, maybe I don't care about it that much. I think these are the questions that determine whether there is an injury in a sense that is meaningful. Otherwise, it's all kind of arm-waving about either, you know, people -- either biased people that do not represent anyone or people say things because, you know, the context they were asked.

They say, oh, yeah, by the -- since you asked about data sharing, sure, I care about it. You know, what does that mean?

MS. HAN: Okay. So I wanted to build on both Omri's and Katie's comments. Oh, actually, Katie, did you want to --

MS. MCINNIS: Yeah, I just wanted to quickly respond, if that's okay.

179

MS. HAN: Sure.

MS. MCINNIS: Yeah, I take your point that survey data isn't indicative of what a consumer's actually going to do, and I hear your logic behind that, and I just want to counter that maybe there's something else going on here. Maybe consumers just don't feel like they have the tools to effectively make sure that their preferences are appreciated across many devices and services they use, which is one of the reasons why the FTC launched their IOT -- their IOT contest last year, was to try and come up with a way that consumers can control the different devices and services they use within their home to make sure their preferences are respected across platform and device.

I think that consumers feel like they don't have a lot of tools, which is -- to be fair, the industry is hugely fragmented. These policies are really, really long and hard to compare, and even when you are in the market for a privacy protective service, like a virtual private network, it's sometimes hard to know what kind of services you're actually receiving, and CDT's recent work and complaints to the FTC highlight that.

Even if you're presented with these assertive statements of what is going to be protected, you're really unsure if that's actually going to be followed

180

through on the back end, and that's one reason why these policies and privacy statements are so important, is because that's in some ways the only way we know what companies are doing with our data and information and the only way to effectuate those choices.

MS. HAN: Thanks.

So I wanted to back up just a second because Omri and Katie both talked a little bit about context, and I wanted to give the other panelists -- and also them, if they have additional thoughts -- to weigh in about how the context or kind of data being shared matters when you're thinking about how consumers are making those decisions.

MS. GLASGOW: Do you want me to start?

MS. HAN: Go ahead.

MS. GLASGOW: I don't think, other than just sharing sounds scary, I don't think most consumers really think about it, because I don't think they understand it. Data sharing goes on in all kinds of ways in every industry. I mean, if I think about the doctor sharing data with a specialist they've sent me to, and they send your regular medical records over or whatever, and every industry has a different kind of sharing practice.

It may be with a third party who's out -- you've

181

outsourced work to. It may be with a delivery service that's going to deliver the product to the consumer. I mean, it goes on all the time. So there are some things that the consumer touches and feels and understands, that FedEx is going to see the package that I ordered from Amazon or wherever, but when it comes to what goes on behind the scenes, I think they are pretty clueless.

MS. FREUND: Yeah, if I could add on to that, I would just say, you know, I think that consumers probably have a different understanding. I think, you know, when we get consumer complaints to our site, you know, we get, you know, I got a targeted ad, stop using my Social Security number.

You know, so the concept of what kind of data people have, I think the "sharing is scary" comment is really important, because I think consumers feel like, you know, for instance, if you get a targeted ad, you must know my Social Security number, my address, my credit card, and my Nordstrom frequent shopper card when, you know, really it's a binary decision made by usually an algorithm that says, do I serve this consumer this ad or this ad?

And when you're talking about the concept of injury, you know, on top of that, I think it's important to remember, you know, consumers -- and I think it's

182

important to try to continue to educate consumers as much as we can, but I think that's where the airplane analogy that I used earlier and the trust factor kind of factors in. If you're on a website or using a provider that you trust or that you feel like you can exercise your own choices with respect to that, then you're going to have trust and you're not going to have to understand the complexity of the ecosystem to feel like your data is safe. And security I know weighs into that very heavily, so...

MR. GOURLEY: Yeah, we can talk about that, too. The security piece, depending on what your business is, there is already security architectures in place to exchange data securely with consumers, especially in the finance world. You know, if you have done a mortgage recently, you have to get documents to your mortgage provider. There's secure ways to do that.

In my business, we used to do a lot of work by email. There has been a big shift in business recently towards very secure channels -- Signal, Telegram, Wicker are key providers with secure capabilities -- where you make sure that nobody can eavesdrop on your business chatter or communications with other businesses. I think that's going to grow significantly and begin to reach into consumers, too.

183

So the consumers won't have to share information by email, which is really vulnerable unless you're using a big provider like GAFA. GAFA does secure email. I trust Google, Amazon, Facebook, Apple, when it comes to email, but other than GAFA, if you are using any other provider, it's unsecure.

And so you're sharing information with a company. That company is storing it in some server somewhere, and it could get breached or lost or, you know, some employee could hit a button and send it to someone else. There's great room for new architectures and new solutions that will improve consumer privacy and protect business information through these secure solutions, like Wicker.

MS. HAN: Thank you.

MR. WOOD: So we've been on this topic for a little bit, but what other obstacles do consumers face in evaluating the benefits, costs, and risks of information sharing?

Omri?

MR. BEN-SHAHAR: Well, I think the main obstacle is complexity. Everything in the consumer world is complex, not just data policy. I mean, you can make things really simple for consumers if you really wanted choice, you know, every app, every website could have a

184

very clear button, "Do not track. Shut down all data collection and everything." That would make things very simple on the face of it, but, of course, it's not, because many other elements of the transaction depend on the fact that these businesses can harvest information and provide some -- sometimes it's for functionality, other times for profitability and the cost of the service, and now you have to -- if you shut down the data collection, if the consumer simply clicks that button, other things pop up.

Well, actually, now you have to choose. Do you want the premium or -- version, how much you want to pay, things like this? There is no way around the complexity of this aspect.

And then once you throw in that every one of those consumer transactions has other elements that -- as I mentioned, other audiences, not in this room today, but think -- other audiences think they are the most important ones, and the democracy depends on them, like class action waivers, arbitration clauses, and the like.

Then, you know, you have -- how many buttons do you have to give consumers to click and unclick and how often do they have to do it, because every website and every app has to have this. So, you see, it's becoming an infinite task of choice, and I think any reasonable,

185

emotionally sound, and rational consumer would say, "Please don't burden me with this world of autonomous choice."

MR. WOOD: Jennifer?

MS. GLASGOW: Yeah, I just -- I totally agree and I would like to add two more factors to the what do they understand. One is, it is constantly changing. A company that didn't have good security practices may have decided that they wanted to go use AWS and is now under very good security practices. So the fact that a situation -- you make a decision at a point in time to either do something or not do something, you might want to re-evaluate that on a fairly frequent basis.

The other piece I'll put on the table, which I think is going to be more and more prevalent in the coming years, is the whole idea of big data and very analytically driven business models, where it is the computer that's making the decision. And we obviously have seen a lot of that in the health area, and it's fabulous results. We see it in cities. We see it -- we see it in advertising. We see it in everything we touch.

But I think it's going to make it harder and harder for the consumer to understand how you went from point A to point B, and I will give an example that goes

186

back many years. There has been a very, very high correlation between insurance claims and your creditworthiness, and across the states, many of the states, the legislators did not understand that correlation.

Whether or not the industry did a good job of explaining it, I can't say, but they didn't understand the correlation. So many, many states have passed laws that restrict the use of credit information relative to insurance claims. And so that is a valid statistical piece of information that would help. It may not help an individual who has bad credit, but it would certainly help those that do to differentiate good practices from bad practices, but we're barred from doing it today by law. So that's what you've got to be careful that you don't -- the trap that you don't step into.

MR. WOOD: Katie?

MS. MCINNIS: So as far as consumer obstacles in this space, we've talked about the law and policies and how it's difficult for consumers to evaluate those. We've also discussed the difference between a consumer-facing organization and an organization that has no direct relationship with a consumer, where there really isn't this array of opt-in or opt-outs even available to the user. So those are two huge obstacles

187

to consumers.

Another one is that they really do have a hard time prioritizing future risk of disclosing this information when they're facing immediate problems setting up a new device or a new service that they really want the service from.

I think part of that is also due to the relative immaturity of our market. IOT devices offer a huge amount of functionality to users, and they definitely want to take advantage of that, but a lot of times they are not able to even assess the security or privacy concerns within those devices or adequately assess the privacy policies as compared to another device, which is one reason why Consumers Union/Consumer Reports, launched our digital standard last March to begin evaluating products and services under privacy and data security, but also in connection to the kind of services that these products can provide to you.

So I may be wanting a new TV and I will be able to assess the kind of color and richness and use of the TV, along with the security and privacy of these TVs, and compare that to other models, and that really does allow the consumer to effectuate some choice, where we're taking into account not only the service of the product but also the kind of data and privacy security

188

concerns.

MR. WOOD: Omri?

MR. BEN-SHAHAR: I want to say this, maybe I will pose it as a question to Katie, but my impression is that there are services in the market, when we had Professor Alessandro Acquisti here from Carnegie Mellon, so Carnegie Mellon, nonprofit, made an effort to rate the privacy performance of mobile apps, privacygrades.org, and this information is all boiled down to a single score, and yet I don't have the impression -- again, maybe I'm wrong -- that consumers are, you know, swarming to privacygrades.org to get that.

If these kind of efforts are failing, what does this tell us about the underlying question that we are all sitting here today to discuss, which is what is the consumer's injury from the existence of these practices?

MS. MCINNIS: I don't think that we're failing. I think that we're really just trying to catch up. Bob pointed out that a lot of these tech companies are moving much faster, and I think Leigh and Jennifer also mentioned this as well that tech companies are moving far and beyond faster than regulators and other kinds of checks and balances on this industry are, and that unfortunately is just how it is so far, but we're trying

189

really hard to catch up.

And just evaluating the privacy and security of apps to me doesn't create a full picture, right? Even if I know that this might have bad privacy and security, I'm only using it for some limited use, so I might not really care; however, something that has more personal information or affects my everyday life, like a calendar app or a fertility app, I would much more highly prize privacy and security in that case.

And I do think that you have to present a holistic picture of these products and security -- products and devices, which is what we're trying to do with the digital standard, but you can tell how hard it is because it's really -- not only do you have to effectively assess the privacy and security, and for our part, we can't assess what's happening on the back end; we can only look at the device itself. So we're -- even then, we're not even presented with a full picture of what the company is doing with your data.

So I don't think it's failing. I just think that we're trying to catch up, and it's extremely hard, especially with the number of devices and products that are in the market. So we hope that we will be able to change the marketplace, you know, watch the space, and we'll be talking about some of our results in the coming

190

year, but just because it's hard for consumers and hard for intermediaries to do this kind of work doesn't mean that we don't have an interest in it.

MR. WOOD: Okay, great.

So let's step back a bit and ask a maybe related question. Is there a robust market for privacy products and services? Why or why not?

MS. MCINNIS: So if I could just jump in on that, yesterday Citizen Lab out of the University of Toronto released a new security planner tool, which is a personalized experience for consumers to go through and answer questions based on not only their concerns with their online data but also the kinds of products that they use quite often.

And then it presents them with a hierarchical list of what you can do to help protect yourself online and also gives an assessment of how much time and money it's going to take for you to implement these different choices. That just released yesterday. We already have over 5000 unique hits to the site, which is showing that people really do want these tools.

It's just really hard, first of all, to effectuate the use -- the concerns and preferences across all of your devices, but also it's really hard to get a handle on your security online, and in some ways

191

it's good to just highlight a few areas, which they're doing on their security planner tool.

Also, the number of individuals that the survey results pointed out doesn't necessarily mean action, but people are very, very concerned, and they have been pushing for better protections online. We saw the backlash after the Congress reversed the Broadband Privacy Rule under the Congressional Review Act. People really care about their privacy. They just feel like they don't have a way to effectuate these concerns effectively.

MR. WOOD: Bob, did you have something to say?

MR. GOURLEY: I do, I have a couple of comments. One, I agree, this is getting really hot, because, you know, all of us are buying all these devices with chips in it, you know, the internet of things. You know, the average home three years from now may have 600 devices in it that have chips that are communicating with your WIFI, with Bluetooth, with each other, by Zigbee, and these things have vulnerabilities, and we are going to need tools like this to understand how to protect our homes.

There are commercial products available now that are aiming at the consumer in the home. The three biggest, as far as I know are, you know, Symantec has

192

something call the Core, Norton Core. Bitdefender has something you can put in your house, and then there's something provided by -- let's see, Norton, Bitdefender, and CUJO is another one. And what these things do is look at all of your internet of things devices and see what's normal.

They report back, so you have to opt in to information sharing and you have to trust that company. You pay them a couple hundred dollars a year and they have a team of people watching your devices. How many people are going to pay for that couple hundred dollars a year? That remains to be seen.

I would also say that -- Jennifer reminded me of something, and that is the analytic tools that are out there now may very well be a key threat to privacy. Right now we -- many of us have either the Android smartphones or Apple smartphones. If you get a call, sometimes it will say this call is possibly Tina, because of an email you received a year ago that had that phone number it. That's, in a way, creepy.

That's not my contacts. They're going through my email and reading it through machine learning, or you start up your car and you look at your phone and it says it will take you 44 minutes to get home from here. How does it know I was going home? Well, because that's

193

what I usually do that time of day.

Now, these are all examples of very discrete, little machine learning artificial intelligence solutions, but this stuff is growing now like it's on rocket fuel. If you look at Pinterest and what they can do now, if you pin an image, they now have machine learning tools that can look at that image and say that's -- not just that's a brown handbag, but here's exactly the type, and you may be interested in the following shoes. And did you know that when you were putting the image up there? And just watch this space of artificial intelligence and machine learning over the next three years, it really is on rocket fuel, and there are going to be privacy and security concerns that none of us have thought about.

That also applies to, like, your medical data. There's going to be analytical tools that look over all your medical data, and all of a sudden you get called in for a meeting with the doctor you weren't expecting, and maybe that's good, maybe it's not. So there's so many of these issues that we just haven't thought through yet.

MR. WOOD: Omri? I think Omri was first.

MR. BEN-SHAHAR: Well, as Bob mentioned, there are options that are sold in the markets to enhance

194

one's sense of privacy and data protection, maybe also for security. I'm not an expert on that, but I have the strong sense that if there were demand, there would be supply.

I recall that when the previous FCC was enacting its privacy rules, I looked at it, and I noticed that, you know, companies like Comcast and AT&T are offering no data collection packages, premium packages. It's just $25 more per month, and yet, if I recall correctly, not many people were purchasing these options, suggesting that, you know, one of two things can happen.

Maybe people really want it, but they don't understand, or people don't want it. I don't think that we can just proceed by saying -- and, you know, no matter what the evidence is, we'll say, oh, people just don't understand. If they did, they would want it. I think we have to consider what is this -- which of these two explanations is the right one, and that would have to be an empirical answer, do they understand or -- do they want it and don't understand or do they not want it?

And one other quick observation is, in a study that was done, my colleague looked at the privacy practices of different websites, and they found that websites that deal with more sensitive issues have

195

heightened privacy practices. For example, adult sites don't share information; cloud computing sites have heightened security measures relative -- it's all relative, but there is some -- the only explanation for this is that there is some response to what these sites perceive to be priorities of consumers. So I would call that some form of a market response.

MR. WOOD: Jennifer?

MS. GLASGOW: Just real quick, I would like to differentiate privacy products that are directed to the consumer versus privacy products that are directed to the business. I'm a little skeptical that we're going to ever see really widespread adoption of the consumer products for all the reasons that we've been talking about, but I think the business community is very hungry for privacy-enhancing technologies and products that they can build into their products, because it may be far more tedious to develop the same kind of encryption or other type of activities that you would want to make more automated. So I encourage the development of commercially oriented privacy products.

MR. WOOD: Okay, great.

So I think we are going to move on to audience questions. The first one, which is maybe to the business folks -- but feel free to answer it, anybody --

196

is how do the panelists define or quantify reputational harm?

MS. GLASGOW: Well, I'll jump in there. I think reputational harm is a harm that's evolving, and it's evolving at a pretty rapid rate. If you think about -- it -- but it varies whether we're talking about harm to an individual or harm to a business entity, because reputational harm can come from both.

We have had reputational harm in the business community from security breaches, although I'm sad to say -- and I don't have statistics to back this up, this is just my personal assessment -- that I think people are getting a little immune to security breaches, because there are so many of them, and, therefore, it is less of a differentiator from a company that has had one or hasn't.

And I think that shows the consumer kind of has given up. They don't know what to do about it. They can't fix it, and even if they're -- as Bob described, if they're a victim, then they're not really sure whether there is anything to get panicked about or not. So I think that we've got a lot of work to do there.

MR. WOOD: Anyone else? No? Okay.

So the next question is, do all consumers suffer the same informational injury? For example, Professor

197

Ben-Shahar said that no rational consumer would want to be burdened with privacy and data security choices, but if consumers' desires are actually more diverse, might some consumers actually be injured by the deprivation of that choice, even if others were not?

Do you want to --

MR. BEN-SHAHAR: Sure. I want to correct the understanding of what I said or what I intended to say. No consumer -- many -- some consumers might want to be "burdened" with these choices. I think few, if any, would want to make choices on all aspects of the consumer transaction, and there are many, many of them, and there are many such transactions.

It's just -- they're just not -- there's not enough time -- and people have studied it -- not enough time during the day to make these kind of choices affirmatively. Some things you need to let go and not make -- now, for some people, maybe the important things to choose are data privacy, and for others, it is -- it might be other aspects of the deal, maybe not so -- not necessarily that.

I don't want to say about what the ratios are, because I haven't seen anything credible about that. I mean, people say that they want, but, you know, they don't behave as if they do. Any attempt to try to

198

simplify the entirety of the data privacy -- of the privacy policies into something like nutrition labels, to make people -- to make it possible for people to choose privacy like they choose food -- and to say in parentheses, I can't resist -- there is no evidence that nutrition labels changed people's diets in any meaningful way, but there is a strong perception that they do and that this is a model to choose.

I tested that in my own work to try to create these kind of labels and to put people in a very nasty privacy setting, where they should really worry, to see whether they behaved differently when they are treated to these really friendly warning boxes as opposed to the very cluttered privacy policies, and unfortunately found, in a very large study, no effect. So it adds to my concern that people generally view these kinds of decisional aids as burdens.

MR. WOOD: Okay.

Katie?

MS. MCINNIS: I would take Omri's burden and reclassify that as agency. A lot of consumers, as we've stated, don't feel like they have any control over their data, and some of the only ways that they can even try to effectuate their choices are through these opt-ins and opt-outs. And in some cases we have even seen, as

199

we saw in the recent Courts article, sometimes those preferences aren't even followed by the company.

So consumers really desire these tools. This is -- yes, I think that the array of products that they interact with in daily life, that really does tire the consumer out to make all these decisions. I would agree with you on that, but I also think that this is one of the few ways that you can really try to have some agency over the data that you're sometimes sharing without your permission.

And I would also emphasize that while, for many consumers, the privacy of what they're doing online or sharing with companies may not be a huge concern for them, members of commonly persecuted groups or outlier groups definitely have a huge interest in the privacy of their communications and actions online. Especially when we're looking at social organizers or protestors who are looking to effectuate change in the larger status quo, you definitely do have an interest in your privacy and your security online.

I think if you didn't have that, we would see a huge chilling effect online. The NTIA did a study and found that since consumers feel so unsecured online, they have actually changed their practices online, right, and that's even before we got rid of -- before,

200

you know, net neutrality was taken away, before broadband privacy was taken away. So consumers are concerned. They just feel like there's a lack of control, and it's up to us and the regulators to provide more control for the agencies and the consumers, not less.

MR. WOOD: Okay. I think we have time for one more question and then we'll leave a few minutes for last words from the panelists.

How can industry and the FTC manage data mishaps? Audits? Who is the auditing agency? For reference, my teenager weighs risks daily, yet often makes decisions based on what she will get away with.

Anybody want to take that?

MS. MCINNIS: Unfortunately, the FTC has kind of a retroactive authority to act on these matters, except under a couple instances, such as COPPA. My personal dream is for the FTC to have more rulemaking power so that we don't have to act after the fact and after an injury has occurred in many cases.

And I think that although we have a fragmented privacy regulatory environment in the U.S., as we have seen under COPPA and other regulations, we've also mentioned some of the constraints on financial transactions, I think that consumers really desire more

201

privacy and security over the data, and I wish that we had better -- although the FTC has had a great track record and I am not diminishing that, I think that we do look for better regulations and better protections at the federal level, especially since consumers have so few tools at hand.

MR. WOOD: Okay.

Anyone else?

MS. HAN: Okay. Then let's move on to final thoughts. I will just go down the row, and you can take about a minute to give us any of your concluding remarks.

Let's start with Omri.

MR. BEN-SHAHAR: Well, I'll say one thing that, you know, the words that I've -- you know, the ideas that I shared today were largely skeptical about when -- you know, what evidence do we have about injury. I don't want to sound like -- you know, to deny the possibility that good evidence will demonstrate that there is injury.

I just want -- you know, I guess my main contribution, what I wanted to be a suggestion to the FTC, that to the extent that you identify something that needs to be done, that choice of regulatory technique is not transparency, informed consent, give consumers

202

control, improve the privacy policies, improve the format, give it in real time, real -- all of these things have been tried endlessly in so many other areas, including in privacy, and the performance of these tools is abysmal. They don't work. So if there is an injury to worry about, please, let's not worry about it through the tools of transparency.

MS. HAN: Leigh?

MS. FREUND: Yeah. Thank you again for having us. I think, you know, I would like to kind of focus us back on the purpose of the workshop here, which was trying to define what informational injury is and whether or not we can do anything about it, either at the FTC or in other places.

I think, you know, we heard wildly different opinions from experts on what was an injury on the last panel, and I -- you know, and I want to really kind of focus on, you know, are we using data in a responsible way? Are we giving consumers power over their use of data? And how are we doing that in a way that mitigates both the risks and the benefits in that not only do businesses benefit, but consumers benefit.

And I think, you know, kind of weighing those risk factors and the benefits and the risks together, you know, injury is a really big word, and it's very

203

difficult to quantify, but I think, you know, we have gotten some insights today into, you know, how serious something needs to be in order to call it injury, and the rest of it we put guardrails around to make sure that people are using data responsibly in our industry.

MS. HAN: Thank you.

Jennifer?

MS. GLASGOW: I would kind of summarize by saying, first of all I think, when having these conversations, you have got to deal with security issues in one bucket and you have got to deal with appropriate use of information in another bucket. I like to think of the appropriate use of information as introducing ethics into that, and we have lots of models in various industry sectors.

The legal sector -- many of you may be lawyers -- are familiar with ethical approaches to things, and I think that there's a play there that we can begin to adopt when it comes to various uses of information as opposed to writing hard and fast rules about what you can and can't do. I agree with Leigh. Let's identify the really serious stuff and deal with it, and then -- but there's a lot in the gray.

The last thing I'll say is I think businesses are going to have to step up to doing more -- this ties not

204

just into the concept of security, but maybe, more importantly, into the concept of how they use information and/or how they share information and what they've done to satisfy the expectation of the consumer and what they've done to give choice to the consumer, because that is the entity that will make those calls in the end.

And I think we're going to have to take some things off the table from the choice scenario. I don't want to have to give 50 choices when I buy my connected car because that's how many sensors are in the car. I want maybe three or four choices, and the rest of it, I want the car manufacturer to stand behind their decision to allow it or only use it in certain situations and so on.

So I think we have to be cautious about giving consumers choices where there are varying differences in opinion and just helping them make the right choice when they're not.

MS. HAN: Bob?

MR. GOURLEY: Okay. So first I would say I'm with her, Jennifer, I really believe all that stuff. It's very important to consider both the security aspects and the inappropriate use and appropriate use of data aspects. It's very important. She's also the one

205

who first brought up the analytics and the future of analytics and that artificial intelligence is really going to be critically important.

So many other things we didn't have time to discuss today, maybe it's going to be on the next panel, but things like tort law. Let's wait and see how these lawsuits around Equifax come out and are there other cases where companies should be sued because of negligence, and what will the courts do to shape the future of this dialogue, I think is a very important topic.

Also, we didn't have time on this panel to really discuss best practices for businesses. There's so many of them out there that are captured and that need to be contextualized for businesses, but there are important best practices out there right now that more people need to know about.

And the same for home, us people at home -- Jennifer mentioned cars, of course, but cars, homes -- all of your devices. What are best practices for those, we didn't have time to talk about. We did have time to talk about the big guys who are really managing our data and working hard to protect it, GAFA as I said, what would GAFA do, we need to keep looking at that. Those are my concluding thoughts.

206

MS. HAN: Thank you.

Katie, you have the last word.

MS. MCINNIS: Thank you for having us on this panel. I found it very interesting and it was also great to meet many of you in person.

I just wanted to take a page out of Bob's book and mention that there's one thing that we didn't talk about, which is another injury to consumers, which is a loss of consumer power. If tons of information is being collected about me, with or without my knowledge, and that leads to first-party price discrimination, that is an injury to consumers, and that's one that I think we're overlooking and haven't paid enough attention to at this time.

I also want to emphasize that Consumers Union and Consumer Reports, which is the same thing, is really looking to try and provide consumers with more agency and more ability in the marketplace really to decide what kind of products and services you want to use based not only on the services that these products provide but also based on the security and privacy of those services.

And we hope that we can help change the marketplace so it is easier for consumers and so that we don't have to get tired out by all these different

207

disclosures, even though I personally see that more as an agency issue than a burden.

And I wanted to point you all to the security planner from the University of Toronto's Citizen Lab which can help each and every one of you effectuate some privacy and data security protections while online. It can be tailored to you.

Thank you, Cora and Dan, for organizing this panel.

MS. HAN: Great. Thank you. So please join me in thanking our panelists.

(Applause.)

MS. HAN: It was a great discussion. We will be on break until 3:30, when our last panel on measuring injury will begin.

(A brief recess was taken.)

208

PANEL 4: MEASURING INJURY

MR. SMITH: So, hi. Good afternoon. My name is Doug Smith, and my co-moderator on this panel is Jacqueline Connor. So welcome to the last panel of the day, measuring injury.

As I think, you know, several times has been alluded to on the previous panels, there are particular empirical issues, or challenges that come up in trying to assess informational injury, and fortunately, today, we have a wonderful group of panelists to help us understand them.

MS. CONNOR: Yes. As Doug said, we have five wonderful panelists with us this afternoon. So I am going to give you a brief introduction to each of them.

First we have Garrett Glasgow, who is a senior consultant at NERA Economic Consulting, where he specializes in market competition, consumer protection, intellectual property, and environmental cases. He has presented and published on several topics related to damages stemming from violations of privacy or misuse of personally identifying information.

Next to him we have Ginger Jin, who is a Professor of Economics at the University of Maryland, who was on leave at the FTC in 2015 to 2017, serving as Director of the FTC's Bureau of Economics from January

209

2016 to July 2017. Most of her research has focused on information asymmetry among economic agents and ways to overcome the information problem.

Next we have Lynn Langton. She is the Chief of the Victimization Statistics Unit at the Bureau of Justice Statistics at the Department of Justice. She is responsible for overseeing the collection and the dissemination of data from the National Crime Victimization Survey, a large-scale survey of U.S. residents that serves as one of the two measures of crime in this country.

Catherine Tucker is the Sloan Distinguished Professor of Management and Professor of Marketing at MIT's Sloan School of Management. Her research interests lie in how technology allows firms to use digital data to improve their operations and marketing, and then the challenges this poses for regulations designed to promote innovation.

And last, but not least, we have Josephine Wolff, who is an Assistant Professor in the Public Policy and Computing Security Departments at the Rochester Institute of Technology. Her research interests include cyber security law and policy, the economics of information security, security metrics, incident reporting models, and insurance and liability protection

210

for computer security incidents.

So first we would like to have each panelist talk a little bit about the work that they have done relating to measuring informational injury. So I'm going to read a brief question, and then if each of you could go down the line and kind of give a quick response to that, that would be great.

So starting with Garrett and I guess working down, can you please give us a brief description of the research or work that you've done related to trying to measure injury in the context of privacy and data security issues?

MR. GLASGOW: Yes. Thank you for having me here. In terms of privacy and misuse of personal data, I'm working in two areas. One -- which I call the easy area -- is misuse of information by companies, so where a customer and a company might have some kind of business relationship, certain levels of privacy or data protection are promised, and then the company for whatever reason doesn't live up to their promises.

I've been working on, survey methods to try to value how much of the price the consumer pays for a product, bakes in this data protection or privacy. That's one area where I feel like we've made some good progress at NERA.

211

The other I call my blue sky research, and that's trying to determine whether there's an intrinsic value of privacy, so not the theft of data or the misuse of data, but do consumers, do individuals actually obtain some kind of utility or value from just knowing that their data is being kept private and is not being pried into by unauthorized people?

This is an idea that comes from environmental economics, which is another area where I work, where we talk about the intrinsic value of the environment. It's not something that has value because it's the hub of a lot of economic activity, but it has what we call passive use value, where just we gain some kind of happiness or utility from just knowing that there's a pristine environment that hasn't been damaged, and if it is damaged, somebody's owed some kind of compensation.

And so I've been doing research in that area as well to try to determine if privacy is similar to an ecosystem that might be damaged by bad acts or by carelessness.

MS. CONNOR: Thank you.

Ginger?

MS. JIN: Thank you. Thank you so much for having me. I have to confess that I have done little research in this area before coming to FTC in 2016;

212

however, while I was at FTC, I dealt with a number of cases in privacy and data security which prompted me to think a lot about information injury.

One crucial question is whether we should measure consumer harm in ex ante or ex post perspective. I know the second panel has already covered on this, but I still want to put in my two cents. The ex post perspective is quite intuitive. Somebody misuses consumer's information that results in say identity theft or a fraudulent transaction, and we can measure that by the amount of time, money, and effort that consumers lost because of this.

However, the ex post perspective I would argue is narrow-minded because a lot of harm may not happen yet, but there's a risk there, or the harm has happened, but we cannot link to a particular firm who have done that data practice. So in these situations, we observe zero traceable harm, but that does not mean there is no consumer harm.

I would also argue that, in fact, emphasis on ex post harm would end up encouraging overuse or misuse of data, because firms who engage in those misuse of data, they don't need to account for the negative externality they are imposing on the consumers. So in my view, that's inadequate.

213

In contrast, the ex ante perspective would emphasize the increase of risk to harm consumers even if that risk has not been realized or could not be traced back to a particular firm.

So I heard a lot of panelists talking about Equifax. Myself is a victim of Equifax. I worry about my record on the black market. I end up paying something to try to freeze my account. I even sort of welcome the inconvenience, I have to lock and unlock on my account in order to get some credit in between.

And to me, that's -- it spends my time and effort and money, and that is harm to me, even though that's -- we haven't seen other harm coming back in my way, and it's probably very hard to trace that back.

Similarly, if I know my favorite retailer has some bad data practice, even if that company has not had data breach, I will be worried. I would want to probably do something to reduce that risk. So in that sense, the ex ante perspective would say there is some harm there.

And to follow on Garrett's suit, we can make an environmental analogy here. So, for example, a firm has polluted my neighborhood and exposed me to a higher risk of cancer. I would say the company should be responsible for that action even though I have not

214

developed cancer in my body yet, because it's increased the risk I'm exposed to. So in that sense, I believe the ex ante perspective is more appropriate.

Another point I want to make is that the biggest difficulty in measuring harm is not measuring harm itself. It's measuring harm attributable to a specific firm. We can measure the extent of identity theft or other things, but exactly how to tie that back to a particular firm, I think that's the most difficult question.

MS. CONNOR: Thank you, Ginger.

Lynn?

MS. LANGTON: Thank you for having me here today. As Jacqueline mentioned, I oversee the National Crime Victimization Survey, which is one of two measures of crime in the U.S. It dates back to the 1970s when it was developed to be a complement to police statistics because of the recognition that a large portion of crime goes unreported to police, and so if we just focus on those crimes that are reported to law enforcement, we're missing a big piece of the picture.

And this is particularly true when you're talking about sensitive crimes and when you're talking about more of what we would call sort of white collar crimes or emerging crimes -- though I don't think "emerging" is

215

really the correct term anymore -- so crimes like identity theft and fraud that are highly unreported to police. So the NCVS is a household-based survey. It's conducted for BJS by the U.S. Census Bureau.

We go to an incredibly large sample of households and interview all persons age 12 and older within those households about their experiences with victimizations. So our sample size is -- right now we go to over 200,000 people every year, and we get response rates in the high seventies. So if you do survey research at all, you know that that's pretty remarkable at this point in time, still -- and we ask questions about their experiences with victimization and also the nature of that victimization and the victim response to that victimization. So as Ginger was just speaking about, this is the ex post perspective very much.

The National Crime Victimization Survey focuses specifically on violent and property crime, but we do conduct a number of supplements to the survey, and the one that I think is most relevant for the conversation today is a supplement that we have been conducting for a number of years on identity theft.

Of course, identity theft is just one type of informational injury. Again, ex post, very specific harms associated perhaps with a data breach and often

216

just because of misuse of personal information through other means, but I think there are some important things that we can glean from the identity theft supplement that are relevant to the broader conversation here today.

So the supplement collects information about the attributes of victims and victimizations and the response to victimization, but it also asks questions about the harms experienced by victims who suffered from identity theft, and on top of that, it asks victims and nonvictims about behaviors that they engaged in to prevent any type of misuse of their identifying information and also whether they have experienced any kind of data breach in the past 12 months.

So when we look at that information, we can see that among individuals who have experienced a data breach, the risk of identity theft or the prevalence of identity theft is double that of those that have not experienced a breach -- not that surprising, but putting numbers behind what we know to be true or assume to be true. And then among identity theft victims, about 20 percent of the victims in our survey say they also experienced a data breach during that reference period. So there's certainly a correlation there, not surprising.

217

I think the other point that's important to take away, though, is that even among identity theft victims that experience no financial losses -- so they have no financial losses whatsoever -- there are still harms. So we asked them questions about not just their losses and the time spent, but also about how distressing they thought that the incident was, and even among those identity theft victims that didn't experience financial loss, about 30 percent still found the incident to be moderately to severely distressing. So that suggests that, in and of itself, the experience of having your information misused, having your information out there, has some harm to victims, and so I think that's relevant for our conversation today.

MS. TUCKER: All right. So I'm Catherine Tucker, and I am an empirical economist, which means that I'm going to behave in a predictably economist way and say two things.

The first thing is a lot of my research has been focused on how hard it is to actually measure this, and, in particular, my research, which is focused on the use of algorithms and AI, suggests this is just going to get harder and harder as the boundaries of an individual's data gets fuzzier and fuzzier.

The second thing which I'm going to say, which is

218

predictably economist, is that my research is really focused a lot on documenting the trade-offs which come from trying to protect consumers and their privacy, and let me tell you a little -- give you sort of three typical unintended consequences that I've found in that research.

So the first unintended consequence I want to highlight is often when we have privacy or data breach regulations, we're very focused on the idea of big companies and the idea that they need to be provoked to do something better, but my research suggests that it's actually smaller companies, startups, which tend to be most aversely affected in terms of the costs of this kind of regulation.

The next piece of research of unintended consequences I want to highlight is some research I did into data breach notification laws and, in particular, what happened when they gave exceptions to encryption. Now, in this research, what we found is that actually there were more data breaches as a result of these exemptions. And why was that? Well, they just focused on encryption, and it meant you had a whole lot more doctors losing laptops from their car as a result of people focusing on just one dimension of data security.

Now, the last thing I just want to highlight as

219

an unintended consequence is some research we've recently done about patient data and hospitals which highlights that actually hospitals are using data security and privacy regulations to actually stop patients themselves getting access to their medical data records they need if they try and transfer hospitals. And you -- some people in the room, I see you nodding, right? You have experienced this in trying to get your own medical data out of the system.

MS. WOLFF: Thank you.

So I've worked on a couple of areas related to thinking about the injury and the types of harm that come out of data breaches and cyber security incidents more generally, and the first time I sort of encountered this was in a project looking at the aftermath of security incidents and trying to trace who ends up suing whom and who ends up paying whom, and a lot of this is at the corporate level.

A lot of this is the credit card companies and payment networks suing breached retailers and demanding certain amounts to settle the fraud costs, and then there are also a number of the consumer class action suits that you heard talked about a little bit on the panel before lunch, and for all the reasons that they brought up on that panel, those are often very hard, to

220

see sort of go through or even get heard by courts, and I got interested in this question of sort of how sort of different types of policy agencies, different parts of the government, and different parts of industry think about and understand these types of harms or injuries.

And so I did another project looking across sort of the different ways that people had tried to measure them and bucketing a couple different kinds of data. There's a lot of survey data in this space, a lot of, you know, sending out reports to different companies and saying, how much did you lose due to breaches in the last year, and putting that together as sort of self-reported data to think about how much we think these breaches cost to the companies, to their customers sometimes.

One of the big sources that some of you may have seen is the Ponemon study that comes out every year, again, focused on self-reported numbers and trying to stack those up against some of the other data sources that we sometimes use to think about quantifying these types of harm. One of those was insurance claims. NetDiligence has a data set of companies that have filed cyber insurance claims for harms due to computer-related losses and looking at sort of how the numbers change if you look at what they file for insurance versus what

221

they report in a survey about these losses.

And then, finally, looking at a lot of the analysis that's been done around stock prices and market value of breach companies and trying to understand sort of how that tells a slightly different story if you use that as a proxy for thinking about the costs or the consequences of breaches.

And then to come back to sort of thinking about the consumer injury and not just the losses to a company, but also to individuals. I spend a lot of time looking at the firms that are offering cyber insurance policies now and how they price those and how they think about the kinds of harm or losses that they will or won't cover, and increasingly, we're seeing a lot of them get into the still relatively small market of personal cyber insurance, selling policies to you individually or to your family to cover certain kinds of losses and trying to understand how it is that we understand the ways people are purchasing those types of policies, what kinds of things we're covering with them, as a way of thinking about the types of losses that people are concerned about.

And also, coming back to sort of some of the skepticism you heard on the previous panel from Omri about, you know, what do people actually do, not what do

222

they say they care about, but what are they actually willing to pay for? What can we see, when we look at this emerging personal cyber insurance market, about the types of injury customers seem most concerned about and are most willing to actually pay for protection?

MR. SMITH: All right. Thank you, guys.

So now we will sort of jump into the discussion. We are going to start off, I think, by talking about kind of trying to measure preferences, right, trying to understand what consumers value. And so I think a good place to start that would be sort of how to measure -- just what consumers say about what they value. And, Garrett, you talked about doing surveys on this, so maybe if you could tell us a little bit about how that works.

MR. GLASGOW: Sure. So, yeah, my work in measuring privacy, measuring the value of data, has primarily been survey-based. Before coming to NERA about 3 1/2 years ago, I was a professor in the Political Science Department at UC Santa Barbara, where I did lots of survey research, so maybe it was natural that I fell into a survey-based approach to looking at privacy and the value of data, and there are two different approaches that I've looked at.

There's what's known as conjoint analysis, and

223

then there's contingent valuation, two different survey type approaches, and I'll just briefly talk about each and then the strengths and weaknesses of each of these approaches.

Conjoint analysis is well known in consumer research. It's accepted in court as a reliable method for uncovering truth in a lot of contexts. And what conjoint analysis involves is we ask our survey respondents to make choices, as if they were in the market making choices, from among some set of hypothetical products, and these products are going to have different features or attributes, and so we can see, if they make enough choices, which attributes they seem to value and which ones are unimportant.

I'll give you an example from my own research. We did research on streaming video services, and we present people with, say, pretend you're in the market to update your streaming video service. Here are some services that might be available to you. They might have different streaming speeds; some might have high-definition available; some might have more television shows; of course, different monthly prices and so on.

We just present these different services and say, well, pretend these are the three that are available in

224

your area or that you could possibly purchase. Which one of these services would you purchase? And they can make those choices.

Now, when we want to bring this to privacy, one way we can do this is just regard privacy as just another feature of a product. Obviously, this is only going to work if we're looking at a situation where it's a consumer engaging in a market transaction with some company that's made some kind of promise about privacy or how they are going to treat your data.

So we can say, what we did with this paper that I'm talking about, is part of that streaming video service was, and then there's different possible privacy policies that these services offer. They might say that we never share your data. Others might say, well, we'll collect your viewing habits, not your personal information, but we will collect your viewing habits and package those up and use those to help content providers decide what kinds of shows to make. And then the third option is maybe we collect your personal information as well, and we might share that with third-party marketers, and so on.

And what we were able to do with this survey was see what value do people place on protecting their data, or conversely, what kind of discount do we need to offer

225

consumers to get them to agree to share different types of data?

The strengths are this is a really straightforward, reliable way to measure consumer preferences. It's widely used. It's widely understood. I regard its main weakness as narrowness. As I said, these are consumers engaged in a market transaction. What do we do in all the privacy cases where there isn't a market transaction, like a data breach?

I mean, I suppose we could design some kind of conjoint analysis where we say, now pretend there's a certain percent chance that this company is going to have a data breach in the next year, but I don't think most consumers think about their purchases that way. It's a strange hypothetical to pose to people, and I think you would get really strange results with a survey like that. Maybe it's possible to do, but I think that's one of the main weaknesses.

But at least in areas where we can apply this, this conjoint analysis lets us measure the value that individuals are placing on their privacy or on different data-sharing policies, gives us what we call a willingness to pay to protect their data, and then we're talking about quantifying damages. If a company doesn't live up to that promise, we can use that willingness to

226

pay to calculate damages, that we can use that as a basis to say a certain percentage of this price was privacy protection, you didn't live up to privacy protection, you owe consumers some refund or some amount of damages based on your failure to live up to your promise.

I mentioned earlier, there's easy cases and hard cases. I think those are kind of the easy cases. I think we have made some headway in terms of measuring the value of information, the value of privacy there.

Hard cases are all the other ones, things like data breaches and so on, and one possible approach to that is what's known as contingent valuation, and to really -- you know, a really rough description of that would be, well, asking people how much compensation they would need if some bad event happened. This is very common in the environmental setting, in areas where there isn't a market transaction. There isn't a market for the environment, but there could still be damage done, and there might still be compensation owed.

The classic example of contingent valuation came from the Exxon Valdez spill. So the Exxon Valdez was a big tanker that spilled a whole bunch of oil in Prince William Sound up in Alaska. It's a place most people have never seen. There is not really much economic

227

activity going on there except for the tankers going through, but people still felt aggrieved, they were still damaged.

How do we measure what kind of compensation Exxon might owe to the American people? And one way that this could be done would be contingent valuation. They went out and conducted surveys. They said, well, suppose we had to increase taxes next to year to create a system, maybe a tugboat system, that would make sure these kinds of spills didn't happen again? How much would you pay to prevent this from happening?

And people would give us answers, and you could tabulate that up and say that this is the value of that environment to these people. You can easily see taking this over to a data breach environment or a privacy environment, where we say a data breach has happened. How much would you pay to do business with a company that has some lock-down system where they have reduced the chance of a data breach to almost nothing? We can see designing a survey that would be something along these lines.

The strengths here are that this is -- unlike some other surveys, this is something we can at least apply to something that's not a traditional market transaction. A data breach is not a traditional market

228

transaction. The weaknesses are that these are kind of strange hypotheticals that we're presenting to people. It's -- and people -- especially with issues that are a bit emotional, things like environmental damage or privacy or data breaches, people often give what are known as socially desirable or protest answers.

How much would you pay to clean up Prince William Sound? We will have people say, I'd pay $50,000 to clean up Prince William Sound. Well, their income was $45,000. They are just trying to signal to us that they're angry about what's happened. And I would imagine with doing surveys on data breaches, you would get the same kind of anger and the same kind of protest answers, where people want to signal that they are angry at companies for allowing a data breach to happen or for leaking their data, and it might not be very useful in terms of measuring damages, but that is one possible survey approach to dealing with data breaches. I'll talk more about that later on, but those are the -- that's an overview of two major survey approaches.

MR. SMITH: Great. I want to open this question up to all of the panelists, but first, I know, Catherine Tucker, you have done some work in this area, so I was wondering if you could sort of comment on what you have done and found.

229

MS. TUCKER: Oh, yes. So just as I previewed, it is going to be about the difficulties in this, and the work I did was basically investigating something we call the privacy paradox, which is that if you ask people in a survey how much they care about privacy, everyone cares about it, but then we see all this actual behavior which is inconsistent with that stated preference.

Now, what we did in this study is we -- it took place among the MIT undergraduate population, and we asked them to take part in a study, and as a reward, they got $100 in Bitcoin, and this was in 2014, so those of you who can do the math, this is a lot of money right now. We made a good call, that.

So when we did this study, one of the questions we asked them for was some data on some really quite personal and sensitive information, which was data about their friends' contact details. Now, in the ordinary setting, when we asked this question, a surprising number of our MIT undergraduates really decided they were not going to answer this question, but instead, they were going to give us fake information, and we knew it was fake because of their use of swear words and their use -- and the use of swear words employed to describe what they thought about us for actually asking this question as part of the email. So it was

230

definitely deliberate. Now, that was the first setting.

And then for half of them, we changed the context somewhat in that we gave them an offer of cheese pizza, and if you're wondering how much cheese pizza, sort of a slice of cheese pizza. And what we found was that when there was no cheese pizza involved, people tend to behave in a way which was consistent with what they said in a survey about their privacy preferences, but the moment we gave them the offer of cheese pizza, even those people who said they really care deeply about privacy started giving away this quite sensitive data.

Now, you know, in terms of interpreting it, it's really quite hard, right? One thing you could do is sort of a typical economist response, which is that stated preferences, it's hard to use them to measure anything, and we should use revealed preferences. Another interpretation of the study is, though, oh, my gosh, if MIT undergraduates behave like this, people who should really understand technology, maybe we need to really protect people to actually get them to behave in line with their stated preferences. So no conclusions, just some difficulties in terms of using survey responses.

MR. SMITH: Thank you, Catherine.

Ginger, do you have any thoughts about --

231

MS. JIN: Yeah. I just want to follow up with Catherine's comments on the privacy paradox and how actions seem to differ a lot from stated preferences. I think the biggest question is to distinguish different explanations behind this. Is this because consumers don't know the risks that they are willing to give away, or because they know the risks but they somehow believe the benefits dominates the risk, or they see the extra risk of giving away this to one more person, given that their data probably has already been breached multiple times, is so small that they feel helpless, and so, therefore, they give away?

I think those explanations would sort of imply very different policy actions. It would be -- I will be really interested to see if future research will be able to distinguish those explanations.

MR. SMITH: Thanks, Ginger.

Lynn?

MS. LANGTON: So as I'm sitting here thinking about measuring actions, I'm realizing that there is more that we could potentially be doing with our survey. So we're not getting at consumer preferences in any way, but we are asking a series of questions about behaviors that they engaged to -- again, this is in the context of identity theft -- but behaviors they engaged to to

232

prevent the misuse of their personal information, and I think, again, this speaks to what they're actually doing.

So, you know, these are pretty basic things, but I think the proportion of people that do these things regularly is still pretty surprising. So if you ask them if they check their bank and credit statements, okay, 76 percent say they do that, which makes me wonder about the other quarter of people, I have to say. Who doesn't ever check their bank statement? But that's a different subject.

But then when you -- you know, when you move into things like do people regularly change passwords on financial accounts, about 30 percent say they do; checked their credit reports, that's about 40 percent; purchased identity theft insurance or a credit monitoring service, again, maybe this gets more into that willingness to actually pay, to spend money on something that might prevent this from happening, and that's only 5 percent of the residents that are participating in the survey that say they actively do this.

So now I'm thinking, what could we ask? Could we ask some questions related to data breaches specifically that would get at their -- the behaviors that they

233

engage in? So even just how often do you consider whether a company has had any sort of breach that you can find information about before you provide your information to them, just to kind of gauge whether or not there are actions being taken to actively avoid potential situations.

And, again, that -- companies can change over time, and so that would -- you know, those are challenging things to measure, but I think there are still some actions that we could think about concretely getting at that would address this issue a little bit more.

MR. SMITH: Thanks, Lynn. And I think, you know, talking about that, you sort of broadened the conversation to looking directly at people's actions or at least their purported actions.

So, Josephine, would you like to comment on either sort of stated preferences or revealed preferences or both?

MS. WOLFF: Sure. Well I think, you know, a lot of what we have when we look at the data people have analyzed and collected around the injuries and costs associated with data breaches is very much stated preferences. Most of what you see, especially, say, in industry reports, is going to be kind of self-reporting,

234

this is how much this breach cost, or this is how I was affected, these were the injuries.

And so I think that we start out from a position of basing a lot of our estimates, basing a lot of our understanding of these types of costs on those stated preferences just for practical reasons, that that was sort of the -- and still is in a lot of ways -- the data that's easiest to collect in large volume, and it's definitely true, and several of the other panelists have brought it up, there's a huge discrepancy when you start trying to compare that to what people are actually doing.

It's very hard, as Ginger says, to get at why exactly that is, and I think one of the things that's complicated when I, say, look at legal cases around these class action suits and the aftermath of data breaches, is that getting involved in one of those class action suits is not always just a sort of straightforward question, do I care about my data? Do I value my privacy to some amount of time?

And the strongest example of this I think is probably the Ashley Madison class action lawsuit, where the judge actually decided if you want to go forward as a member of this class action lawsuit, you are going to have to use your real name, right? So it was this sort

235

of complicated decision for everybody whose data had been stored by Ashley Madison and breached, that, okay, on the one hand, I valued that data and the loss of privacy that came, with everybody in the whole world potentially knowing that I'm a member of this website for adulterous affairs, but on the other hand, if I now move forward with trying to pursue that as an injury in court, I am going to lose this privacy in a slightly different way.

That kind of decision, I think, is not actually as straightforward as saying, oh, I guess nobody cares about the injury that was done to them in this data breach because they are not willing to pursue this lawsuit, but more about trying to understand what are the actual decisions that people are making when it comes down to acting on what they think are their preferences, on what they think is what they value.

MR. SMITH: So in some sense understanding sort of the costs of the trade-off, the benefits of privacy, it's sort of important to figure out how to measure it?

MS. WOLFF: Absolutely, and I think it gets at some of that discrepancy we see, which would be sort of how much I think I value my privacy and the things I'm actually maybe willing to do in practice to protect it or pursue my losses.

236

MR. SMITH: Okay, thanks, Josephine

So do people have any further thoughts about sort of measuring consumer preferences through their actions?

MR. GLASGOW: Yeah. Actually, I think we have had heard several different stories now about a mismatch between a consumer's stated preferences and consumer actions, and this is a regular feature of privacy research. We see this a lot, that people seem to take actions that contradict what they have stated matters to them. And I think that it's often used to dismiss the concerns of consumers in what I think is an unfair way.

It's a complicated issue, whether you -- people might be willing to take some actions in some contexts but not in others. I'll give you one example from my research, that conjoint analysis I mentioned on streaming video services. About 10 percent of our survey respondents said I would never choose a streaming video service that collects any information about me, whether that be movies about -- which movies I watch or personal information or anything like that, but one of our screening questions to get into our survey was, do you have currently have a streaming video service? People had to say yes to get into the survey.

So 10 percent of people already had a streaming video service that was collecting information about

237

them. I don't know how many of you saw the Netflix tweet from yesterday or the day before, but Netflix sent out a tweet that said for the 57 of you that have watched The Christmas Prince 18 times in a row, are you guys okay? It was something like that. So they clearly are collecting information on who's watching what, and they even sent out a tweet kind of teasing people about movie choices.

So what's happening with this 10 percent of people in the survey that said I would never choose a streaming video service that shares my information, yet they already are doing business with one? One possibility is just lack of information. They weren't actually aware that the streaming video services are collecting this information, and that's one possibility for this mismatch between action and preference, although if somebody is really a privacy hawk and says I would never share information, they probably know who's sharing what.

Another possibility is -- this is that signaling of importance that I mentioned earlier, where you give a protest answer. So maybe those 10 percent of people, even though we tried to be careful, they figured out this isn't a survey about streaming video services. This is a survey about privacy, so I am going to send a

238

signal that I value privacy by saying I never share anything, and that will let those researchers know that this is really important to me. That's another possibility for the mismatch. It's not that the consumers are being inconsistent, just that we failed as researchers to try to hide what we were trying to get at and honestly measure that.

Then the third possibility is the context. In some situations, you're willing to share, and in some situations you're not, and consumers would like to pick and choose, or in some cases you might perceive a threat and in other cases you might not.

And I'll end with another little story. This is from about three years ago. An artist baked some really fancy cookies that had the Facebook logo and the Google logo and so on and went out into New York City and gave them away to people -- well, didn't give them away -- didn't sell them, but they said you have to give me personal information if you want this cookie.

And about half of the people were willing to let their photograph be taken in exchange for a cookie. Half of the people were willing to give what they said were the last four digits of their Social Security number. They -- a lot of these people were willing to let her write down information off their driver's

239

license in order to get a cookie. And one-third of people were willing to give up their fingerprints for a cookie.

So you say, well, how do we reconcile that with people who say they care about privacy? Well, maybe there's a context issue here, that an art student probably isn't going to take those fingerprints and then go rob a bank using them or something like that. There is probably a very low risk there, but if a data hacker gets the same kind of information and it winds up on the black market, that's a very different risk.

So I don't have answers to this necessarily, but I would say there's often a mismatch between preferences and actions when it comes to privacy. It doesn't necessarily mean that consumers don't know what they're doing. It's a complicated issue.

MR. SMITH: Thanks.

Ginger, do you have any further thoughts?

MS. JIN: Yeah. I just want to emphasize that, as an economist, I often believe actions are louder than words, but in this particular area, I really believe the survey approach and the revealed preference approach should be really complements to each other. We actually have already seen some interactions here while Catherine has been talking about sort of MIT student action in the

240

lab, while Lynn was trying to think about how you can design a better survey question to get at that sort of behavior in a national representative sample.

I guess the feedback could go the other way, that if we know, for example, consumer lack of information or feeling helpless is one of the primary reasons for them to behave a certain way, maybe Catherine can design a very clever lab experiment to determine or to really tease out people's actions. I really see the positive complementarity between the two approaches.

MR. SMITH: So we're kind of going long on this topic, but, Lynn, do you have any sort of last thoughts on this?

MS. LANGTON: I'm sorry?

MR. SMITH: Do you have any last thoughts on this before we move on to outcomes?

MS. TUCKER: Yeah. I would just like to also add, to make it worse, that actions don't always translate to actions in the way we would like, and if you sort of move along in the Bitcoin study I was talking about, we then tested really how much laziness affects people's privacy preferences, in that we asked them what wallet they wanted them to put their Bitcoin in. We gave them a variety of wallets, some of which protected their privacy, some of which completely

241

didn't, and the biggest predictor of what wallet they chose was not their privacy practices. It was, instead, how far up on the page that wallet was.

Then, we then said, okay, well, maybe it's just like a lack of information, so half of the students saw lots of information about the wallets, basically everything you could ever want to know, and there was some good news in that that reduced, shall we say, half of the laziness, but I guess the bad news is there was still half of the laziness there even for people expressing a lot of privacy preferences.

So in other ways, you can actually change people's actions a lot simply by the context in which you get them to make the choices, making even this action data hard to judge, I'm afraid.

MR. SMITH: Josephine, any last thoughts on --

MS. WOLFF: Well, I think -- I mean, I think what Ginger says about needing both the survey and the revealed preference data is really important because I think there are certain types of injuries, and you heard a lot about that this morning, all the different ways that people can be affected by data breaches.

I think there are some of those injuries that we really have no other way to get at, especially when we're talking after a breach has occurred, when we are

242

not necessarily going to say that all of the different types of emotional or psychological harm you may have suffered are things you can obviously mitigate, right?

It's one thing to say, well, people didn't go and change their passwords or people didn't go and file a class action lawsuit, but if all of your emails have been released in the Sony breach, then it's not obvious what the action you would take would be to demonstrate that that really affected you, that that was a real kind of injury.

So I certainly think it's true that the self-reported data has a role to play in this. I think where it becomes tricky is when you're looking at really kind of crisp quantification for things like legal remedies or policy interventions, whether there's any way to turn that into something that can be calculated precisely enough for those types of remedies.

MR. SMITH: Okay, thanks, Josphine.

So we are going to have to leave the topic of preferences and move on to, you know, an alternative, which is to just straight up measure outcomes and try to understand things that way.

So, Lynn, I'd like to hear more about sort of how your study allows us to get at that question.

MS. LANGTON: So, again, we're talking in the

243

context of identity theft here, so beyond just any sort of data breach or privacy-related issue, but we ask a whole series of questions related to both tangible and intangible harms -- again, ex post harms -- associated with the misuse of personal information.

So, of course, we ask about financial losses; that's an obvious one that you have to include. We also ask about the amount of time that an individual had to spend clearing up the issues related to the victimization. And then we ask a whole series of other questions trying to get more at some of these intangibles.

So I mentioned already that we ask about distress. We also asked about whether the incident resulted in any problems with family, friends, work, school. And then we asked questions about whether they experienced any credit problems related to the incident; whether they experienced any legal problems related to the incident; whether they had to deal with debt collectors. So these are some of those more intangibles.

And then, you know, I think to get more at those intangibles, too, and, like, really the impact of an incident on an individual -- and, again, this goes broader than identity theft -- you know, the other thing

244

that we can do is then sort of cross these different types of outcomes and different types of harm.

So when we look at, for example, how much time an individual has to spend dealing with an issue, and then we look at the level of distress, I mean, there's a positive linear relationship there that's pretty strong. So when an individual has to spend six months or more dealing with this misuse of their information, you know, a large portion of them say they were severely distressed, whereas if they're able to resolve it in a day or so, you know, a smaller portion say that it was severely distressing.

So, again, it's indirectly translatable to data breaches, but I think sort of the same ideas, and you can use survey research to sort of tap into these ex post responses and harms that victims experience as a result of these incidents.

MR. SMITH: Thanks, Lynn.

Catherine, do you have any thoughts about this, about sort of measuring outcomes, what data we might look for or how we might take things that are sort of less tangible and convert them into something we can sort of quantify?

MS. TUCKER: Well, so, my thought actually was that really when it comes to measuring injury, I think

245

we lack a good model of the supply side of people who want to do injury. And what do I mean by that?

Well, I've seen some really intriguing research at the Workshop on the Economics of Information Security, which is all about your potential for actually being injured as a result of a bad actor on the dark web. And you should go and look at this research, because it's fascinating, in that these great computer scientists go out there, and they actually try and work out, if I'm trying to do identity theft, how do I do it? And so it's all very interesting.

But what they found out, which really surprised me, is it's just not a very scalable profession, and as a result, basically they take so long to do their identity theft, there tends to be 100 people each time who are injured, and it doesn't vary as much with the numbers of records that are left or lost as you might think.

And so I think, you know, if I was sort of -- as Lynn is collecting such wonderful data, I think if we were to really try and have a measure of identity theft, the risk of it, sort of the first point of measuring the injury, I would actually like to sort of highlight this research.

Which talks about the supply side and some of the

246

surprising insights that it's not very scalable, which may explain why sometimes the incidence of identity theft from a data breach is less than we might think.

MR. SMITH: Thanks, Catherine.

Josephine, any thoughts on these measurement questions?

MS. WOLFF: So one of the things I think is really interesting about outcomes -- and I come across this a lot when I'm both talking to people who sell insurance and people who buy insurance in this space -- is that outcomes are much easier to tie to an individual person than to an individual incident.

And so if we're talking about, you know, what were the consequences of this particular breach with Equifax or whoever else, that's a very, very hard question to answer, because your information has been stolen so many times. Earlier today, they referenced they think most of the Equifax information had probably been available on the black market even before that breach, so it's hard to say this incident of identity theft or this particular -- even sort of emotional or psychological toll associated with that was the fault of this particular company and this particular breach.

It's much easier -- though it's not easy by any stretch -- to say this is the outcome associated with

247

this individual person. This is the amount of financial fraud. This is the amount of time. All of those things are easier to measure in terms of people. And the reason that's important, I think, is because all of our thoughts and all of our sort of ways of understanding how policy can come in and try and correct these externalities have to do with what are the levers, what are the pressures we apply to the companies and say your breach caused this amount of damage, which is actually a much harder question to answer than what were the outcomes for this individual person as a result of the sum of all of the incidents that they've been involved with?

And I think that's a big part of the reason that we're starting to see a lot of the solutions -- or at least the insurance market certainly thinks of them as solutions -- coming to center more on the individual than on the companies. But I think one of the things that is perhaps problematic about that is even though that aligns with the ways that we can perhaps calculate or collect data more accurately, it means that all of the entities that have the most power to decide how this data is being protected are not necessarily the ones who have the real incentives to be executing that in the same way.

248

MR. SMITH: Thanks, Josephine.

Garrett, do you have any thoughts on this?

MR. GLASGOW: Sure, just briefly. I keep using the words "easy cases" and "hard cases" if we're talking about outcomes. I think in some cases it's -- I wouldn't say easy, but it's easier to measure harm. I've talked about market transactions with companies that don't live up to a promise. So, say, for example, you do business with a company that has a certain privacy policy. They say we will not share your data with marketers, but, in fact, they do, and that's the extent of it.

Measuring the outcome there and measuring the harm that came from that outcome is pretty feasible. I think that that's something that we can do. It's sort of a benefit of the bargain type argument. I paid you a certain amount. Some of that amount was for you to protect my data. You didn't live up to your part of the agreement.

It's the same kind of techniques we would use for a defective product, say, in a class action case. My washing machine is leaking. One of the deals is when you buy a washing machine, it doesn't leak. So you owe me some kind of compensation for that. One of the deals when I bought your streaming video service was you

249

wouldn't share my information with marketers. You did, so you owe me some compensation for that.

Where it gets a lot harder is when we start to introduce a lot more uncertainty, which is you didn't sell my information to a marketer, you lost it to hackers, and it's now on the black market somewhere. The possible harm that might happen, it could be well beyond the value of the product. I'm paying $10 a month for a streaming video service. If a hacker steals my information and drains my bank account, that's probably going to be a lot more harm than $10 a month.

Now, I will say some researchers have tried to measure this sort of harm by looking just at the value of the information, which is how much can you sell somebody's personal information for on the black market? What does it cost to buy a Social Security number? There's actually prices out there. You can buy someone's information.

I don't think that's a measure of harm. I think how much a criminal would pay for a record to open up a fake bank account or take out a payday loan and then run away, that amount is much less than the harm caused to the consumer than the criminal is paying for. Just like if somebody breaks into your car, they break into your car and they steal your sunglasses, well, you have that

250

loss, plus you have got to fix the window, and your car is in the shop for a while and so on. There's a lot of additional harm that's not measured by just looking at the price that that information can be sold at. So that I think is the hard case.

Once we introduce this uncertainty -- and we can measure things like the time spent changing records and changing passwords and so on, but that's -- I guess I'll leave it there. There's easy cases and hard cases, and unfortunately, most of them are hard cases.

MR. SMITH: Thanks.

So, Ginger, I'd like to get your thoughts on this, but also maybe if you can transition us into our next question, which is basically this question of causality. You know, when information has been stolen, how can we measure the risk and how can we kind of maybe sort of back out causality, what events may have caused that, or just get any kind of sort of way to clarify sort of the harm that happens in these kinds of cases?

MS. JIN: Yeah. I started the panel by emphasizing the difference between the ex ante perspective and the ex post perspective. In fact, estimation on the ex post harm I would say is extremely important for us to understand the ex ante expectation of harm, because we want to know what's the financial

251

loss of identity theft, and we also want to know how many percent of people actually suffer that loss, how many do not suffer that loss, right? We want to know the distribution of those losses across the population. That would give us sort of a big picture distribution so that we can form an ex ante expectation on this, okay?

But I will say that probably ex post harm estimation is not enough. We also want to sort of quantify the risk. I mean, I talk about -- the third panel, too, also talked about if a firm have risk, that a firm's data practice would increase, right? If there is an increase of risk, we want to measure that. We also want to measure how many people got exposed to that risk.

So I just want to point you to two sort of research directions I see as quite promising in doing that. We know it's very hard to tie back to a particular firm's data breach or data practice. One promising direction I have seen is a study by a group of researchers from Google, from Berkeley and from the International Computer Science Institute. They actually studied the dark web, okay? They track -- they are passively monitoring the dark web for a year, and they were able to identify 1.9 billion user names and passwords stolen from the previous breaches.

252

They are also able to -- if you look at their paper, they are also able to give a list of the top 20 leakages, which identify the companies that got the data breach and, therefore, got the records on the black market. So this suggests there is a way to sort of link a data breach to what kind of records are on the market.

Of course, we need other link from that records on the market to some real fraudulent transactions or identity theft or other sort of tangible or intangible outcomes that eventually happen to consumers, but obviously this is sort of a good one step forward to connecting the dots in the dark web.

Another study that was fascinating was about the blockchain. I think this might be a technology solution in the trackability of data flows. If blockchain can be used as technology to track Bitcoins ownership, maybe that technology could be used here to track sort of how the data change from one hand to the other.

I'm not a computer scientist, I don't know exactly how to do that, but I have seen other people, both computer science and economists, sort of try to work in this area. To me, it seems pretty promising direction to really do some research on.

MR. SMITH: Thanks, Ginger.

Lynn?

253

MS. LANGTON: So both Ginger and Josephine made the point that, you know, causal ordering is really difficult to establish when you're talking about individual incidents, and I would certainly echo that, and it's something that we've wrestled with quite a bit with the supplement. Are there ways that we can more directly try to tie experiences of data breaches to individual incidents of identity theft?

And the reality is, I mean, when you're collecting survey data, the data are only as good as the responses you get, and we know, because we already ask respondents if they know how their information was obtained, that the majority of victims don't have any idea. So we have about 30 percent of our victims that say they have some idea, even if they're not sure, about how their information was obtained.

So, you know, using a survey like the NCVS to try to get at this causal relationship between data breaches and identity theft is really not the best vehicle, unfortunately. I think you have to look for sort of these other technological, computer science-based solutions.

I mean, among those that do say they know how their information was obtained, about 20 percent of our victims say that it was obtained through personnel or

254

other files, personal information being obtained through a company that had their information, but, again, that's such a small percentage of the 30 percent that knew how their information was taken that we can't really use that to draw any conclusions. The causal ordering issue is a big problem for that.

MR. SMITH: Thanks, Lynn.

Catherine?

MS. TUCKER: Oh, no, I just -- I think, you know, Ginger and I completely actually agree in that, you know, what she was saying is that to measure a causal link, we have to have a better understanding of what actually causes injury to happen, and I think if you look at, you know, too much of research, we're trying to come up with an average effect rather than understanding what leads to those few incidents which are really bad that happen. So I think we're in agreement. I don't need to take any time.

MR. SMITH: Great.

Josephine?

MS. WOLFF: I would just say I think the flip side of that, which is also an important area for more research both from academia and from government and industry, is what prevents harm from happening, which I think is also a pretty underdeveloped area.

255

To take two examples that you have probably encountered, one are the chips in your credit cards where we had a big liability shift in this country a little more than two years ago about the kinds of credit cards you use and implanting those microchips in there to help prevent future data breaches from being able to steal those numbers as easily.

Another one you have probably encountered at some point is multifactor authentication, maybe for some of your accounts now you log in not just with a password but with another code or a tap on your smart phone or something, and I think that trying to understand what the impact of those types of defenses actually is in practice, right? Does it actually cut down the harm? Does it cut down the number of data breaches? Does it cut down what the injuries look like when data is stolen?

I think we have a long ways to go still in trying to understand which of these different mechanisms actually work for preventing the types of harm we care about, and especially if we're moving away from being able to hold companies liable directly for things like identity fraud, because it is going to be hard to trace back to individual breaches, I think being able to have very clear expectations about what the ex ante

256

protections should be, based on empirical evidence, is going to be really important in that space.

MR. SMITH: Thanks, Josephine.

Garrett?

MR. GLASGOW: Sure. Just briefly, I think if we are going to understand the -- understand causality and understand the increase in risk from any particular data breach, we need to understand the baseline risk, and I don't think we have a good handle on that right now.

The Identity Theft Resource Center has reported over 8,000 data breaches this year alone, more than 20 a day. Some are large, some are small, but there is this constant background of leaks of information out there, and it's perhaps impossible at this point to know, if your information shows up in some identity theft case, where exactly it came from.

To apply this to my -- I have this same analogy of privacy as Prince William Sound and the Exxon Valdez spills oil there. Maybe imagine we go -- Exxon Valdez runs aground, spills oil, the cleanup crews show up, and it turns out 20 other tankers ran aground there overnight and were towed away by the owners, but left all the oil behind, so, you know, how much damage did the Exxon Valdez do if it spilled oil where 20 other tankers have already dumped a whole bunch of oil?

257

Maybe that's a really negative viewpoint, but, you know, when we talk about the intrinsic value of privacy, maybe that value has been severely eroded by the sort of permeating background of radiation and continual data leaks, and that's one thing I think we need to get a handle on, is what's out there, how maybe in the aggregate we can see identity theft increasing, but, you know, how much of this is due to breaches six months ago? a year ago? We just don't know. So we don't have a good handle on baseline risk, and that makes it very hard to establish causality or increases in risk.

MS. CONNOR: This has all been incredibly interesting, and we have a few audience questions. The first one actually I think could implicate both consumer preferences and revealed preferences or stated versus revealed, potentially.

So the question is, are you considering studying the increase in purchases of ID theft products or cost for freezes, lifting freezes, delays to obtain credit while waiting for a freeze to be lifted?

And I don't know if anyone, in particular, would like to take a first stab at that.

MR. SMITH: I think the question could just be rephrased as would that be a way to get some kind of

258

measure of how much consumers value or how much they're harmed from ID theft?

MS. WOLFF: I think that's possible. I think one of the things that's complicated about purchases of those products is that most people who have them have them through a breached company. So the people who are usually sort of the big purchasers for those are the entities that get breached, who then purchase a contract for hundreds of thousands or millions of their customers, and it's then usually offered free of charge for at least a period of time to those individuals.

So there's some interesting work to do sort of looking at how many people actually take the company up on that offer and how that's changed over time, but I'm not sure it's exactly an economic decision on the part of the individual consumer.

MS. CONNOR: And, Josephine, I wonder if your answer changes when you're talking about personal cyber security insurance. I know you discussed it quickly at the beginning, but I know I didn't really know about it before I spoke to you for the first time, so maybe you could kind of explain it and whether or not a measurement of that would really change your answer to this question.

MS. WOLFF: I think that it could. I think at

259

this point, when we talk about personal cyber insurance, we're talking about probably about tens of thousands of people in this country, so it's a very, very small population of people who have actually purchased these policies. They are mostly fairly high net wealth individuals, so it's not something that low-income people are looking around to purchase or even most middle-income people are looking at.

I do think it's a different measure because it is something that people are paying for out of pocket. So it is them deciding, you know, this is a type of risk I'm concerned about, and I want to have some protection from all of these different ways of thinking about it.

I think the other reason it's interesting is if you look at how those policies are being structured, at least right now, they encompass a lot of things that I at least would not necessarily have associated as risks of sort of computers and cyber security. For instance, we're seeing families purchase insurance to cover home schooling costs in the event that their children are cyber bullied and don't want to continue going to school.

So I think it's this very complicated set of costs that's totally dictated by what people want and are willing to pay for as opposed to what we think the

260

company should be responsible for protecting you against, and that when you shift the question and you say, okay, what is it that we think people are actually willing to spend money on, it turns out there are a lot of different kinds of harm or injury that people might be interested in trying to insure themselves against and that at least some people are willing to spend money on, which I think does give you a sense of what people really care about and want to protect themselves against, but, again, at this moment, we're talking about a fairly small population.

MS. CONNOR: Okay, great.

And we have another question, and maybe, Catherine, given that you opened up talking about data breach notification laws, you might want to take a first stab at it, but the question is, can panelists discuss the multitude of state breach reporting rules and how that complicates setting a national standard of harm or injury?

MS. TUCKER: Well, I'll just start off by saying, so as part of this study we did about hospital data breaches and data breach notification rules, we spent a lot of time trying to decode the different texts of the laws, and there's an amazing amount of variation.

And the other thing I would say is that in my

261

very lay opinion, many of the laws seem rather inconsistent if you know anything about technology, in that there were exemptions which make no sense, and it looks like, I don't know, people, you know, were just sort of taking a random word generator sometimes to actually try and describe what they wanted to happen.

Now, maybe this is an opportunity for something better, but certainly I had a certain amount of disquiet, having seen the lack of standardization in these laws, and I'm an economist, not a lawyer, you know, so that's probably a bad thing.

MS. CONNOR: Ginger?

MS. JIN: I think given the discrepancy we see across states, there is definitely a value to standardize the notification law just to make sure that firms know what to do after a data breach, but I want to ask probably a harder question, that we've got a concern about what happens after notification, okay, if the firms sort of meet all the obligations stated in the law and have disclosed the information they know at that moment.

And so what -- that's -- we're relying on the consumers or the media or the public somehow respond to it, so that they would feel embarrassed and, therefore, improve their data security? I mean, that's -- to me,

262

that sounds like a pretty wishful thinking. As we know, if sort of the harms we have in mind cannot be traced back to the individual firms, and they already done their duty in notification, it's almost like, okay, I have done what I can, right? The rest is up to you.

So it's up to the consumers and the vigilance of sorting having their own preventive measures or other things. I think that question has got to be coupled with about the data notification law itself.

MS. CONNOR: Did anyone else have any thoughts on that question?

MR. GLASGOW: Just on reporting requirements in general, I'll say, you know, of course, we want -- if companies lose control of our data, of course, consumers need to know about that, but I think -- I have nothing in particular to say about any particular state's requirement, but one thing that we have looked at at NERA is the correlation between media coverage and both class action lawsuits and enforcement actions.

If your case goes viral, you can expect lots of class action lawsuits, and you can expect a lot of attention from regulators. So I think one thing we might want to think about is are we creating any perverse incentives with the way that these disclosure laws are written? Do companies want to try to game

263

them, to try to do things to minimize media coverage or meet the letter of the law without providing any information that can be seized upon by the media?

It's something that -- I mean, I'm not sure that we have created a perverse incentive, maybe it's just one of those things that goes viral sometimes or the biggest cases tend to pick up the most coverage, but that is one thing that I know that a lot of people in the industry are aware of, is the data breach is bad enough, but if the media gets wind of it and then the class action lawsuits and the enforcement rolls in, that's a whole bunch worse. So that's one thing to think about in terms of what kinds of reporting we require and how we require that reporting.

MS. CONNOR: Okay, thank you.

Well, we have a few minutes left, so I think as a final topic, we would love to hear what you all think should be the focus or the focuses of research on informational injury going forward.

And I know you have sort of answered this in your answers already, but maybe specifically what should researchers and the government work on to improve our ability to measure and assess informational injury. And maybe this time we will start at the opposite end of the table, so Josephine.

264

MS. WOLFF: I'm trying to think of a new answer that we haven't touched on. One thing that we haven't talked a lot about that I think is important for an agency like the FTC is thinking about sort of which injuries people have some protection from, right? So back when data breaches were mostly about payment card information, we had a lot of policy protections in place to say, if somebody's using your credit card fraudulently, you're not going to be liable for at least most of those charges, usually any of those charges, because your bank and your payment card network are going to want to keep your business. They are going to want to cover that for you.

Now, as we have seen, say, a strong shift towards ransomware and other types of cyber crime that are hitting individuals more directly, that they have less insulation from the direct economic costs of, I think there's an advantage there that sometimes it's easier to calculate those costs, because if you're talking about something like ransomware, you can put a price tag on it more directly.

But there's another question that comes up about how much do you sort of weight the fact that individuals are going to be bearing those costs directly and does that change the kind of injury or the kind of remedy

265

that you want to have in place, depending on why that ransomware infected my computer in the first place and was it my fault or should my ISP have had some inkling of that, and how the question of sort of distributing the costs in that way is going to change the incentives of the large, centralized, powerful actors who might have the most potential to implement widespread security controls or defenses?

MS. CONNOR: Thank you.

Catherine?

MS. TUCKER: Well, so, if I've got to say what you have got to research, it's just like a wonderful thing to be asked. I think I would flip the focus of research in that there's a very natural tendency to try and think, well, how can we measure most accurately an average effect? I might encourage you instead to focus on the question of how can we identify the occasions or incidents or people where there is no injury, where there is no effect, because I worry that we may be going towards a world -- which is very natural to try and understand how do we measure an average from average baselines and so on -- but I think this is such an intriguing field because so oftentimes we don't see injury, and trying to understand and pinpoint those occasions seems to me incredibly both exciting and

266

valuable.

MS. LANGDON: Interesting perspective.

So I think there have been a lot of issues raised today about how we can measure different aspects of data security and the harms associated, and in some ways, the NCVS, the National Crime Victimization Survey, is limited in that it can't address these full range of topics that we would want to measure and harms that we would want to measure, but I think as a government researcher, we want to try to measure as best we can what we can measure.

The thing that we can do is look at level and change over time in terms of identity theft, which is obviously a direct harm, and then we also have another supplement that I'll just make the pitch for that we just implemented, it's in the field right now, looking at financial fraud, which is another potential outcome of a data breach.

And so just to be able to track over time whether we're seeing any changes in the prevalence of these particular types of outcomes that we know are correlated, though we can't look at the direct causation and the direct causality to data breaches, but we know that there's still a correlation there, and to see if the risk of experiencing some of these outcomes and the

267

nature of some of these outcomes is changing over time.

So that's really what our focus is right now. We will be putting out data from our 2016 identity theft supplement in January, and then also in January, our 2018 supplement goes in the field. So, again, we're trying to measure this consistently over time so that we can track those trends.

MS. JIN: So one thing we haven't touched much on in this panel but has been touched in the previous panel is some similarity between the problem of privacy data security and the problem we have seen before in, say, food safety, drug safety, product liability, and tort laws, and so I would like to see probably more interdisciplinary research to sort of summarize the lessons we have learned from those areas and to see to what extent we can sort of apply the insights we have learned from those to the market of privacy and data security.

MS. CONNOR: Thank you.

MR. GLASGOW: And I think an important topic that isn't well understood, yet that we should push forward on, is maybe a theoretical or definitional issue of what we mean by informational injury. Does informational injury spring from the content of the information itself or is it how it's treated by, say, a company that you're

268

doing business with?

You know, so here's a thought experiment. Suppose you're doing business with a company and they have a data breach and some of your personal information is stolen and is now out there on the black market, but then you find out the previous week there was a another data breach with a different company and all that same information was already out there a week ago. Were you harmed by that second data breach or not?

And whether or not you were harmed might depend on what you think of as informational injury. Is it the fact that this information is now out there? If that's the harm, the harm is already done. Or is it this company is mistreating customer data and is not being fair to customers, is dealing unfairly, maybe that's a different kind of informational injury, and in that case, you could be harmed twice by the leak of the same information.

Both of those things could be sources of harm. It's entirely valid to believe they are both sources of harm, but I think that's something that is important to distinguish when we think about what kinds of enforcement we want to do and what kinds of harm we're trying to measure.

MS. CONNOR: Thank you, everyone. It looks like

269

we are pretty much out of time, unfortunately, but thank you to our panelists for such a great discussion and a wonderful way to end the day.

Andrew Stivers will be coming up shortly to give the closing remarks. So thank you.

270

CLOSING REMARKS

MR. STIVERS: Well, I'm disappointed to discover that they just -- they noticed the flaw in the agenda and actually put an end time to my closing remarks. In the earlier agendas, it was 4:45, that's when I start, no ending, but somebody discovered that, and they know me, so...

But most of you don't know me. I'm Andrew Stivers. I'm the Deputy Director For Consumer Protection in the Bureau of Economics at the FTC, and what that means is I head up the consumer protection mission for the Bureau of Economics. And the Bureau of Economics, in terms of our consumer protection mission, we advise the Commission on its actions and its policies, and we do economic research to help increase our understanding of the market practices that help or hurt consumers.

So from that perspective, this has been incredibly interesting to me, and each of the panels has built beautifully on the previous one, and I sort of felt like, wow, they are not going to be able to top that, and the next one comes up, and they, you know, reveal all sorts of great issues that we need to grapple with.

I, unfortunately, don't have time -- because they

271

stuck this limit on me -- to go into all the really interesting and important topics that were discussed today, so I want to pick sort of one issue from each of the panels to highlight, to leave you with, and that doesn't mean that these are the only things that you should pay attention to. I'm sure everyone has their own favorites, but these are things that called out to me and hopefully will help continue the conversation and the research that is clearly needed.

So the first thing, on the first panel, that I thought was incredibly good, especially for the FTC, is we're experts in markets. That's our mission. Our mission is to make sure that markets work for consumers. We have over 100 years of grappling with these issues, both on the consumer protection and the antitrust side of the house. So we really understand the kinds of injuries that occur contained within the marketplace.

But what's really interesting about data security and privacy is we see that this -- there's this flow between commercial applications and personal and private and social applications that muddy those waters a little bit, and in the first panel, they noted a variety of fallouts from the commercial space, the collection of data for commercial purposes into the personal and social space that cause these harms that are incredibly

272

important and that we need to grapple with as a society.

The other way is also true, that we see personal and social transactions, conversations between people, that because they're mediated through commercial information products or because those conversations are able to be monetized in some way, it really creates some complications for how we think about injury.

So an example of the latter would be in terms of our work in revenge porn. So there are these terrible interactions between former intimate partners, and that transaction, that kind of information being broadcast or those harms being visited are sort of amplified by the commercialization and the ability to monetize that information on some of these sites.

The upside for us, of course, is that we can actually go after that because it's in a commercial space. So I think that was really incredibly important for that first panel and highlights some of the research and thinking that really needs to continue.

The second panel in my mind really highlighted the definitional issues. So there were some hypotheticals put forward. The panelists had often very different reactions to the hypotheticals and sort of how they thought about those things. Some of those reactions to me seemed to be based on definitional

273

problems.

So Paul, right at the very end, at least tried to articulate a definition of what he was thinking about with injury. The moderators tried to say, no, no, you know, this is what we're thinking about, but sort of our interaction in this space or our place in this space often really informs what we think of as injury, what we sort of privilege as injury, and what we don't.

And so I think we need to do some more work in making sure we understand what it is that we're talking about, and I'll highlight one issue here. Is Geoff still here? I hate to call him out. This is the one thing I am going to kind of harp on a little bit. He said that risk was not injury. If that's true, somebody needs to tell the insurance industry that they have got it all wrong, because insurance is exactly about offering people the opportunity to mitigate a future, potential, unrealized outcome ex ante, before it actually happens.

So people are willing to take expensive actions to mitigate a future outcome, a potential future outcome. So to give Geoff his due, I think he was really thinking about, you know, what's the kind of injury that the FTC should go after? And maybe you agree with that, maybe you don't, but I think it's very

274

important to think about if a consumer, if anybody is willing to take an expensive action ex ante to mitigate the risk of a bad outcome, that's potential risk. If a firm takes an action, has a practice that motivates me to go out and spend that money, that's harm to me. So I think we just need to be a little bit more careful about how they think about those issues.

And then when we are clear about that, I think we can really see across the entire panel that was up here that where we choose to fall in terms of what's an injury really becomes a policy question, and that's this whole stew of the questions about, you know, are we addressing privacy and data security as citizens, as consumers, as all of the above, or both of the above? So there's definitely an ongoing policy debate that needs to continue.

The third panel brought up a variety of interesting issues, and it was hard for me to pick one because they -- I think of the panels, they really covered the broadest range, but the one that I will highlight, which is one that's really interesting to me, is this question of internal misuse versus external misappropriation, right?

So to use an example that I'm familiar with from some of my past work, when we think about economic

275

adulteration versus disruptive adulteration in the food industry, right, so economic adulteration is I want to water down some milk, right? I have a profit motive there if I'm doing that, and the last thing I want to do is get caught. So there's going to be some limits and different incentives in terms of how I approach that -- if I am an adulterer -- that is very different from if I'm a disruptive adulterer, right? I want to create some huge disruption in society. I actually want to get caught. The whole point is to let people know that I've done this terrible thing.

So it's not an exact analogy, but it at least highlights that there's very different incentives, risks, and potential remedies when we think about internal misuse of the data, internal privacy -- internal to the company uses of the data that somehow fall afoul of consumer expectations or deception or whatnot, versus these harms that arise potentially because the firm has collected this data and not secured it, but is really being mediated through these outside actors, these criminal actors.

So we need to think very carefully about what the incentives are here, what possible remedies might be put in place, and the liability associated with these things.

276

Lastly, the final panel -- and I could just applaud again, because as an economist, this was exciting and really good and points to all sorts of really interesting and important research challenges that we need to grapple with, and I hope that you folks, as you have access to research skills and funding, really try to push the envelope on -- but I will kind of sum up some of what I heard that highlights sort of the really important research questions and measurement questions that we have here.

We need to know the kinds of injuries, the kinds of actions and outcomes and practices that we're dealing with here, and I think we addressed that in some of the earlier panels, and we're getting a better sense all the time of what this looks like, but we need to be able to measure that pretty carefully.

We need to measure the severity and the impact of these outcomes. That has all the complications that the panelists today talked about. We need to know the incidence of outcomes. That was discussed as well. But we also -- and this was touched on, but sort of the last two points are, I think, of equal importance -- we need to know what confidence we have in the causal link between the practices that we're attempting to address and the outcomes to consumers.

277

So the more confident that we are that there's a link between those practices and the particular outcomes -- or even the practices and a risk of a particular outcome -- the better off we're going to be. So in the world of, say, drunk driving, we have really, really good data on the risks associated with drunk driving, right? So we are very confident that there is a causal link between an increased risk of death and drunks on the road.

Now, there's a really small probability that you'll actually die in an accident that's caused by a drunk driver, but the causality in terms of the increase in the probability, in terms of the risk, is really well understood, because we have really good data, and we have been collecting data for a long time. We need to make sure that we have as much confidence as we can get in the causal link between the practices and the outcomes.

And the final thing that I'll mention here is we need to have some sense of what the market incentives are. Do we think that the market is going to take care of a problem? Do we think that there are incentives in the market for actors to mitigate the risks, to somehow prevent these harms from occurring, or at least mitigate them to the extent they can be, or is there some sort of

278

market failure that would really call for a regulatory approach?

So I want to thank the panelists. I want to thank all of the participants for lending your time and your expertise to this endeavor. We continue to improve our understanding, and we want to continue to mine your expertise. So to that end, there is a 45-day comment period, and it's going to be ending, I think, January 26th -- is that correct? -- okay, thank you, so that ends January 26th, and I encourage your input. You can go to the FTC website to find information about how to submit your comments if you're interested in that.

Remember that our conference highlighting new research in the privacy and data security area, PrivacyCon 2018, is coming February 28th. So please mark your calendars and show up for that. And the other thing is that PrivacyCon 2019, which I hope will also happen, is going to be coming up sooner than you think. We start asking for submissions, assuming we do this, in the late summer. So if you're inspired to do some research now, think about how you might apply to get that research put in front of us for PrivacyCon 2019.

All right. Finally, I want to thank the Acting Chairman and her staff for pushing us on this issue, for instigating this workshop, because I think it's been

279

fantastic. And one final time, I want to thank the boots on the ground, the team that put this all together. Jacqueline Conner -- is she here someplace still, there she is, waving -- Cora Han, Doug Smith, and Dan Wood. Thanks very much.

(Applause.)

(Whereupon, at 4:58 p.m., the workshop was concluded.)

280

CERTIFICATE OF REPORTER

I, Susanne Bergling, do hereby certify that the foregoing proceedings were recorded by me via stenotype and reduced to typewriting under my supervision; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were transcribed; and further, that I am not a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of the action.

s/Susanne Bergling
SUSANNE BERGLING, RMR-CRR-CLR

B

C

D

N

P