FTC Informational Injury Workshop
December 12, 2017
Segment 3: Panel 2
Transcript

MANEESHA MITHAL: --take their seats. We're going to get started again. 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's 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 of 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, in 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're going to be talking a little bit more about the responsibilities of commercial entities that collect and store your data.

And so what we're going to be doing is we're 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 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's 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.

NEIL CHILSON: Thank you very much, Maneesha. Thanks to our panelists for being here, and thanks to all of you. So yeah. So We're going to do a hypothetical here. And when the panelists, as I read this along, 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'll be discussing why you identified injury at that particular point.

So with that, onto our privacy hypo. So in this hypothetical.

MANEESHA MITHAL: I'm sorry. While we're getting the technology cued up, I just want to give one disclaimer, which is that we're really not here to talk about the law and the legal ramifications of Section 5. We're really here to talk about injury as a policy matter. So again, when people raise their hands, this is not what qualifies as injury under Section 5, but this is when do you think injury has occurred.

NEIL CHILSON: Right. 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 onto the privacy hypo. Of

A pharmacy uses retail tracking in its stores to determine the most effective way to display greeting cards.

JAMES C. COOPER: I got to stretch. My arm's going to be up for a while.

MICHELLE DE MOOY: We have to hold our hand up.

ALESSANDRO ACQUISITI: Oh, keep.

NEIL CHILSON: Yeah.

ALESSANDRO ACQUISITI: Oh.

NEIL CHILSON: 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. The marketing company then uses the data to target advertising to identified consumers, including Carl Consumer.

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

MANEESHA MITHAL: James, is your hand up or down? It was up.

JAMES C. COOPER: It was up, and now it's down.

MICHELLE DE MOOY: It's resting.

JAMES C. COOPER: It's resting. Yeah. It's resting. I'm very weak.

NEIL CHILSON: The ad mentioned that Carl Consumer recently purchased this product. A local insurance company gets this information and raises rates for Carl Consumer. Carl Consumer's employer sees one of the ads and fires Carl.

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

JAMES C. COOPER: --said anything.

NEIL CHILSON: Yeah. Practically. 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 run down the line here and have each of you explain why you raised your hand when you did, starting with Alessandro.

ALESSANDRO ACQUISITI: Well, this is my thinking. Clearly if you're defining injury or harm specifically as realizing in quantified economic harm, I guess I suppose that most of us, even the ones who raised their hands at the first scenario there, would agree that there was no realized quantifiable economic harm.

However, that would be a very reductionist definition of injury which would ignore over 50 years of 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 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 broken. And when the boundary has been broken, well, that can be considered an injury.

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

So we jump 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 think 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 we can get into it later. I probably can pause here, let others talk, but I would like to go back to the issue of why quantified economic harm is so hard.

NEIL CHILSON: Great. James?

JAMES C. COOPER: Thanks. And thanks for inviting me. It's great to be here. So I raised my hand. I went up and down a lot. And so one, two, and three, we still have that aggregate 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 it's not really known-- nothing's been revealed about me.

And I'm willing to even entertain the notion that you may want to keep your interest in greeting cards private. I mean, I'm not here to dispute-- this isn't about well, that's obviously innocuous, who cares. I mean, someone could legitimately have utility loss from having people see the greeting cards they look 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-- and I put my hand up here-- and I think it gets to be a closer call because at that point, you're taking this aggregated information and somebody is saying OK, well, now I want to find out more about Carl.

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, if we're thinking about the kind of harms that can flow to privacy or not. Talk about, I think, 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, just to sum up, 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 harm, the things we think about with privacy.

Now, when we get into number five, at that point-- and I guess it was maybe unclear with the hypo. This is the kind of thing one of my students would come to me. It's like, I don't understand. It wasn't clear. That's why I missed it.

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

JAMES C. COOPER: Yeah. Yeah. Well, there are in my finals. Yes.

So anyway, here there are two potential harms. So 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 with the greeting cards, maybe there's not a harm.

With the HIV test, I still think it's a targeted ad. 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? Like the Target baby ad thing that people talk a lot about.

But I have a harder time with targeted ads in general. I mean, it's an intrusion into seclusion, maybe. But I have a harder time there. And the same thing with number six here. It's like, OK, 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 are not necessarily linking that-- maybe somebody's light bulb goes off and says, oh, I'm getting this because I'm friends with Carl. I think that's maybe too tenuous a link. But my hand, of course, then goes up 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've made a prediction that Carl is purchasing these HIV tests. And it's certainly something very sensitive. And you're telling all of his friends and his network, 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 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 knowing something about you.

Now, finally, when we get to eight and nine-- and this is where I make the 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 you. I've written about this a little bit. I call it intrinsic privacy harm.

But it's essentially that dignitary harms, loss of autonomy. All the loss that you feel when someone knows something private about you or you lose control of your information. Now, then there's the outcome that when a third party knows that about you, they act on that information.

And I know this is maybe something we'll 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'd like to pay a lower insurance rate. And the insurance company now knows that.

Now, it's a consumer harm to Carl, yes. He's paying higher insurance rates. But it's a reduction in adverse selection. So in some ways, surplus for society as a whole is getting larger. Now, I know this sounds cold-hearted. I'm not saying oh, well, you know, poor Carl.

And then 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 that, legitimate, illegitimate reasons. But again, if they're acting on truthful information, but I wouldn't categorize that as a privacy harm. I mean, there are very legitimate reasons why we as society prevent classification 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 there's a distinction in that, again, we can use discrimination law to get at these. We can use other means. And this panel is not about law. But I think that to me, in my mind, there's a distinction between those two. And I'm probably going on too long.

NEIL CHILSON: It's fine. Michelle.

MICHELLE 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 funny in a sense, it's not that I think that this person has been injured in a physical 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. 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. A lot of the discussions around privacy, and particularly I think in FTC cases, 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 benefit for this transaction?

Also, did he have access to or understanding, awareness, of what was occurring. This goes along with expectations and consent. 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 in 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 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 my hand would be raised for the rest of them, of course, but for different reasons, similar principles but different reasons. For example, the aggregate to me is meaningless. The fact that it's aggregated, it's one method but it's very meaningless when it comes to protecting information and protecting it from identification.

For example, how much data was there in this collection? Maybe there were three people who bought an HIV test, and therefore Carl is pretty exposed in that aggregate data set. So 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. I don't think that sensitivity of data is everything at all by any means. But I do think when you're talking about health information, and particularly highly sensitive health information like HIV status or 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 raises more risk in terms of harm.

I think finally, the idea of whether or not a person has recourse. 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 information to be marketed, or 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.

NEIL CHILSON: Great. Geoff?

GEOFFREY MANNE: Thanks, Neil. And thanks everyone for having me here, and for coming and 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've been thinking as I've been listening to people talk is that people are identifying something as injury, not the sorts of things that we would all clearly understand as injury, in ways that 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 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 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 all the way up until at least number seven, my sense here is that anyone who says 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 company advertises HIV tests to friends and associates in Carl Consumer's social network.

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 tests were revealed to people who could identify him, to people he knows, or something. The fact that the marketing the test to friends

and associates in Carl Consumer's social network-- and I'm assuming someone with the name Carl Consumer has lots of friends. So not going to be clear that it's him.

That may indeed increase the risk that someone will be able to figure out that he 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 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 benefits, and talking therefore about net injury. James started to talk about this a little bit. And I don't 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 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 careful about that. But again, 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 it. That's a little bit harder.

NEIL CHILSON: Great. Paul?

PAUL OHM: So all four of the co-panelists did such a great job dealing with this. And I'm going to spend most of my time just responding to things that have now been put on the table. Nobody, I think, if I recall, tried to 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're 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 every single one of these, we can point to something that is an injury. Now, I think a lot of Geoff's comments betrayed the idea that they may not be injuries that we want the legal system and the legal apparatus and an enforcement agency to be able to remedy. But that's answering a different question than has an injury occurred.

So let me just address some of the things that were said. 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 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 context. 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 in broad economic terms. Layer on top of that in the way Alessandro urged us to 100 years of writing about emotional distress and anxiety, the 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 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 is itself an injury. And I say that 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 Michelle capably rebutted-- the idea that the word "aggregation" is some 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 is talking about is this 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 unusable 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's so much rich information about your transaction or transactions of a few people that we're going to be able to sell it to advertisers and insurers and employers, then that also means that the risk of privacy is still latent within 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 remember 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 older people who go to a particular pharmacy. I think that should factor into the way we assess the harm.

MANEESHA MITHAL: So there's a lot to unpack here. And Geoff, I'll give you a chance to respond. But let me just tee up the question I think everybody wants to talk about, which is, OK, you've 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'll 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.

GEOFFREY MANNE: Can I just respond to one thing that he said first before we get to that point--

MANEESHA MITHAL: Sure, sure, sure.

GEOFFREY MANNE: --that I think will help, just because it's not about that. I do want to just 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 injured by whatever might happen, befall that drugstore. So I don't think it can be that.

I also think that it's essential-- and I like that Paul actually tried to identify a little bit of the kind of disutility that one could experience from even a risk of harm. This idea, the knowledge of risk can create 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 yeah, 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, 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. And 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 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 that 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 to 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 separate risk.

PAUL OHM: So I'm 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. The courthouse doors are being closed to 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. And 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 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 and balances.

I think it behooves the FTC to step up and fill the gap of the closed courthouse doors that I'm referring to. And so let me just go through two really quickly. HIV, right? HIV is not only a 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.

It's also tied intimately to sexual behavior, right? And so to 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 you lingered by the HIV test shelf for a while, or maybe if you tie 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 of tread into significant sensitive information.

And Michelle 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 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 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. So that's number one. Let me give you one more.

We think a lot about 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 or arguably violates some law. I think it's become common for retailers and stores to look at Mac addresses emanating from smart phones. I know there's been some activity about this in the FTC.

That is an easy violation of the Pen Register Trap and Trace Act, right? Congress in its infinite wisdom has said that this is a misdemeanor crime. Now, there's 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 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.

MANEESHA MITHAL: Geoff.

GEOFFREY MANNE: OK. So just very quickly, I think one of the things that Paul said, and that is at issue in this hypothetical with the retail tracking 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 course.

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 over-enforcing the law and over-deterring investment and innovation. And so one of the big problems I see with over-enforcement is over-deterrence of experimentation in the 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. That they may actually prefer. That they may be willing to pay for. Any number of relationships 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 for related reasons 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, even the increased risk of 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 collecting data, we're going to over-dramatically deter that, we're 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 where we actually can really identify that there are viable harms here.

And again, in some cases we're 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. It's a great start. But I think there's years of work to do beyond this before we should start identifying the government should be deterring these examples of data collection.

MANEESHA MITHAL: OK. Michelle.

MICHELLE DE MOOY: I just want to push back on one thing that you said earlier, that there is a relationship here. And I think that that's debatable. And in most cases, I think it's debatable whether the-- usually a relationship involves at least two parties, and I'm not sure that Carl is aware that he's in a relationship here, right? Maybe it's a stalking relationship. I don't know.

But so I think that that's an important point to make, that his expectations, his understanding of the situation, is probably different from the tracking of 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 the level of risk.

And I think also, of course, tied to consent and the person's expectation. The government already intervenes when it comes to sensitive information. So I agree with Paul that sensitive information should trigger obligations. 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 the assessment for whether it's 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. 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.

And then just generally speaking, I fall, I think, in the category of the idea of Professor Kalo's rubric on this, which is that there is subjective harm and objective harm. So subjective harm is the perception of loss of control that results in 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.

JAMES C. COOPER: Thanks. So I think 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 Carl. Well, again, as I mentioned, he may be harmed if people know about his greeting card habits.

I mean, legitimately. I know it's hard to say that with a straight face. But by the same token, I mean, there's no accounting for taste and everyone has their utility function and economists are--

MICHELLE DE MOOY: Maybe he's buying it for his mistress.

JAMES C. COOPER: Right. Yeah. Yeah. Good. Yeah.

MICHELLE DE MOOY: Thank you.

JAMES C. COOPER: Shop up on next year's final. But when we think about regulation in government, we can't make-- at least not in this world yet-- we can't make individualized rules, right? So we have to look at distributions. We have to look at the distribution of where we draw the line.

So I think that in this case when you talk about the greeting cards, we think, well, 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 habits. But it would be hard to see for me government intervention, especially some sort of aggregate, or even individualized. 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, you seem to like these greeting cards. Try this.

Because again, there may be some people way out on the tail. But we can't individualize our rules. And what I'll echo is something that both Paul and Michelle, and I think Geoff, maybe-- maybe not. But I do think that 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-- that's very, very sensitive information.

And I think here when you think about what's the right enforcement posture, where would we have government intervention, I think you have to balance a lot of things. I mean, first there's just the direct utility harm from Carl. OK. So would this data about my interest or actual purchase-- again, not clear yet-- in HIV testing is out there. That's legitimate loss.

There's also dynamic losses. And I think this is maybe the flip side of what Geoff-- I like what Geoff had said. OK, this information is out there. And 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 protected behaviors that keep me from learning that valuable information. So there's that dynamic part, and it's related to autonomy benefits from privacy. How do you act under observation versus not observation.

So these dynamic things you have to consider. But they're also beneficial. The data being out there, 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. Speaking as economists, I mean, you want to unite information and control the person who is 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 actually acquiring the information about his HIV status, which can be quite beneficial.

I also agree here-- I want to echo something Geoff said-- as I do think that it's better to focus on uses rather than collection of data, if we're thinking about a regulatory posture. So I think that if we're concerned, say, about numbers eight and nine, what the insurance company is going to do with this data or what the boss does. I mean, 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 you're doing a lot of sensitive things in drugstores. So no retail tracking in drugstores. That could be one, because there could be this cascade of risk.

But we could also just say, if we're concerned about HIV status, look, as an insurance company you can't base insurance rates on HIV status. If we're concerned about that outcome, we could say that employers are not allowed to fire based on HIV status. If we were concerned about that. We as society make cuts like that all the time, that there's certain information you can't act on.

I mean, I think rather than suppressing truthful information that it's better to just prevent the uses of that information we can. And I think that has the two benefits in that. So one is that if we don't want insurance companies to use this information, but we don't 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.

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

And so there have been a couple Ban the Box, which is you can't-- there are about 25 states have them. And they're in different forms. But it basically says you can't look at whether you have a criminal record on a job application. And there's a good study by Agan and Starr, then another one by Doleac and Hansen, basically finding the idea Ban the Box is incredibly well meaning.

It's let's break the cycle of I go to prison, I get out, no one will hire me because I've been in prison. It's this catch-22. So I just go back to criminal activity. So let's break the cycle. 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 engage in Ban the Box? Again, two really good studies that are out there. So that, well, discrimination against African-Americans goes way up.

I mean, 10 times lower callback rates New Jersey, New York, after Ban the Box than before for African-American males between 18 and 25. So I say that by there are costs to suppressing

truthful information. People want to make decisions based on info. So the way to go is not say, can't collect it. Just say you can't use it. So anyway.

MANEESHA MITHAL: OK. Alessandro?

ALESSANDRO ACQUISITI: 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, and needs to go back as I often do when I work in this area, to the seminal work on the economics of privacy coming from Chicago school scholars in the '70s such as Posner, Stigler. They pointed out that privacy protection is inherently redistributive. 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're going to affect the distribution of wealth. If you do not, you're 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 a case of privacy they work that optimally. For first, we have ample evidence which we have described in the Journal of Economic Literature Review we published last year. We have [INAUDIBLE] and [INAUDIBLE].

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

And 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 render 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' interest 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 matrix

for government intervention. If a majority of voters think that privacy is important, perhaps we should listen to them.

MANEESHA MITHAL: OK. 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. We'll do the same exercise. We'll read out a sentence from the hypothetical, and just raise your hand when you think that there is injury, OK?

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. Two. Unauthorized access occurred, but confirmation that no consumer data has been exfiltrated. Unauthorized access has occurred and it is possible that consumer data has been exfiltrated.

Unauthorized access and consumer data from company A has been found on the dark web, but there's no evidence it has been used for fraudulent purpose. [INAUDIBLE]. And then finally, unauthorized access and consumer data from company A has been used for fraudulent purposes.

OK, so let's see. So why don't we 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. Drill down.

JAMES C. COOPER: OK. 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.

And certainly number five, you're there. But one through three, at that point, it's too speculative to me to say. Again, 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 legal intervention, or where we are in the hypo. But I would say one through three isn't sufficiently cognizable in my view.

MANEESHA MITHAL: Why don't we go--

MICHELLE DE MOOY: Am I next?

MANEESHA MITHAL: OK. Go ahead, Michelle.

MICHELLE DE MOOY: OK. So where I land on number two is of privacy violation. Perhaps lesser on an 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 whom consumers have no relationship do unfair conduct, practices prey on vulnerable consumers involve coercive conduct or create significant information deficits."

And so my assumption in 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 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 I think it just depends on which way you lean, which way you decide to assume. And I think if you come from privacy as a core principle, you lean towards the protectionist idea. And so therefore, number two would not necessarily trigger government intervention or laws, but that it might trigger some awareness, some 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.

GEOFFREY MANNE: Let's see. I raised my hand with the last one. And even it is actually somewhat questionable, not that there's injury. There It's defined in terms of injury. But in part to respond to what James said, I just want 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 hand either way. So no one thinks it's injury? 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. The language they used in one place was we actually think it won't make any significant or noticeable difference.

Again, let's just take that as true for the moment. It may very well not be. 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. And number four too.

But it really matters-- to me, anyway-- 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 problematic here, especially in the data security context where, again, we don't know anything about any of the conduct here. We know that unauthorized access occurred, for example.

Well, I can tell you that there is a non-zero 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 again, in the narrow confines of this initial question, that doesn't actually tell you anything about whether there's been an injury or not. 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 if we're going to be talking about risk, for example, the risk of injury as being an injury itself, I don't think 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 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 things are really problematic then. So just to clarify, so maybe five. Probably five. But nowhere before give.

MANEESHA MITHAL: So Paul and then Alessandro and then James.

PAUL OHM: So I've warned the moderators that I'm going to fight the hypo. I'll fight it very quickly because I know we're short on time, and then I'll answer it not fighting the hypo.

Let me put another hat on. 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," and "it's possible that--" then you know that serious violations have occurred because those are naive statements. Those are impossibilities.

You 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 sysadmin said we have an ids those are the companies that I guarantee you are 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.

But let me not fight the hypothetical. So another way to respond to Geoff's comments is do we find a world in which Equifax has occurred to be an acceptable state of the world? 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.

Before Equifax, a company a week would demonstrate to us time and time again that, to speak like an economist, there are externalities that need to be internalized and they're not under whatever mechanisms we have. The state of data security in this world is horrid, is horrible, and it's causing concrete repercussions to everybody in this room. I like Elizabeth Warren. I like saying I like Elizabeth Warren.

Spent three hours on the phone and on the web site trying to 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 that I've thought. People face financial ruin. People self-chill. They change their behavior. They don't apply for jobs they might apply for because they're worried about their credit report.

They don't try and buy a house because they know they're not going to get a home loan. There are documented cases time and time again. And Geoff has impressed on us several times that we should look for documented cases. They're 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. 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. 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. They've 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 in making these hard choices we should think about how good was your security at the time.

ALESSANDRO ACQUISITI: I raised my hand at scenario one for reasons similar to the ones Paul mentioned. And very similar to the ones I brought up under the previous scenario. This idea of boundaries, and whether the consumer even knew that the company, company A, stored his or her SSN, whether the company A had a right to in fact have this information. How did it acquire this information, why and how it is using, et cetera, et cetera, et cetera.

Now, again, I would agree that not necessarily quantifiable realized economic damages occurred. But the possibility of and in fact the increased risk of a downstream cost has emerged, which allows me to go back to a promise I made earlier. 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'll 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 the other 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 those benefits.

And then there are also other differences, other nuances which again make it incredibly hard to pinpoint quantified harm. There are costs which are exceedingly small, but upping 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 long period of time amount to a huge waste of time due to the fact that when I load a page on the internet, the page is loading slowly because of the tracking going on behind the browser.

On the other hand, on the opposite end, there are the costs which are catastrophic with very low likelihood, such as catastrophic medical identity theft. And then differences between the harms which occur immediately after some privacy invasion has occurred-- in your example, scenario

one, Carl's employer firing Carl immediately after this information has arisen-- and the harm which may happen months or years after the fact, such as maybe someone suffering from the Equifax breach one year out 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 that the idea that we can create just a simple matrix to capture them all and their simple formula that the regulators can use to decide oh, yes, I should intervene or no, I shouldn't intervene, is really hard. It seems almost too haphazard and too untenable to me.

NEIL CHILSON: Yeah. 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?

JAMES C. COOPER: Yeah. Well, I guess since I put my tent up, a lot has been said. First, I'm going to agree with both Paul and Alessandro. I mean, these economists, as far as internalizing the externality, it's hard to think of how this can happen. But for all the link between I committed-- you think of the normal tort. I drove carelessly. I hit somebody. They got hurt. That's easy causation, so you calibrate tort law to internalize that.

That can't be done here. I mean, it is really, as far as linking, 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 goes off on something Michelle had said, and a theme throughout of there's data everywhere. We don't know if there's a breach where that data came from.

I mean, what ultimately is a system. If this is something we just live with, I mean, do we think of this as a first party insurance world where we just all either self-insure or buy insurance policies against cyber risk?

And then we reduce some of the incentives that the tort could potentially bring. But it's hard, for the reasons that Paul actually talked about in the previous case. Many or if not most data breach cases get thrown out either for lack of standing or lack of meeting, pleading harms sufficiently. So the tort system, it's unclear.

So does it make sense to, rather than being insured through the tort system, to have 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 amplify something that Geoff had talked about, is risk. And I do think that this is the big question, as Geoff said. I raised my hand at number four. But you raise a valid point, that number four, even number five. So there's harm. I mean, you're fighting the hypo. It was unclear. Again, next year we'll write the questions more clearly.

GEOFFREY MANNE: You need six pages for each of them.

JAMES C. COOPER: Yeah. Yeah. Well, that's how I want--

GEOFFREY MANNE: And then we'll--

JAMES C. COOPER: But the idea that it's going to be really hard to link up. So it's on the dark web, but 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. And so I think about this.

And like probably everyone in this audience, I think about this as like a Bayesian updating problem. It'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 start off with some prior view of the world for the odds of my data being misused, part of a breach. And then 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 has engaged in likely to lead to harm.

And so you think about you update your 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, how much does what we know about the likelihood that this conduct is related to harm change our priors. And it could be that this, the delta, the change in the odds of harm, are really, really high. It could be 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, it moves from the odds of harm from this breach being 1% to 3%. So at that point, do we look at the delta, which could be really large? These are like epidemiological studies where you start with a really low baseline of some kind of condition, and then there's a drug. And the drug reduces that condition by four times, but it just goes from 3% to 2.5% 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 is a case whose name can't be spoken up here that that's one 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 we haven't really gotten into this, but the extent to which the conduct has been out there and it hasn't happened for a while. I think that does inform. But anyway, I see other people with their tents up, so I'll go on.

MICHELLE DE MOOY: Am I next? Or Geoff?

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

MICHELLE DE MOOY: Thank you.

GEOFFREY MANNE: I think there's a big problem in--

MICHELLE DE MOOY: --epidemiological?

GEOFFREY MANNE: Yeah. No, no, actually.

MICHELLE DE MOOY: OK.

GEOFFREY MANNE: No. I think 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 risk of harm, 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 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're willing to pay, you haven't yet established that we're actually at a point where we should be intervening more or identifying more. Looking at changes in risk of from 1% to 3% is actionable. 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.

MICHELLE DE MOOY: So I reject that, and I'll get to that in a second. Just a second.

GEOFFREY MANNE: It's tautological.

MICHELLE DE MOOY: One thing I just want to say, just to push back to something that you said, James, about tort, which is that the tort of assault requires imminence. So I think you were saying that that wasn't a part of that assessment when in fact it is.

JAMES C. COOPER: No, in fact, I thought I said that 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--

MICHELLE DE MOOY: OK.

JAMES C. COOPER: So I think we're in agreement.

MICHELLE DE MOOY: I just want to finish my thought. 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. I mean, we can look at social harm in the 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 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 fatalistic feeling that we can't actually--

GEOFFREY MANNE: Optimistic. No, I'm saying we have to do it. We should do it.

MICHELLE 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 that 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, under the FTC Act, "substantial injury cannot be reasonably avoidable, is not offset by benefits." So all of the 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's not necessarily possible.

Many, many people that I spoke to had no idea that Equifax existed or had data on them. So I mean, the information asymmetries, the lack of a level playing field, I think, is absolutely crucial. That you cannot just go past that and say that 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.

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

MICHELLE DE MOOY: Oh, sorry.

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

GEOFFREY MANNE: I have more to say.

MICHELLE DE MOOY: Reserving my right.

NEIL CHILSON: So one thing, tying together the responses to the two sets of the hypotheticals, while Paul openly admitted that he was pushing back against a 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 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 laid out a boundary framework that when you cross a boundary, that's a type of harm. And I think it

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

So I just want 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.

MICHELLE DE MOOY: Please.

ALESSANDRO ACQUISITI: Well, my two points on this would be that there are violations which may not arise economic harm, but other forms of harm, on autonomy, on dignity, on 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 without materialized damage is enough for intervention or not. That's a fair point to debate.

And further, 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 essence my argument. And the last point I would make, if there are a few more seconds, is that I would suggest that as important this workshop and this type of panel is, 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 intervene to protect.

But the burden is twisted around, and puts 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. Are only to you, or to consumers, to society.

So fundamentally 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 and used, and then societal level. Because as much as we like to believe in big data and analytics, much of the internet data economy right now is a black box where we do not exactly know what is happening. We know the 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.

NEIL CHILSON: Great. Michelle?

MICHELLE DE MOOY: I think that was well said. The only thing I would add is that 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, like Alessandro

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 violations to harm. But I think that's absolutely what occurs, and I think has been borne out in example after example.

NEIL CHILSON: Paul. We have some questions from the audience.

PAUL OHM: Go ahead. I'll find a way to say my answer in response to whatever you ask me.

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

GEOFFREY MANNE: And while we're on that point, that reminds me of some unrelated thing.

NEIL CHILSON: So we have a couple of questions here. Some of them have been somewhat addressed since I got them, so I'm 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--

PAUL OHM: [INAUDIBLE] panel.

NEIL CHILSON: --do you all accept the notion that privacy is a critical component of democracy, as Michelle stated. A big picture question, but if you can tie it back to harm.

PAUL OHM: I totally have a segue to the point I was going to make. I absolutely do.

NEIL CHILSON: Start with Paul, then.

PAUL OHM: Yes. So let me start small, and I'll end up at the question. So I wanted to respond to the Bayesian brothers, the idea that we're going to examine the delta.

JAMES C. COOPER: I didn't know I had a brother.

PAUL 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 that if you're in a space where there really isn't much harm and there's a lot of responsible practice, and then there's one really bad actor who's below the standard of care, then the FTC has jurisdiction. But then as the world goes to hell in a hand basket and we end up in a cesspool where all corporate actors for whatever reason, malevolent or benign, are not protecting us, are causing anxiety, causing the kind of fear I'm talking about, then suddenly you strip 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 of times, but I think because he's an economist and because he's a polite Italian hasn't made quite forcefully enough, which then leads to your question. Which is Alessandro has repeatedly said that the economic toolkit can be very helpful when it talks about harm, but it should not be considered complete.

And don't mishear me. We still should be empirical and we still should be rigorous. But I think in many ways the economic toolkit 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 measure results that come from those other fields with what the economists say. And if you're only looking at the economists, you're thinking of this too 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. But there are ways to say that when we're talking 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 a responsible agency would take advantage of that and try and protect consumers in the way that Congress had in mind. So thanks.

NEIL CHILSON: Geoff.

GEOFFREY 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 we're talking about private entities here and our fellow citizens. And I think often-- not always, certainly-- but often the two are elided and they're extremely, extremely different in my mind.

And we do a real disservice when we say something like the drugstore in the first example knowing something about me is-- not that anyone said this, but one could say it's just as bad as the government knowing this about me. The next hypothetical after the insurer and the employer could be the government. And I think it's crucial that we keep those things separate.

With respect to the ability to keep information private from other people as being crucial to democracy, I don't even really know where to begin to answer 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 supersedes every legislature in the country and every other statute in the country.

And indeed, you could say that's what Congress intended. I mean, you'd be wrong, but you could say it. And the idea that trying to implement some idiosyncratic principle like democracy at the level 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're few and far between, really. That doesn't mean we shouldn't care about them. It 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, 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.

NEIL CHILSON: James.

JAMES C. COOPER: Thank you, Geoff, for going on long enough that I may get the last word.

MICHELLE DE MOOY: I assume I don't have to answer. But I will just reaffirm what I said before and say that I don't think that anyone was suggesting what you just said.

NEIL CHILSON: James raised his card. That's the only reason I moved past.

MICHELLE DE MOOY: Oh, I see.

JAMES C. COOPER: No, no. Go ahead. And then I'll say the last thing.

MICHELLE DE MOOY: OK. You get the last word. So I would say I don't think anyone was suggesting this lawless world of immeasurably damaging our democracy through the FTC's unfettered-- I don't think anyone is 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. And a lot of the data ecosystem violates these principles in different ways. And so therefore we can look deeply into those to figure out 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 is much blurrier than I think you were painting it to be. The government is acquiring commercial data all the time. They are working with private contractors all the time. So I think it's not possible to make that distinction per se. I mean, in law we can maybe. But in this discussion, especially when you talk about health data, which is more my area of expertise, the government is constantly selling and buying commercially generated information about people.

NEIL CHILSON: James.

GEOFFREY MANNE: To seconds James.

JAMES C. COOPER: OK. And Paul's wrong. No. No, no. I'm just kidding. The only thing I would say directly to Paul is you'd said that the FTC needs to incorporate a lot of other things other than economics. I think actually that's one of the issues is, as you say, legal scholarship needs to be incorporated. I think that there's really 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 the balance is certainly more on the other side. But the last thing I'll say, I agree with the question about democracy. I mean, I would agree with Geoff, I think. When it comes to privacy, it's vis a vis the government, not really vis a vis private, corporations.

And I'll leave with this. I mean, I think one of the big picture questions here is I completely agree that there are information asymmetries here. Alessandro's great body of work has shown a lot of this contextual dependence, a lot of biases, endowment effect exists in this. But asymmetric information and behavioral biases exist across a lot of markets.

The question, I think, the big picture question here-- and I'll just end on this-- is 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 the more it's informed with empirical literature, I think, the better. So I'll just leave it at that.

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

[APPLAUSE]

SPEAKER 1: Just a couple of quick announcements about logistics of lunch. if you leave the building to get lunch, you will have to--