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 arodehe bi-2 (A (-2 (l)-2 ((e)4 1 (he)m 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 hypothetica

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

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 ACQUISTI: 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 is 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

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 -- this isn't about well, that's obviously innocuous, who cares. I mean, someone could legitimately have utility loss in 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 clear 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

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, assert 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 whey. So

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, 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, 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'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

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?

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 -(s)(. H)(t)-22 (t)-2 (')10 mattt(i)-2 (l)-125(t)-22 (t)-2 ((i(?)-16 (H)-2 (r)5 0 (

that so you never see this enforced by anyone. But Congress said in the 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.

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 sen-1.14 Td [(J -5.24 -1.(g)10)-2 (he)4 rnmsen-14ntr (nk a)4 that up i fd tJ -5.2indihre raa-2 (e)4 (d)-2 bJn-14ntatonmcndirhret(he)(4-2m'v

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 econo-mists are

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

So I think that it's a different balancing. Speaking as economists, I mean, you want to unite

truthful information. People want to make decisions based on info. So the way to go is, not s
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 should 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 privac

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 fe(e)4 l-4 (r)-see- ct. We have the data se fze(hor)3 (t)-2 (, a)4 (nd (ha)4a(1(po a)4 (ho)( )-110.)4 y(d4 (h)10 (t (he)4 10 (y)h(ve)4 h)10 (t (he)e4 (nt)-12)4 (i)-2 (s)h A)2 (

And so my assumption in this and this could be incorrect

without talking about

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

one, Carl's employer firing Carl immediately after this information has arisen

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 the 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 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, 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--

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 some of 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

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.

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.

NI 4 >>BzrLe aate an re10 (r) EM7Tc 0I nd trjnks.

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.

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 bias, 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--