EMERGING TRENDS IN HEALTHCARE TECHNOLOGY:
PRIVACY ISSUES WITH HEALTHCARE TECHNOLOGY

SPEAKER:
PROFESSOR CHARLOTTE TSCHIDER, *LOYOLA UNIVERSITY CHICAGO SCHOOL OF LAW*

[edited for reading]

FEBRUARY 19, 2021

Casey Goggin: Next up we have Professor Charlotte Tschider. Ms. Tschider is current an assistant professor at Loyola University Chicago College of Law. She was previously a visiting professor at the University of Nebraska College of Law and the Jaharis Faculty Fellow in Health Law and Intellectual Property at DePaul University College of Law. In 2017, she was named a Fulbright Specialist in Cyber Security and Privacy Law by a Fulbright Scholar program. She received her J.D. from Hamline University School of Law where she was a member of the Law Review. She received her M.A. from the University of Minnesota Twin Cities, as well as her Bachelor's degree. Her primary scholarship is information privacy, cybersecurity law, and artificial intelligence, with a focus on the global health care industry. She has written or spoken about many topics, ranging from data collection in the medical industry and internet privacy to global data protection. She is the author of *International Cybersecurity and Privacy [Law] in Practice*.[1] And with that, I'm going to hand it over.

Charlotte Tschider: Well, hello everybody. It's just a pleasure to be here. I was actually supposed to present last fall and it's amazing how much can change in just six months. I think that this first panel was excellent in demonstrating how technology is really changing things for a variety of people in a variety of locations. And the investment in artificial intelligence and big data use has really transformed that to even a greater extent. And so, I'm hoping today that I can illustrate some of the challenges related to the current privacy regime.

But I'd like to start by talking a little bit about the technology. And before I jump into the slides, I always think it's fun to talk about certain scenarios that I've been faced with from a consulting perspective that kind of put things into a little clearer focus. So, one of the devices that I have consulted on in the past actually came out of Finland. And for many of you who might have some familiarity with the EU data protection directive that preceded the GDPR and the GDPR, one the challenges we faced in those spaces is related to data sharing, data use, and data reuse. And we know that data has become tremendously important for both the treatment of health conditions and for the technology that is associated with the treatment of those healthcare conditions. In this case, I was talking with a company and they had produced a really amazing type of technology where you can actually look at somebody's cornea and do it without the puff of air that many of us are familiar with from getting eye exams. And they said to me,

[1] *See* CHARLOTTE A. TSCHIDER, INTERNATIONAL CYBERSECURITY AND PRIVACY LAW IN PRACTICE (Wolters Kluwer, 1st ed. 2017).

"Well, what do we do if we want to use the images we've collected for other purposes?" And I said, "Well, what is the value in doing that?" And they said, "We think that it might be possible to diagnose early-onset Alzheimer's before there are ever clinical symptoms simply by analyzing the images using AI technology, and in something like 75% of the cases, we would be effective in doing that." And unfortunately, even just talking about this from an EU privacy perspective, which we know tends to be a little bit more, I'm going to say, advanced, perhaps, in the privacy space than what we see in the United States. Even in that scenario it was tremendously difficult to identify a justification for reusing those data. And yet we know that it could have enormous impacts on our ability to diagnose very serious diseases, get individuals on the right pharmaceuticals or other treatments needed to prevent the further progression of some of these diseases. And it kind of got me thinking, what do we do in the United States? How might we evolve our privacy models to better provide support for these types of technologies and other data uses?

And so, I'm going to start today by talking a little bit about the technology and then go into, what are the primary privacy considerations we have. And then how might we think about evolving those models, both under HIPAA and outside of HIPAA. And it was great that we had this previous panel because there was a great introduction into HIPAA and some of the considerations for that. So hopefully I won't have through too much detail there. I'm going to share my screen and hopefully you can see that okay. Alright, can everybody hear me? I just want to make sure you heard my, did you hear my introduction?

Paige Goodwin: No, I think you cut out as soon as you screen shared.

Charlotte Tschider: Oh, excellent. Okay, well I didn't say anything after the screen share, so I think we're alright. I wanted to illustrate at least for you where we're seeing these considerations around big data and AI in healthcare, at least initially. And then we'll go into some specific examples related to the technology implementations and some of the challenges associated with those technology implementations.

So, at least initially, we know that data are tremendously useful for operational support. Whether that is the efficiency of operating in a large health system, the cost and value analysis that goes into reimbursement calculations for Medicare and Medicaid, we know that there is a huge focus on quality. And understanding

how provisioning healthcare is going to increase or decrease quality is certainly a huge goal and a goal that actually is incentivized by a lot of our other healthcare laws related the ACA, MACRA, and others. And then there is sort of this benefit potentially in AI in administrative automation. So, if there is the ability to automate more administrative tasks, we might be in a scenario where we can repurpose staff, use hours in different ways, and certainly provide better quality healthcare to individuals because we're able to sort of shift things around.

But we also know that big data is heavily, heavily used in the diagnostic medicine area. Especially around imaging, we've seen incredible developments related to x-ray image evaluation and others. And we have diagnostic AI now that applies to treatments, you know, what is the best treatment for an individual, given characteristics of that individual that we've seen in other types of individuals who have received certain types of treatments? And we know that diagnostic medicine, for example, is often developed using base data and that base data usually come from health systems that are, what I would say, high-resource types of contexts. They're the types of situations where, for example, you have the best machines, and you have people who are specifically focused on certain types of cancer diagnosis who may be actually identified as the best in the world in doing that. And so, when we create this type of AI it's really wonderful. But how do you take that and apply it to new contexts? One of my colleagues, Professor Nicholson Price, has written a great deal on this concept. And certainly dovetailing from the rural health conversation we just had, think of the myriad of ways where using those types of diagnostic tools in rural contexts where you may not have access to highly specialized cancer diagnosticians. That might be tremendously valuable. But at the same time, you need to make sure that the data you have behind the algorithm is going to represent those new populations. And we've seen it in a variety of scenarios, not just in rural settings, but also in big cities, and situations where you potentially have more diversity of living conditions, housing conditions, and individuals from a variety of different backgrounds. So, we know that data are tremendously important in making sure that those algorithms are actually going to facilitate better treatment and facilitate better diagnosis.

And then finally, and this is the area that I spend a lot of my time focusing on, is artificially-intelligent-enabled Internet of Health Things. If you see IoHT, it was a term that was coined quite a long time ago, to represent healthcare technologies. So not consumer technologies, but consumer technologies that are really

oriented towards health. And the interesting thing and the distinct thing about AI in these technologies is that it actually drives the functioning of medical devices. So, the data you have actually informs, for example, what amount of insulin might be recommended for someone with an insulin pump to use. Or, what is the charge that we need, for example, for a brain stimulus device to reduce pain in an individual. All of those data can actually inform not only privacy concerns, but also safety concerns regarding their functioning. And ensuring that data free flows back and forth is tremendously important for the effective functioning of those types of devices. And often those data are considered a personal information, whether they fall under the de-identification safe harbor in HIPAA, or if they properly are identified as protected health information. So, we know that data are tremendously important.

Okay. Just switching slides here. Just give me a minute. It's a little bit slow. Alright, so as I figure this out, I'm just going to stop the share and reshare here a minute. My apologies, I'm having some network issues due to a lot of snow. So that makes this a lot of fun. So hopefully you can see my screen again. Let's see if we can get it to switch.

Alright, so instead of switching to the next screen for now, what I'll explain is that, when we're talking about medical devices, we're not just talking about the thing that is implanted in somebody's body, you know, the implanted pacemaker, for example, that enables somebody's heart to function properly, or an insulin pump that is pervasively attached to somebody's body, or a hearing aid that somebody wears regularly. In those situations, actually, we're not just talking about the physical thing that connects with the person's body or is inside a person's body. We are also talking about applications. So, applications though a mobile device or another type of user interface that's available to the individual. I know with insulin pumps, for example, there's usually a user interface that's sort of attached to the pump that someone uses to actually make decisions about how much insulin to deliver to their body. Something that I think was actually in clinical, the third range of clinical trials, or the third stage of clinical trials, was the artificial pancreas. And the artificial pancreas doesn't have a user interface in the same way that an insulin pump does. It's generally designed to function almost independently of the user.

Now, whether you have a user interface or whether you don't have a user interface, usually individuals, especially individuals with health conditions, trust the technology, or we're expecting

them to trust the technology. So, for example, if you have instructions that are delivered to an individual about an insulin pump that says, "this is the amount of insulin that you should deliver your body based on the value we've ascertained of your blood sugar," to what extent do we really expect that individuals will challenge that kind of a direction. Probably not, right? Individuals tend to believe what the technology tells them, so we really need to make sure that the information we have behind the technology, that are often stored in offshore locations, that are stored in big data implementations, such as Amazon Web Services and the like, and which use machine learning technologies in those locations, it can get a little bit challenging if the data are not correct or we don't have enough data, and if we don't have really any control or ability to influence third parties and their practices with regard to those data from a cybersecurity perspective. So, because we have this broad distribution of what a medical device means, we have additional challenges related to how HIPAA typically manages these types of scenarios.

Okay. So, let's try this one more time. Alright, I am seeing a chat, but I cannot get to it. Feel free to just jump in. Okay, can you see this screen now? I believe maybe you can. Aha, excellent. Alright, we are back in business.

So, from an historical perspective when we look at privacy, there are really four categories of privacy considerations that we have. The first is from a notice and consent perspective. In the EU we call this "lawful basis." In the United States it's just generally "notice and consent." That is the primary vehicle that we use across most privacy laws. So yes, there are additional requirements in any privacy framework that you have from a legal perspective, but notice and consent tends to be the most powerful, all in all. I'll talk about here in a second why that is maybe not the right focus for any privacy framework, including HIPAA. Although in HIPAA we have notice with a kind of a reasonable acknowledgment, at least at the federal level. Most of the states have an additional consent that's sort of added on to that. But outside of that, so under general Federal Trade Commission jurisdiction, and what we're increasingly seeing at the state law level, is that consent is usually required. And I think that there are some limitations to that, both in terms of data usage, and the practicality of managing those processes, as well as just the efficacy of consent and how that works.

We also have this focus on data minimization. So not collecting, using, retaining data in a way that is exceptional to the purposes that are disclosed and the purposes for collecting them.

Now that in and of itself is also a little bit challenging and we'll talk about that first here in a second. And then we have identifiability issues. So, as I was just mentioning, under HIPAA we have the de-identification safe harbor.[2] The de-identification safe harbor is used primarily to reduce risk to individuals. So, if, for example, an organization wants to reuse data, they take a certain number of steps, usually it's removing 18 identifiers or obtaining an expert opinion from somebody who is a statistical expert, to determine there is very, very low risk to an individual person. And that renders the data non-PHI. So, it's no longer Protected Health Information when it has been de-identified. The problem, of course, is the larger data set you have. When you take data sets from a variety of places, say a public record from another organization, say an insurer, and from what you've collected as a medical device manufacturer or as a health care provider, suddenly you have a variety of data that are tremendously useful, but nevertheless may actually be more identifiable. And even removing those 18 identifiers could indeed result in identifiability of the individual. So, there are some interesting challenges related to that.

And then finally data subject rights. Data subject rights with AI are, I think, for the most part still intact. The challenges related to data subject rights, though, are related to that technology model that I was just talking about. When you have a variety of different third parties, you potentially have a variety of different partners, affiliates, or just customers, if you happen to be selling data. It may be very difficult for you to undo what you've already done. We know that data flow pretty easily. And so, for example, if somebody wanted to restrict further processing of their data by revoking their consent, it can be very difficult to get those data back. So, some interesting challenges here.

These are just kind of the primary privacy challenges. But as a backdrop in the medical space, at least modern medical technology today, we have other issues that complicate these. One of them is market concentration. So, for example, there are only, I believe, two manufacturers in the world that manufacture insulin pumps. What that means functionally is that where we expect the market to jump in and for individual customers to sort of choose their options and choose an option that might be better from a privacy perspective, if they desire that, there just aren't that many options. And further, there are, you know, additional challenges because a lot of these devices are prescribed or recommended by physicians. They're not the type of thing that an individual is likely to go out and choose on their own, so they are really depending on

[2] *See* 45 C.F.R. § 164.514.

the expertise of another individual. And there aren't a lot of alternatives. Many of these devices are moving towards this kind of a digital footprint with AI and other types of functional technologies behind the scenes, which ultimately means that an individual would not, number one, might be less likely to choose an analog device because the technical features are so much more superior in a more digital or connected or algorithmic type of an implementation. But additionally, they may not even exist. So, in this movement, individuals who are already reliant on these types of devices for either just the ability to live if we're talking about pacemakers, or for quality of life if we're talking about hearing aids. There is inherent coercive bargaining. What we mean by coercive here is that we have contracts of adhesion that apply. So anytime that somebody is actually signing up for the mobile application that helps their technology run or to kind of keep them in the loop – those are not the types of things that individuals can actually bargain about. You know, there's one form, there's one piece of information. And it's sort of like a supersized coercive bargaining because, again, you have individuals who are dependent on these technologies, either to live or for quality of life.

There's a disproportionate knowledge barrier here. And I don't just mean between the patient and the manufacturer. I mean, that's a pretty big chasm. But often we have disproportionate knowledge between the physician and the manufacturer. A lot of physicians don't actually understand how a lot of these technologies work, but are trying to find the best technology fit for their patient. And so we have pretty much one organization that knows a lot about the technology and what's happening with it, and you have an individual downstream that is really trusting in their doctor and trusting in the manufacturer to ensure that it is going to be a safe and privacy-rich type of functional technology. And I would also argue that when somebody has to choose between their life and their quality of life versus privacy, usually those first things are going to win out. And they are more likely to give up their data for purposes that are beneficial to an organization but maybe less beneficial to themselves.

Alright so data minimization, I wanted to just show you an illustration of what artificial intelligence can look like. And this really illustrates why it is very, very challenging, for example, to adequately inform somebody at the point of a privacy notice. Data is tremendously useful, we know that. Data reuse is tremendously useful. And it can be used in a lot of different products. But that use can continue indefinitely. And, again, we may have data that functionally are de-identified, but actually are tremendously

identifiable that are used. At the same time, data go into each one of these layers and you can see a picture here of the input layer, hidden layers where calculations are happening. For example, if you look here on the right, the skin cancer diagnostic app, some of you may be familiar with this. I was fortunate enough to present with somebody who actually created this app. And they told me that they have 1,000 hidden layers. So, at a thousand points, there are different calculations, different weighing, that happens between those data points. And there might be additional injections of additional data in each one of those layers. You can imagine how difficult it would be to explain to somebody, especially a downstream user, how the calculations are happening or why certain data points are going to useful in a particular calculation. Describing the purpose and use at the point of forming a relationship with that individual is tremendously difficult.

Identifiability. So, we talked a little bit about the need for big data in AI implementations. But additionally, we're dealing with a personalized kind of medicine. So, the entire purpose why we have AI diagnostics and AI technologies is that we believe they will be more effective than the alternative. They will be more personalized, they will be more effective. And so, for that reason, actually, if you want to facilitate personalized medicine, it usually requires more collection of personal information and less de-identification or anonymization of the of the data sets you have. And AI can be used to be used to identify and create inferences and so usually it's very difficult to achieve things like de-identification and anonymization. And as I mentioned before, HIPAA's current de-identification safe harbor is not really a great fit for this kind of a model. So, we're kind of in a difficult position.

Let's make it a little more complicated. I previously wrote a paper on the concept of consent and why it is tremendously difficult to achieve in the healthcare environment in particular.[3] But often, from a legal perspective, we position notice and consent as sort of curative. And I'll that even from a Federal Trade Commission perspective, if somebody files a complaint and they look at the notice and consent and the person consented, and the notice was reasonably informative, we're often in a situation where it's almost a rebuttable presumption that what they did was legal. But there are a lot of problems with the function of notice and consent just functionally and logically.

[3] Charlotte A. Tschider, *The Consent Myth: Improving Choice for Patients of the Future*, 96 Wash. U. L. Rev. 1505 (2018).

First of all, we have again the voluntariness problem with contracts of adhesion and coercive care. When I say coercive care, I don't mean somebody is forced to have healthcare, but rather that they don't really have a choice. When a person is seeking healthcare or seeking use of a technology, they don't really have a lot of better options. The choices really are to live or have some quality of life or not. And because a lot of these technologies are functionally more effective than some of their analog counterparts, your choices are "Do I get the less effective technology?" or "Do I give all of my data away and use this more effective technology?" And it's not really a fair calculus.

Secondly, we have what's called a structural problem. In the structural problem we have privacy policy fatigue, I think most of us are familiar with that. When an individual is forced to go through privacy notice after privacy notice all day long, they can actually stop paying attention. And there was a study that was done that estimated that if somebody read every single privacy policy that they were confronted with, it would take seventy-six full-time days of the year to do it.[4] We know that individuals just simply don't have seventy-six days a year to look at everything and we would presume that they care enough in a healthcare context to look at it. But the reality is that the way that privacy policies or privacy notices have been written historically, is in some ways to kind of provide the formality without tremendous information even being offered.

Then we have a "cognition problem," so privacy as risk. When we're in a situation where privacy is something that somebody has to think about in terms of whether or not they're going to agree, they have to be able to think about it from a risk-to-themselves perspective. But the concept of privacy harms and what kind of challenges a person might face if they give too much data away are highly attenuated and very, very difficult to imagine in a really visceral and specific way. And then we what I call an "exogeneity" or "abstraction" problem. And I was referring to this in the beginning when we look at the technology implementations. It's very hard from the position of a patient to imagine all of the third parties who might be two or three steps back in these technology implementations. And, in fact, when I work with organizations, often they don't really even know what the practices of their third parties are. And they haven't even functionally agreed to appropriate terms from a contracts perspective. So, if you're dealing with, you know, an organization or manufacturer that probably doesn't even

[4] Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4(3) J. of L. and Policy for the Info. Soc'y 543, 543-568 (2008).

know their third parties are doing, how can we expect health care providers or individual patients to do the same? And then finally, we have a temporal problem. When we provide a privacy notice and we offer consent, usually that is based on the purposes that have been specified in the information that's been specified in the notice. But AI actually benefits from more and different information that's presented along the way. And for that reason it becomes tremendously difficult to ensure that somebody knows what data are going to be used for at the time when they consent. It's just almost impossible.

Alright, and I'll skip "Data Subject Rights." I kind of mentioned this, but, it's very difficult to actually get information about your data when it's handled by third parties throughout the process. And I talked a little bit about the market dynamics and coercion piece so I'm going to skip past that.

So, what might this look like functionally? Well, a HIPAA-compatible privacy model would relate to, number one, "minimum necessary" still being in place. Reliability, safety, and efficacy purposes might be justification for keeping data for a period of time that is reasonable. But we refocus it from a legitimate interest perspective, so it's like another kind of lens in which we evaluate "minimum necessary." So "minimum necessary" is not necessarily what is needed right now, but what may be needed overall in the course of the life and the improvement of these AI. De-identification and retention are positioned more explicitly as an ambit of "minimum necessary." So, one of the problems I see with a lot of organizations is that they only use de-identification when they want to do something with the data that probably the patient or the doctor might not like. But they often have almost no retention practices whatsoever, in that data are not securely deleted on a regular basis. Perhaps we can kind of bolster that side of it, while at the time offering a little bit more fluidity in data use and reuse.

And demonstration of reidentification risk might be a way to bolster the de-identification space, so, shift from an 18-identifier a model to a model where we really do focus on expert determination as the basis for de-identification. And then finally, reevaluate this concept of an "information fiduciary." This is something that's actually been raised in the privacy community as a way of refocusing towards organizations that are taking on responsibility in receiving data and creating a fiduciary responsibility to the individuals whose data they have collected. Now, I don't necessarily endorse a broad model like this, but in something like healthcare when we're talking about manufacturers and downstream patients,

this might actually be a really nice model that we could look at a state level if HIPAA is not in a position to expand in any meaningful way over the next few years.

And then finally, and I'll kind of go to this interest balancing because I think it's probably the most interesting part, is that instead of focusing on a notice and consent model, perhaps we instead refocus towards a legitimate interest model. And I know from talking with a lot attorneys that many of you don't like balancing tests. From a court perspective, for example, they can be a little bit challenging, especially in the criminal law space. But perhaps we put the onus on organizations to actually conduct a risk assessment to determine if the benefits to the individuals, whether they're a class of individuals or just individual people, would actually be advanced by processing the data further. You can see some examples I've included here and how you might do legitimate interest balancing. But the overall function is that the organization would have to demonstrate and would have to record and document that, if they're going to use data for additional purposes, that the interests of the individual, the users of these devices or the subjects for diagnostic tools, would actually benefit more with further processing. It's a way to sort of reformulate how we're thinking about the concept of notice and consent. Thank you so much for your patience with my technology issues and I look forward to your questions.

Casey Goggin: Thank you so much. That was a fantastic presentation. It's unfortunate that we're running out of time, I would love to do questions.