Is Our Privacy Law Ready for the Age of AI?

Transcript

Dr Caroline Green

Welcome to Accelerating AI Ethics, the podcast of the Institute for Ethics in AI at the University of Oxford. My name is Dr. Caroline Green, and today I have the great pleasure of having one of my very own colleagues from the Institute, Professor in AI Law and Regulation, Ignacio Cofone with me. Welcome, Ignacio.

Professor Ignacio Cofone

Thank you, Caroline, for having me here. It's a great great pleasure to chat with you about this.

Dr Caroline Green

Yeah, it's really great. And we'll be diving into your recent book, which is called The Privacy Fallacy, Harm and Power in the Information Economy. So we'll be talking about some of the main kind of premises that you describe in this book, but we'll also be talking a little bit about AI law and regulation more widely, if that's okay. But before we start with that, I'd love to get to know you a little bit better. So, tell me a bit more what drove you to work in this particular area of law, AI law and regulation?

Professor Ignacio Cofone

That's a great question. I guess we stumble into ideas that then we end up finding interesting and they end up kind of driving us. I started writing about privacy a while ago, and then I realized that it didn't make any sense to write about privacy or data protection without writing about AI. So, around 2017 or 2018, I started writing about AI, and then I realized that I found AI... broadly quite interesting and lots of the pressure points that AI has for the law relate to the regulation of information. So when we look at how we regulated and we didn't regulate data in the last couple of decades, there are lots of lessons for AI, kind of two ways. First, because data is a good pressure point to regulate AI and prevent some of the AI harms that we face. And second, because we can learn from all the mistakes that we made in regulating data to not make them again in the regulation of AI more broadly.

Dr Caroline Green

And I think, you know, as you say, you are now working on this really important topic of privacy in relation to AI, which is so rapidly evolving. So in your book, The Privacy Fallacy, you show that today's privacy laws are outdated. and in a way even enabling an

ecosystem that produces many kinds of data harms. And so this is really concerning in the age of AI where everything is just developing so fast and where there seem to be, you know, I hear constantly about the opportunities of lots of data and AI using this data for the benefit of us, of course, But I know from my work that people are incredibly concerned about how their data is being used, about the different types of harms. So tell me more a little bit about that. How do you feel about your work right now?

Professor Ignacio Cofone

So one thing that I constantly think about when I write about data and about AI is why do we care about privacy if we care about privacy at all? And I think that relates to the question of what do we mean by privacy? And there are many ways to define privacy. And depending on how we define it, a lot of people would say that they don't care about it. And perhaps they're right that they shouldn't care about some aspects of it, but not others. So oftentimes we think about privacy or data protection as the control over personal information. And a lot of how our rules are set up are focusing on the control of personal information. And I think that leads to many saying, I don't care about my privacy because they might not care about controlling their personal information directly if they are enabling data for the social good or they are getting lots of interesting things in exchange. But if instead of thinking about control, we think about data harms, I think people do care about preventing data harms. People do care about not being discriminated. They care about not having their identity stolen. They care about their democracies working well because people are not manipulated to vote for someone that they wouldn't have wanted to vote. They care about not being harassed and they care about not being exploited. And so if we reorient the discussion away from everyone should control their personal information into how do we have a system that enables technological innovation while protecting us from data harms, I think that is something that people do care about.

Dr Caroline Green

Yeah, I think it's really important how you show that privacy is such a multifaceted concept, value, something that's so incredibly relational. It's about how we interact with each other, how others, even who we don't know, can interact with what's really important to us. And I know you were on a recent podcast with Rebecca Lowe, talking about privacy and what that means. So anyone who's interested in really diving into that, I think that's a brilliant podcast to listen to. But I think, in your book, you have a wealth of stories of where data harms have come about in all its, sort of diversity. Can you tell me about some of the, some of those examples, how people have seriously been harmed? by that state of the digital ecosystem. And tell me some of the examples.

Professor Ignacio Cofone

Yeah, so you mentioned that many of data harms right now are relational. And I think that's a really important point, because the way that the regulation of data has been structured in the last couple of decades is that it has been structured in similar ways to consumer relationships. The assumption is that there's one person who holds data about themselves that has a choice to make about whether to share that data or not with a single data collector or processor, and it's about giving that person power to make that choice about whether to share data or not. But today, as you were saying, most data harms, particularly because of AI, are not so linear. One example is the 23andMe bankruptcy. 23andMe was this company that offered people the service of sending a swab with their DNA in favor of information, in exchange of information about ancestry and propensity to get certain diseases, and lots of people sent that information. And when 23andMe got hacked with that, which then led to a large bankruptcy, part of the problem was that hackers didn't only get access to data from those who send their samples to 23andMe, they also had access to probabilistic information about lots of other people because genetic data by its very nature is shared. So anyone that was related to those who sent samples also had information about them sent. And this is just one example of millions of pieces of information and databases that provide information on more than those who agree to that information being sent. If you were to download, I don't know if you use TikTok, but if you were to download TikTok right now or a different social media app and you start swiping, you would find that after a few minutes, the app is very good at estimating what videos you might or might not like, and that's because... As soon as you download it, even if you don't give it much information about you, won't only have information about your swiping. It will match that information with probabilistic information that it has based on traits, behaviors, trends. All the other information that other users sent it is helpful to inform about you. And that's fine, but that does mean that our laws cannot simply rely on individual choices as mechanisms of protection, because these data are interrelated and the choice of one person also affect others.

Dr Caroline Green

So, I mean, I do use some social media, and I know all of us say that.

Professor Ignacio Cofone

We sometimes speak on social media.

Dr Caroline Green

Exactly. And in your book, you also said something made me laugh. It said, well, terms and conditions, who ever... never reads those. But constantly I'm ticking, oh, consent to this, consent to that. And if I have a bit of time, I'll be like, okay, no, decline this. But, you know, that's all I really do. So do you feel that one of the real issues also is that

people, you know, they know that they're consenting to something, but they don't, you know, they don't actually know what's happening here?

Professor Ignacio Cofone

I think that's right. And I would take it a step further. And I would say that they're not really consenting to that. They're agreeing to that. And that agreement doesn't and really constituted consent. Because it is not just that it would be impossible for a normal human being with the amount of services that we use to participate in social and economic life to read every privacy policy from top to bottom and read again every time the terms and conditions change. It's also that if we read it, we would find information that is at points technically quite difficult for the person who's not trained in data protection law. It doesn't say much about the exact processes, and therefore is not very informative of the risks. A privacy policy, for example, could say we share data with third parties. But sharing data with third parties could mean we have a very secure server that we need to send some data to because we don't have space in our service and they have very high cybersecurity measures. Or it could mean we share data to advertisers, or it could mean we sell your data to data brokers. And those three things are very different. and the risk profiles are very different. So because consent requires having some understanding of the object that one is consenting to, and in this case, object includes potential risks and consequences, I don't think there's real consent in agreement to many of these policy policies, which again, it's not something devastating as long as we don't rely on those agreements as the main mechanism of protection.

Dr Caroline Green

Is that what we're doing though?

Professor Ignacio Cofone

I think that is what we're doing.

Dr Caroline Green

That's the problem, isn't it? So what you're really describing here is because, I mean, the companies who are doing this, they are ticking the box of the law. And by doing so, the law is really enabling this whole ecosystem of how our data is being taken, and how then people also interact with it. Can you tell me more about that ecosystem and the power dynamics at play?

Professor Ignacio Cofone

So yes, and that takes us back to the consumer relationship example, because the system is framed as one in which we operate as consumers. And we as consumers can make choices about whether, for example, We want to give our e-mail to a certain store to get a discount after. But that's not how many, or at least the most significant, data

relationships happen right now. Data is used now to process our data to see whether our CV gets to an interview stage through automated decision making to see whether we get employment. Data is processed to calculate our credit scores to see if we get loans from banks. Data is processed by third parties that get data from those that we give it to, but then package it and sell it to other entities. Data is processed for many daily interactions that we have with each other and that we need to operate in the economy. And that context is quite different from the context of I make a choice about whether to use this particular platform that I really need, that I could use, and then I give this discrete piece of data to. Also, we're used to thinking about data as disclosed data, but most of the data that is valuable and most of the data that is risky is not so much the data that we disclose. but the inferences made from that data. It is not that every like that we put on LinkedIn is particularly risky or particularly valuable, but rather the aggregation of millions of those together with probabilistic trends from data that is provided by others can provide quite sensitive information about people. So the categories that we usually think in kind of break down. We can't think anymore in terms of data of one person, because most of the data is shared, because it's interrelated. And we can think of risksless data and risky data, because even the most riskless data can provide risky inferences. That means that the I agree moments that people have are not particularly helpful. And the strange thing about it is that it might be a system in which everyone loses, because people get tormented with constant requests to agree or not agree to things, including every time they visit a website and they get a cookie banner. They don't get meaningful protection from the risks and harms that can happen from it. Companies then complain that compliance costs are really high because it is truly costly. to create all those consent opportunities. And so new entrants to the market are deterred because the GDPR is hard to comply with. People don't get protected in exchange. And so it is difficult to see why we keep having the system at all.

Dr Caroline Green

Yeah, that's a really interesting point. I'd like to focus a little bit on companies here. What you're describing here just goes to show our data is incredibly valuable. Companies want it. Why? Tell me more about that.

Professor Ignacio Cofone

Why? Well, first of all, it's valuable because it can create lots of profit opportunities. And second of all, it's valuable because of what you were saying earlier, that it creates power. Large tech companies are some of the most powerful companies of the world, not only because of the money that they make, but because of the data that they hold about everyone. And that creates potential opportunities for manipulation, that creates potential opportunities for allocating resources, that creates potential opportunities for shaping public discourse in one way or the other. Tech companies now have actual airports, have had incidents in the definition of borders, have satellites, have eyes into

people's homes. They do things that no company all the way back to the Dutch East India company could do. And that gives a pretty solid indication that the power dynamics between those who hold the data and those whose data is held is different from those power dynamics than the ones in which we choose whether to buy something at a store.

Dr Caroline Green

Yeah, and you're talking, also in your book, you talk about the information exploitation that we're really victims of. And often talking to companies or visionaries in AI specifically who are saying, we're building these systems and these companies to benefit you?

Professor Ignacio Cofone

Right.

Dr Caroline Green

And surveillance, yes, I mean, I don't care about surveillance. Actually, it will keep me safe.

Professor Ignacio Cofone

Right.

Dr Caroline Green

And I'd like to know, I want people to know where I am because then if something happens to me, they can come and help me. Or I want health or companies to know when there's something wrong with my body before I would otherwise, so that that we can prevent conditions or whatever. So here I think that goes back to that point of why should we care about our privacy and our data so much when obviously it benefits us?

Professor Ignacio Cofone

Yeah, that's a really good point. And that shows both all the beneficial things That we can get from technology that are desirable, and we should get them, and it shows why we need to frame privacy in a way that is not just control over information, but is preventing exploitation, having... an app that allows me to show my partner where I am in case they need to know where I am is very valuable. Now, if where I am is then shared with third parties that can find me and stalk me and harass me, then that's not so good. Knowing what's going on in my body to prevent certain diseases is very helpful. Now, if that gets given to third parties who might use it to game insurance premiums, and not give healthcare to those who need it the most, then that might be not so socially desirable. So I think we need to spend less time and attention worrying about types of data collection and more time and attention worrying about different data uses.

Because the same data can be used for socially and individually beneficial things or for socially and individually detrimental things. And we should enable the beneficial treatments while preventing the exploitative ones.

Dr Caroline Green

Yeah, so it's that argument of, well, if I've got nothing to hide.

Professor Ignacio Cofone

Right.

Dr Caroline Green

You know, it doesn't work nowadays, doesn't it?

Professor Ignacio Cofone

Doesn't work because some people that have nothing to hide have a lot to lose. And sometimes we confuse the negative external consequences that we might face with the thing itself that we might or might not try to hide. Imagine, for example, a colleague that might have an invisible disability. And our colleague might be great at their job, they do their job without any problem, they have all the accommodations that they need, they can do their job perfectly. But imagine that their employer has a bias against people with invisible disabilities. And then imagine that the colleague gets outed by a malicious colleague that then tells the employer, oh, this person has this particular disability, and then they get fired. We would think about that person having been wronged in two different ways, having been dismissed unfairly when they could perform at their job, and having been outed by their colleague who didn't have the entitlement to out them and to reveal their personal information. Not because they had an entitlement to control all their personal information, but because their identity shouldn't be exploited for fun or to enable bias. And those two wrongs are different because one of them may be the material consequence that they suffered, but the other one, I think, is the privacy wrong that they suffered. Someone who used their information against them for fun, for feeding a bias or for something else. And if we repaired one of them, if we repaired the dismissal, we would still have not repaired the privacy wrong in and of itself. That's what I think is the privacy fallacy. The privacy fallacy is the idea that privacy is very valuable and we all care about it because we all care about exploitation, but then we end up only focusing on those consequential, tangible things are not on the wrongs of exploitation by using people's information against them, when from the idea that some version of privacy is valuable, follows the idea that someone can create wrongs to that value.

Dr Caroline Green

And that's what the legal framework currently is also based on, that privacy fallacy, right?

Professor Ignacio Cofone

I think it depends, because it depends on how, it depends on how it is enforced, and it depends on how it is proposed. I think the legal framework has lots of problems in that it relies too much on individual control over their personal information. And individual control is only a proxy for preventing exploitation and negative consequences. That sometimes works and sometimes doesn't work. I think saying that someone has nothing to hide would be operating under their privacy fallacy because it is implying that Privacy only exists to hide undesirable things and not to help people not be exploited or no loose things like being dismissed unfairly. But a lot of the ways in which the legal framework is applied does follow, I think, this faulty logic. And we do have places in the legal framework to grab to be able to move past it.

Dr Caroline Green

So just to clarify what we mean when we talk about the legal framework, so what kind of flaws and regulations are you referring to?

Professor Ignacio Cofone

So mostly right now I'm thinking of the GDPR and the UK GDPR and privacy and data protection laws that are similar in other countries.

Dr Caroline Green

And so, you know, taking you up what you just said, so there are actually connection points here to improve and to change the system to be more in line with actually protecting people from these harms and addressing them. What would you want to see?

Professor Ignacio Cofone

I would love to see less emphasis on modes of data collection. Less emphasis on whether people agree to terms and conditions, less emphasis to whether people agreed on cookies, and more emphasis on the principles that are also in these legal frameworks that are underutilized and could be used to prevent exploitative data uses while enabling beneficial data uses. like the principle of fairness or the principle of data protection by design, which we could use to say if someone collected certain data, it doesn't matter that much ultimately whether they went through the mechanics of people mechanically agreeing without reading the privacy policy, but it does matter that they don't design a system to exploit people, or they don't design a system to use people's information against them. Part of this is using the information that one needs and not using high-risk information otherwise. And I think not doing that and not paying attention to the reason for which we might care about how that information is processed sometimes leads us to false dichotomies. For example, During the

pandemic, when different governments and companies were developing contact tracing apps that could have helped us be better at reducing how many people got COVID, this was often presented as a trade-off between privacy and public health. Well, how about maybe it wasn't a trade-off between privacy and public health? How about if an app was designed in a way that was not exploiting people, was not misusing their personal information, and was only gathering the information that was required to contain the pandemic, then there was no privacy harm. And then there was just a public health metric.

Dr Caroline Green

Yes. So what you're explaining, I think what I hear too is, so we have the legal mechanisms that we could change, adjust, do more of, but there also needs to be a change in the way that companies are operating and that they take more responsibility. for the things, the values that we care about, right?

Professor Ignacio Cofone

I think so. And lots of companies do do it. is very difficult to know to which extent each company does it. I mean, for some, we might have intuitions one way or the other. It's difficult to know the internal workings. But to a large extent, companies respond to the incentives that the legal system gives them. There's this trope in tech of move fast and break things. And move fast and break things was initially an idea from Mark Zuckerberg when he did the Facebook IPO that then Facebook quickly abandoned. But other members of the tech industry do embrace. And the interesting thing about the move fast and break things is that it's used in two completely different ways. is sometimes used by sectors of the tech industry to say, oh, look, you do have to move fast and break things. If you're not moving fast and breaking things, you're not being innovative enough. While other techs of the tech industry don't necessarily believe that. And it's also used in policy circles to say things like, oh, look at the tech industry, they're breaking so many things. They don't care about the things that they break. But It might be more effective to match the moral scandal that many policy circles feel about move fast and break things to reasons to not move fast and break things. If companies were asked to pay for the things that they break, they would be more mindful about whether they move fast and break things or whether they move fast a bit more carefully and try to break through things in the process.

Dr Caroline Green

And I guess that's also a point of investors, right? Because companies are often under so much pressure to find some, raise funding and then scale really quickly. So I think that kind of like ethical investorship is also really, really important, right? To say it's a part of the ecosystem in a way.

Professor Ignacio Cofone

I think it is. And investing often responds to the profit potential that different firms may have. I think a bad legal system would make it very profitable to engage in exploitation and that might drive certain amount of investment to firms that may engage in exploitation. A better legal system is one in which exploitation is not encouraged because it becomes expensive to do it. And at the same time, engaging in ethical practices is profitable and is less expensive to do it. Many of the problems that current regulations have in an environment where companies say that compliance costs are very high is that some level of compliance costs are relatively fixed no matter what you do. There are some things that scale in terms of compliance costs, but there's some things that don't so much. You might need to have a data protection officer, whether you're a huge company or you're a small company, whether you have a safe data practice or you don't have a safe data practice. You might have to file lots of paperwork on privacy and data protection impact assessments, whether you have very safe data practices, whether you have very unsafe data practices. A better system is one in which we make it cheaper for companies that have safe and ethical data practices, and we make it more expensive for companies that have unsafe and unethical data practices. And that might create a better environment for investors who do want to invest ethically and want to invest in ethical use of data to actually get a reward for doing that.

Dr Caroline Green

I just want to focus one more time on that, on harm liability. And just to give it a bit more of a, what could it look like in practice? Do you have kind of like a practical example of harm liability, what it could look like?

Professor Ignacio Cofone

I think it could be as simple as a provision in a data protection statute that says, If a company exploits people by privately benefiting from using their information against them or by privately benefiting from exposing them to known harms, they are liable for the potential harms that those people end up facing. That's relatively short, but would make a dramatic change because it would mean that Even if someone didn't breach a very specific provision like acquiring agreement for the collection of data, they would be accountable for the way that they use data. We would put a lot more pressure in different uses of data and on ethical use of data. And if we do that, we could also put less pressure on data collection and on other aspects of enforcement that do have fixed costs that apply the same for everyone. And putting pressure on the users is the maximum variability of making it easier for those that have ethical data uses and making it more difficult for those who don't have ethical data uses.

Dr Caroline Green

Yeah, I see. So I think that also is part of an answer to my next question, which is what would you tell policymakers, politicians, lawmakers, and people. Right now, what could we do tomorrow to make this situation, to improve the situation that we're in?

Professor Ignacio Cofone

I would tell policymakers and legislators, reform data protection law, we could get away with the problems that it had before AI. But in the AI-driven information economy, this is a system that might actually slow down innovation and doesn't protect people from harms. So 2 sides lose. Reform it so that we put less emphasis on the individual control over personal information, and we put more emphasis on the prevention of data harms, including exploitation, when we try to avoid companies from using people's information against them. What to tell people is harder because in a context where people are, to a large extent, powerless, there's not actually much that we can do because we can take certain safety measures. We can try not to share too much. We can try to not share sensitive information. But I worry that if we put too much effort in telling people what they should and should be doing, the risk profile won't change that much because even if people don't share sensitive information, sensitive inferences can be made about them. And I worry a little bit about in the long term blaming the victim. I think of all those cases of non-consensual disclosure of intimate material where you sometimes have An actor saying, Well, should this person really have shared that with a second party? And the reality is that we deserve to live in a world where we engage with others without being worried that others will exploit us. that be a partner with who someone shares an image or a company that invites sharing data to offer a product. We should be able to engage with others and trust that they will not exploit us. And that is a better system than one in which we tell people, don't engage with others, be careful, be aware, be worried, and then bad things could happen to them anyway because of all the things that happen in the background. So if I were to suggest to the average person who doesn't know what to do with their data to do one thing, I would say the best thing that you can do is call your legislator and tell them you need to reform data protection law because that's the biggest payoff that we can get.

Dr Caroline Green

Yeah, I think it boils down to that question of what's the society we want to live in. And we want to be in communities and in relationships. with people where we can trust each other, where we can also sometimes make mistakes, I guess. we're just human without having to worry that it's just going to be all out there. And that we're not exploited, but that we are, that our data isn't used against us and manipulated. And it seems like a lot of people are so worried about that. And so it'd be nice if we could not just half the vision, but move towards that. So I think your book is a really important contribution towards a future like that. So before we close, I just also, I mean, this is a big can of worms, really, because we've got you here as an expert on this, you know, AI

law and regulation. It's one of the big topics out there, when it comes also to AI ethics, as an institute, it's one of our core research areas. And there seems to be this sort of like tension between innovation and law and regulation. So we want low regulation and sort of loose laws and so on so that people can innovate. Where do you see is all of this moving towards? Because we've seen some, you know, big changes, the EU AI Act, for example. How do you currently, from sort of like a macro perspective, how do you see the whole field developing AI law and regulation?

Professor Ignacio Cofone

I think the problem is often presented wrongly because the way that the problem is presented is often regulation versus innovation. allowing for new things or making things safer. But the reality is that regulation and innovation are not opposites. It depends on the type of regulation, and it depends on the type of law. There are lots of regulations that enable innovation. Intellectual property is a form of regulation, and patents enable innovation in all sorts of ways. Contract law is a form of regulation, and it enables innovation in all sorts of ways. Property law is a form of regulation, and it enables innovation. So when people say innovation versus regulation, they're thinking about bad types of regulation that prevent innovation. But it is worth thinking a bit more broadly about what are the legal and regulatory rules that we can design to allow for innovation while keeping people safe. And that is often possible. So when we're thinking about scaling back the AI Act or changing the AI Act, the discourse that we see on newspapers is often, oh, Europe realized that it regulated too much, and now it is scaling back. But the better lesson might be, Europe realized that the type of regulation that it put forward was not the best one, and now it needs to change it. because otherwise we end up with the types of innovation that we may not want, the type of innovations that are exploitative and harmful. What we should do through regulations such as intellectual property, property law, contract law, and better data protection law that prevents exploitation instead of focusing on control is trying to enable and encourage the type of innovation that is socially beneficial.

Dr Caroline Green

And I think, that goes back to what our whole conversation, right? And it's also not necessarily about reinventing the entire wheel. You already, there's already very good, there are already very good frameworks out there. It's understanding where are the weak points, where do we really truly need to innovate in law and regulation and policy and so on and otherwise do more of what's already there. So it's not just a black and white question or, you know, bad or good. Yeah, that's right.

Professor Ignacio Cofone

And it requires connecting with what is it that we're trying to protect. And that is, I think, sometimes why these things get put into false trade-off. Because it is relatively easy to

come up with a checklist of measures that companies need to jump through in order to collect data. It is more difficult to have a serious debate about Why are we regulating data in the 1st place? What is it that we're trying to prevent? What are the legal values involved? And how do we protect those while enabling all the type of innovation that we actually want?

Dr Caroline Green

And I think for that, we also need these open conversations where we do actually talk about these issues of like power dynamics, how are those currently being played out? also developing that understanding for each other and how people are being exploited, how they do perceive harms. So having these spaces where, tech companies, civil society, academia, and so on can come together, legislators, to inform each other, I think that's really valuable and important.

Professor Ignacio Cofone

Yeah.

Dr Caroline Green

So, yeah. Well, thank you so much, Ignacio. Is there something else that you would like to...

Professor Ignacio Cofone

Nothing else, really. Thank you for having me. This was a fun thing to talk about with you.

Dr Caroline Green

Well, thank you so much for coming. And yeah, it's great to have had you at this podcast. I would highly recommend your book. And yeah, well, thank you again.

Professor Ignacio Cofone

Thank you.

Dr Caroline Green

You've been listening to the Accelerating AI Ethics, a podcast from Oxford's Institute for Ethics in AI. If you enjoyed this episode, please subscribe and share. Until next time, thanks for listening.