



The Battle for Your Brain

*A LEGAL SCHOLAR'S ARGUMENT
FOR PROTECTING BRAIN DATA
AND COGNITIVE LIBERTY*

NITA A. FARAHANY IN CONVERSATION WITH PAUL W. GRIMM

Mindreading may sound like the stuff of science fiction, but these days, as they say, truth is stranger than fiction. Employers track employee attention and even moods. Technology users can type a text using only their minds. Neural wearables, using biofeedback, give migraine-sufferers pain relief in real time. Neurotechnology may be the next great tech frontier, but how brain data is accessed — and who has access to it — will be the next great legal battleground, with implications for privacy, freedom of thought, and self-determination.

In *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology* (St. Martin's Press, 2023), Duke Law Professor **NITA A. FARAHANY** provides an in-depth look into the future neurotechnology can provide, outlines the stakes if brain data is left unprotected, and provides a plan of attack amid the battle for our brains. She stresses the need to establish an international right to what she terms “cognitive liberty” and to define a new set of norms that will protect brain data as neurotechnology continues to improve.

Judge **PAUL W. GRIMM** (ret.), the David F. Levi Professor of the Practice of Law and director of the Bolch Judicial Institute (which publishes *Judicature*), talked with Farahany about her book, cognitive liberty, and the rise of neurotechnology. The following is an edited transcription of their conversation.

What people don't realize is that the space they have long thought of as private — what they're thinking and feeling inside their brains, the things not expressed through words or actions — is now up for grabs and is already being accessed by a lot of different companies.

GRIMM: Can you explain what neurotechnology is, describe its risks and benefits, and tell us about the “battle” for our brains that is already being waged? What is the current technology capable of, and what are the stakes?

FARAHANY: People are already familiar with the fact that if they're using digital technologies — which everybody is — their data is being collected. We are the product, and these so-called free services are really about collecting personal information from individuals — mostly to try to understand how we think and feel.

What people don't realize is that the space they have long thought of as private — what they're thinking and feeling inside their brains, the things not expressed through words or actions — is now up for grabs and is already being accessed by a lot of dif-

ferent companies. Neurotechnology is the use of any sensor to try to directly interpret information from our brains.

People are increasingly comfortable with quantifying themselves. We wear Apple Watches that track our heart rate or Oura Rings that track our temperature or sleep patterns. Increasingly, brain sensors are embedded in headbands, hard hats, or baseball caps. Coming soon is the integration of those brain sensors into earbuds, headphones, and watches.

What these pick up are different measurements of our brain. The most common is electroencephalography or EEG — brainwave activity. For anything that you think or do, your brain is firing neurons, and each neuron fired releases a tiny electrical discharge. Hundreds and thousands of neurons are firing at the same time, giving off electrical discharges that can be picked up by EEG sensors. ►

Now, with increased sophistication of AI, those brainwave patterns or electrical discharges can be interpreted to associate them with what you're thinking, feeling, seeing, or hearing.

So neurotechnology is the use of either implanted or wearable brain sensors that pick up brain activity. EEG is the most common because there are different measurements of brain activity, whether it's how light is reflected or the consumption of oxygen. Each of the different kinds of technologies measures something different, but all are trying to pick up brain activity and then use powerful algorithms to interpret that information. What's surprising to people is the many contexts in which neurotechnology has already been integrated into our everyday lives. It's being used in workplaces worldwide as employers track fatigue levels, attention, boredom, or employee engagement. And it's used by governments to do things like interrogate criminal suspects.

But it's not happening at a wide scale across society yet. The big change coming is that all of the major technology companies are investing in and integrating brain sensors into everyday technology to not just pick up what's happening in your brain, but also to become the way that you interact with other technology — what we call neural interface technology. That means, instead of a mouse or a keyboard, you'll think left and right, up and down, or swipe. And brain sensors will pick up on your intention to move or type and will interface directly with your technology. When it becomes the primary way that we interface, all of our brain data will be up for grabs. And it cannot only be interpreted to precisely understand what we're doing but also to try to change our behaviors.

All this will ultimately be in the hands of the same tech companies that have used the rest of our personal data for years.

GRIMM: Is the current neurotechnology refined enough to reliably distinguish between different states of mind, like doing a math problem versus relaxing? Or is that level of knowledge still a long way off?

FARAHANY: There are two pieces to this technology. One is how reliably it can detect brain activity. The second is how well it can associate brain activity with information and decode it. I distinguish between the two because one of the things that's traditionally been problematic with EEG is that it's a very noisy signal. If you nod your head, your muscles twitch, or your eyes blink, those can register on these same EEG sensors in ways that can interfere with the signal.

That's where the algorithms have gotten much better. Instead of treating that as noise, they treat it as information that goes toward interpretation. You have a much more powerful algorithm if you can start to say this is a muscle twitch, this is an eye blink, and this is EEG activity, and then use all of that in decoding to say what the person is thinking, feeling, or doing. There are still problems, but the ability to filter out noise or to use noise to train the algorithm has gotten much better.

GRIMM: And there is an ever-increasing population to which this technology is being applied.

FARAHANY: As this goes to scale across society and is associated with everyday activity, the refinements will

get more and more powerful. I'd say, right now, it's very reliable at detecting brain states. It's very good at fatigue or relaxation or basic emotional or physical states. It's less good at picking up consumer-based EEG and interpreting that into words or images that you're thinking. I expect that gap to close pretty quickly.

Since I published the book, a major study and several others were published using generative AI for both encoding and decoding. The power to be able to decode literally whole sentences or paragraphs of what a person is thinking or imagining is startling. In the book, I was careful to say this isn't mind reading, but I don't think that's true anymore. I think we're now closing the gap to actually be able to do far more than brain states and to do far more that is closer to what we would traditionally think about as mind reading and decoding.

GRIMM: The explosion of generative AI has gone such a long way toward closing the gap between science fiction and nonfiction. GPT-4, for example, is leagues beyond GPT-3, with more ability to contextualize information more broadly.

FARAHANY: That's right, and it gets very good and customized to the individual. A very early version of this, predictive algorithms, was used with Stephen Hawking. It learned from everything he'd ever written. If he's thinking "black," he's most likely thinking "black hole." Or maybe if I'm thinking black, I'm thinking about the TV show *Black Mirror*. When you have generative AI that is customized and learning of an individual, it becomes much faster and much more efficient to being able to predict — Nita is thinking

You have a much more powerful algorithm if you can start to say this is a muscle twitch, this is an eye blink, and this is EEG activity, and then use all of that in decoding to say what the person is thinking, feeling, or doing. There are still problems, but the ability to filter out noise or to use noise to train the algorithm has gotten much better.

black, she's thinking *Black Mirror*. And the recent study I was referring to — that was very powerful and startlingly accurate — used GPT-1. Just think about what that means now that we're at GPT-4 and leaps and bounds forward in accurately decoding language.

GRIMM: You approach research with a scientist's, ethicist's, philosopher's, and legal scholar's background, which I think allows you to recognize ethical issues and concepts that others don't see. You bring these perspectives together with the concept of "cognitive liberty," which you argue should be a fundamental human right with an accompanying set of norms. How did you come up with this concept of cognitive liberty? How do you define this right, and what would its coexisting norms look like?

FARAHANY: I first proposed cognitive liberty in 2012. I wrote a pair of law review articles looking at the emergence of neuroscience and whether criminal defendants would be safeguarded against its use either in terms of the Fourth Amendment as an unreasonable search and seizure or the Fifth Amendment's privilege against self-incrimination. And I found that, based on existing doctrine, it's unlikely that either of those amendments would provide adequate protection, because

we'd never imagined this world in which you could access the mental states of an individual physically and without their willing or unwilling participation — that you could bypass the conscious person to get at what they're thinking and feeling.

It took a long time for me to figure out all of the contours of cognitive liberty. This book really brought all of that together, which is to say: What exactly is cognitive liberty? How would it work? Where do we need to place it to begin with in order to operationalize it in society?

When I started to write the book, it was with this goal of laying out the framework of cognitive liberty and using neurotechnology as the ultimate threat to it. Many threats to cognitive liberty in the digital age go well beyond neurotechnology. Many of our interactions in the digital age have presumed that we have this internal space that no longer exists as a private space.

Cognitive liberty, as I define it, is a fundamental human right both to access and change our brains and to be safeguarded against others doing so. The basic definition of cognitive liberty is the right to self-determination over our brain and mental experiences: the right to access information and the right to change our brains if we choose — whether that's to enhance them or diminish them, to have a glass of wine, or decide to go to law school. In the same way that you have access

to heart-rate information or genomic information, none of us today really has robust and real access to information about our brains. It's all based on how we access them through our internal software. But how do I really react to things? Am I actually stressed or tired? How do I have transparency into my own brain?

The flip side is that the space of private rumination, of being able to cultivate your own personal identity, the space that protects you from being manipulated directly or indirectly by others, is covered by mental privacy and freedom of thought. Privacy is already an international human right, but we've never explicitly recognized that mental privacy is included within that. Freedom of thought is already a recognized international human right, but it's primarily been applied to freedom of religion and belief. And I show how, given the modern threats to both privacy and freedom of thought, we need to expand our interpretations of both those existing human rights to give us the full spectrum of protection that we need in the digital age. So it's self-determination as an individual right rather than just a collective or political right, and mental privacy to protect against the full kinds of access to information in our brains in the digital age, which is a relative right.

It's a balance between individual interests and societal interests and freedom of thought, which is an abso- ►

lute human right, but much narrower than the other rights because it's only about literally what you're thinking or seeing in your mind, and protecting against the interception, manipulation, or punishment of that.

GRIMM: What is the roadmap for achieving cognitive liberty within the U.S.? Is the effort dependent upon international organizations, or can it move forward here on a national level?

FARAHANY: Since the book came out, I have really been explaining how cognitive liberty is a framework that exists on many levels. It's an international human right, and I started there because I really believe that you need a liability regime and not just a set of incentives to help realign society toward technology being for human interest — and not just for technological or tech company interest.

We also need national legislation in many different parts and subparts, which I'll come back to in a moment. We need commercial redesign of products that align with cognitive liberty instead of with extraction of information from people and instrumentation of people to think reflexively rather than critically when they're on platforms and devices. And we need public-private partnerships to incentivize greater transparency and accountability with tech companies and tech platforms that realign us for cognitive liberty.

GRIMM: So, carrots and sticks.

FARAHANY: We need a bunch of carrots and sticks. Nationally, movement is happening that aligns well around cognitive liberty. On both the international and national fronts, a lot of siloed things have been happening

but not a comprehensive framework about what we're really trying to accomplish. One of the things that the concept of cognitive liberty has done is to help name and frame the issue, the problem, and the needed solutions. For example, if you look at the voluntary agreements just made by the Biden administration with generative AI companies, many align well with cognitive-liberty principles — things like explicitly putting into place safeguards against disinformation and misinformation or labeling of generative AI. A big risk of working or interacting with AI products is when you don't know something is an AI product, which increases the risk of mental manipulation. You want to safeguard against manipulation because you have a right to cognitive liberty and it's fundamental to human flourishing and to how humans interact with technology.

Recently, the Digital Services Act in the European Union included a set of rights for individuals that led to both TikTok and Meta enabling people to opt out of recommender algorithms that put people in echo chambers, silo them, and cognitively shape them to choose just what's popular in their region. That opt-out feature aligns with cognitive liberty because it empowers people to make choices when they're on technology and be able to interact with technology in ways that put humans first rather than tech company interests first. In the U.S., a number of design codes have been proposed, including the California Age-Appropriate Design Code, that are trying to achieve many parts of cognitive liberty. What I have found encouraging in the U.S. is that, whether people are on the right or left, they agree with the need for cognitive liberty.

The more we focus on the common-

ality that cognitive liberty gives across the political divide, the more I think we can achieve.

GRIMM: You wrote in a recent *Wired* article about how tech companies ought to adopt particular design principles to help safeguard against mental manipulation. How do you get companies to integrate components of cognitive liberty into a model that's built around making money on the aggregation and dissemination of brain data?

FARAHANY: It's hard. It's a lot easier to see how a company like Apple — whose revenue doesn't depend on aggregating personal data and selling our attention for advertisements — can incorporate cognitive liberty into its product design than it is for a company like Meta, where the whole purpose is to both understand precisely how people react to data and then sell that information to advertisers. Where a company's commercial interest is tightly aligned with the destruction of cognitive liberty rather than its enablement, there have to be, I think, tax incentives and other kinds of government investment to steer them in that way. This isn't unprecedented. We see this in the need to invest in alternative energy resources and strategies, where traditional companies were built on more extractive technologies.

Incentives can be designed to help companies reorient toward something that's more beneficial for humanity. I think that's what we need. You can't expect companies to just wake up and say, "OK, this is how we are delivering value to our shareholders and we're just going to abandon that to focus instead on cognitive liberty."

Freedom of thought is already a recognized international human right, but it's primarily been applied to freedom of religion and belief. And I show how, given the modern threats to both privacy and freedom of thought, we need to expand our interpretations of both those existing human rights to give us the full spectrum of protection that we need in the digital age.

We have to create incentives in society that help companies shift toward something that's more sustainable for human flourishing.

GRIMM: Everyone is affected by this new reality, no matter what part of the political spectrum you're on. You can't avoid using tech in your own life. So when it comes to cognitive liberty, we might have more cause for optimism than in other large-scale threats, like climate change, which tend to be very polarized. We can perhaps already see this affecting us — and our kids may have the ability to bring consensus.

FARAHANY: I think that's right. We've achieved far less consensus across society either with respect to the science or the existential risk that climate poses. Nevertheless, even in that area, incentives have led to significant shifts in markets and corporate actions.

That gives me optimism because where you don't have consensus, incentives still can be a powerful motivator. I think everyone understands and no one denies the impact of technology. Any parent can look at their child and see the way in which they quickly become addicted to technologies, how technology affects their behavior in a way that makes them

worse off rather than better off in most instances.

And with that kind of consensus, I think it's much easier to recognize the existential threat to humanity. Fundamentally, in order for humans to flourish, we need our minds. We need our ability to expand and to continue to improve our mental states, not diminish them or take our attention and put us into automatic reactionary mode.

One area that's been encouraging as well is that a lot of people are really worried about generative AI replacing humanity. The best way we can safeguard against that is to cultivate cognitive liberty in humans to further develop the ability to have better interoception, to better improve mental agility around critical thinking and resilience, and to shore up our empathy and relational intelligence. Those are the fundamental building blocks of cognitive liberty. They are the best safeguard we have to enable humans to actually flourish in the era of generative AI rather than be replaced by generative AI.

If platforms are diminishing and undermining cognitive liberty, I think it's helpful for people to understand there is a counterweight. There is a way forward, and that way forward demands of us that we really start to invest in cognitive liberty rather than allow it to continue to be eroded.

GRIMM: *Dobbs* is a recent exam-

ple of a recurring theme we see in the courts — decisions about how to delineate fundamental rights in a Constitution drafted at a time when no one could have imagined today's current technology. Do you see *Dobbs*, and other discussions of fundamental liberties, affecting the establishment of a right to cognitive liberty?

FARAHANY: I think so. People have said to me, "How can you be arguing incremental privacy in an era where privacy rights were just eroded under the *Dobbs* opinion?" But I think the *Dobbs* opinion is about a conflict of rights, not necessarily about the lack of rights. It's a belief that there is a conflict between two human beings' rights, as opposed to cognitive liberty, where there is no such conflict. We're not talking about a tradeoff between an individual and a fetus. We're talking about a tradeoff between tech companies, extractive processes, and human flourishing.

My right to cognitive liberty doesn't trade off with yours. It is not a zero-sum game. And we are better off if all of us have an expansion of our cognitive liberty. If you look back at the pandemic and how people fought against wearing masks because they saw it as a tradeoff between individual liberties versus collective and group interest, you don't have that with ►

You need both carrots and sticks to really realign human flourishing with corporate practice. You have to start by flipping the narrative, flipping the terms of service, and flipping where the rights lie. You start with a powerful set of both legal and global norms, and then you layer in incentives that help shift people.

cognitive liberty either.

What you have is a right that everybody can enjoy and that everybody is better off if everyone is enjoying. So both individual interests and group and societal interests are met by cultivating cognitive liberty.

GRIMM: Economists might ask whether my enjoyment of a right prevents you from enjoying a right. If you look at the issues that came up in *Dobbs*, as you pointed out, for me to win, you would lose. Cognitive liberty is different because this is not a zero-sum game. And maybe finding consensus here will help us by showing how we can learn from conflicts between the corporate world and individuals who are affected by corporate decision-making.

FARAHANY: To me, it seems like the best example really is climate and to see that many of the corporations engaged in extractive practices are doing exactly what economically rational actors would do, which is to maximize shareholder value. I don't think that makes the companies evil. What it shows us is that it's not enough to have carrots. You need both carrots and sticks to really realign human flourishing with corporate practice. You have to start by flipping the narrative, flipping the terms of service, and flipping where

the rights lie.

You start with a powerful set of both legal and global norms, and then you layer in incentives that help shift people. Look at the political and economic obstacles from eras where practices have emerged that ultimately have been at odds with human values — and the set of interventions that have been necessary to try to steer norms in the right direction. You see it's never a straight path, but it's always a set of liability, incentives, and powerful norms that need to emerge. You have to grow the demand side. People have to understand that cognitive liberty is being eroded and that there is an alternative path. There are ways that they can claim it — and it has to be easy for them to do so.

The burden can't be on individuals to do so. There has to be a framework and a comprehensive solution across society to ultimately lead us to the next phase of our interaction and integration with technology, which is the phase of human flourishing rather than the phase of human diminishment.

GRIMM: We've talked a lot about how we might be able to use incentives — the carrots — to protect cognitive liberty. In instances when the carrots are not successful, what are the sticks?

FARAHANY: The sticks include a very

strong right to mental privacy for individuals. Most terms of service favor the tech company and disfavor the individual. In fact, most tech companies that have launched neurotechnology products say, "The use of our product means that you agree to it."

A right to mental privacy would say, "You can't do that." The terms of service should be that individuals have a right to their brain data and they have a right to keep that brain data private. And tech companies must get explicit consent for each and every use case.

Interestingly, the first case on this has just come out. The Chilean Supreme Court adopted a set of rights in its constitution that aligns with cognitive liberty. One of the people who helped to bring that forward, a former senator in Chile, bought a consumer neurotechnology device and then used that to gain standing to argue that the company was storing and aggregating data in ways that violated the basic rights to mental privacy. The Chilean Supreme Court agreed that the company had failed to get the necessary certifications, had not conducted the needed impact assessment on his mental privacy, and that it needed to have received his explicit consent to use any of his brain data for scientific or other purposes.

GRIMM: When it comes to allowing individuals to enforce their rights, standing becomes relevant. And in the area of intangible

rights, we have not done a good job of figuring out damages and valuing intangibles. What are the various doctrines and issues that lawyers ought to be aware of moving forward?

FARAHANY: One of the challenges is going to be like the tobacco cases. A lot of the evidence was inside the companies, and it was difficult to access the information necessary to bring forward a lot of cases. In the U.S. and most other countries, we don't have the transparency needed to fully understand tech company practices and be able to build cases around cognitive liberty. For example, much has been written about how TikTok recommender algorithms are powerfully leading to the cognitive shaping of individuals. But to truly test that, we need full access to what recommender algorithms are being used and what data is being used to train them.

GRIMM: How does the defense of trade secrets come in?

FARAHANY: That will be a problem and will create a long delay in the ability to trace the effects of all these different technologies. Then it's going to be hard to figure out exactly where the line is, because even if we get access to all of the data, advertisements have forever been designed to influence us. We have to define impermissible forms of manipulation and cognitive shaping, and ways to develop evidence of how a filter used on a social media platform affects the mental health of youth differentially than a magazine picture that has been airbrushed, which has been permitted for years. Each leads to a kind of cognitive shaping and forms of manipulation.

We need transparent access to

tech information to develop the evidence case. Then we're going to have a hard time drawing the normative lines about when we have crossed the line to violate cognitive liberty versus when this is a permissible influence that we treat as part of our ordinary interactions of being human. And you can show a difference in kind, but you can't prove it until you have robust data that really allows you to understand exactly what's being done, why it's being done, and what its effects on human behavior are. The companies have a lot of access to that information, but they're not going to reveal or release any of it.

GRIMM: I'm sure you're familiar with the *Loomis* case out of Wisconsin that dealt with an algorithm designed to evaluate the chances of recidivism. Originally, this tool was meant to be used so that the court could try to minimize the likelihood that someone would recidivate while awaiting trial. But it was later used to determine how to sentence someone post-conviction. So the uses of these tools can change in important ways.

FARAHANY: There's always a double edge. It starts with positive use, with companies saying, "How do we best pair advertisers to what people are interested in?" But then it becomes "how do we just create demand and change what people are interested in?" And those are subtle mission creeps over time, but those subtle mission creeps are really ways that have been eroding human cognitive liberty.

GRIMM: On a day-to-day level, we might be inclined to say, "This is overwhelming. I don't have the

choice, I have to use this app. I have to agree to the terms of use policy. I'm caught in the tide." What should we, as individuals, be thinking about? What can we do on a micro level to protect our own neuro privacy interests?

FARAHANY: First, it is making smart choices about which tech platforms to use. Some companies have committed themselves to greater privacy protections and empowerment of individuals. I'm encouraged by Apple's commitment to privacy and its decision to enable people to have much greater access to information like their screentime usage or to turn off notifications or tracking.

That's simple, and I encourage people to do it. If you are on a device that enables you to turn off tracking, you should do so. Even if you're not on an Apple device, you can download programs that enable you to turn off tracking. This is so important because it allows you to reclaim your attention. There are features on phones, for example, that allow you to set focus time. If you use Microsoft Outlook for email, it allows blocking out focus time and a separate inbox for messages you want to focus on. Every one of those techniques that allow people to reclaim the space of critical thinking is crucial.

The goal is to safeguard your own cognitive liberty, limit the distractions, give yourself focus time, and cultivate your critical thinking skills by spending more time reading an article before sharing it, thinking carefully about information, and pausing for a second before saying, "Yes, I'll watch the next episode of the show."

The most important skill that humanity has to safeguard against manipulation is our empathy. It's about ►

The most important skill that humanity has to safeguard against manipulation is our empathy. It's about our interrelationships with other people and having greater empathy for ourselves and others. The more we can cultivate that empathy, the better our cognitive liberty will be, and we'll have better safeguards against ending up in technological silos.

our interrelationships with other people and having greater empathy for ourselves and others. The more we can cultivate that empathy, the better our cognitive liberty will be, and we'll have better safeguards against ending up in technological silos. So it's about working to reclaim your attention and your critical thinking skills while continuing to cultivate your empathy.

GRIMM: The example from your book that really stuck with me was when you talked about how, before you had kids, when you were writing an article you used to make a huge pot of hummus, dive into the research, and come up with a whole draft a week or two later. Then when your children were born, they didn't fit into that schedule and you had to relearn how to work in smaller bursts. It's sort of a similar thing that you're talking about here.

FARAHANY: Exactly. I mean, I bought a timer. It's just a little square cube that I talked about in the book. It's the Pomodoro Technique, where I would just give myself the focus time. I turn the timer to 20 minutes and just focus, no email, no distractions, no notifications, nothing during those 20 minutes, other than the focus on the deep critical thinking work that I am doing. My kids have learned to respond to the timer. And I have learned to focus in

those ways. That's cultivating my cognitive liberty. That's the kind of thing that people can do for themselves to find ways in an increasingly distracted world to reclaim their own attention and their own focus.

GRIMM: In some ways, we have been kidnapped by our own technology. And if you get a notification on your phone, you instantly think, "I've got to check that. That could be important."

FARAHANY: And it's designed for that. What people don't understand is that when you're off your device for a while, the algorithms are designed to batch notifications to bring you back onto your device. But you can turn that off to create focus time. There are apps that you can download, but most software also has that. I've turned notifications off on my computer for email. It doesn't even show the number of emails that I have in my bar. I attend to my email when I have chosen to attend to my email rather than when the program tries to pull me in.

GRIMM: It's rare when someone is among the first to visualize the necessary paradigm shift across disciplines and find a way to go forward. And in a time when it seems like we are all so polarized, to have something so fundamentally important as this

bringing people together is very special indeed. We're grateful to you for sharing your insights. I'm sure that your schedule will continue to be busy for quite some time.

FARAHANY: If it means that I will be able to make some inroads on what I think is a really important paradigm shift in society, it's worth it.



PAUL W. GRIMM

is the director of the Bolch Judicial Institute and the David F. Levi Professor of the Practice of Law at Duke Law School. Previously, he served as a district judge (and before that as magistrate judge) in the U.S. District Court for the District of Maryland.



NITA A.

FARAHANY is a leading scholar on the ethical, legal, and social implications of emerging technologies. She is the Robinson

O. Everett Distinguished Professor of Law & Philosophy at Duke Law School, the founding director of Duke Science & Society, the faculty chair of the Duke MA in Bioethics & Science Policy, and principal investigator of SLAP Lab.