

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

- Nita Farahany: What we're looking at is major tech companies who have had an approach to the brain, which is use as much information about how the brain operates to exploit it rather than to enable and empower people. I think it paints a really troubling future if we don't reset the terms of service. And so, cognitive liberty really is meant to connect up all of those pieces.
- Aza Raskin: That's Nita Farahany, the author of *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology*. Today we're going to focus on a technological arms race, which we haven't discussed yet on *Your Undivided Attention*. That is new forms of hardware embedded into mainstream devices like earbuds, which can gather data on our most intimate signals, even in our skulls, our unfakeable brain activity, and in ways that benefit others at our expense. Remember the film *Minority Report*?
- Speaker 3: Well, let's not kid ourselves. We are arresting individuals who have broken no law.
- Speaker 4: But they will. The commission of the crime itself is absolute metaphysics. The Pre-Cogs see the future, and they're never wrong.
- Aza Raskin: Luckily, we're not yet in the version of the future where authorities read our minds and predict our thoughts. But now predictive algorithms can predict our behavior and thoughts to some degree. But Nita says we're already seeing neurotech that is far more invasive than we realize, and that in fact we're at an inflection point that we can see is similar to AI and compounded by AI.
- I'm Aza Raskin. Today on *Your Undivided Attention*, the concept of cognitive liberty, how AI is making the hidden signals of the world suddenly decodable, and how this all intersects with an explosion of ubiquitous, cheap, high-definition hardware. The first question I ask Nita is how this is already playing out in our everyday lives.
- Nita Farahany: People are already quite accustomed to having sensors that are in their watches or on their fingers in the form of a ring, or a Fitbit that tracks basic movements and activity. A decade ago, those were just starting out. People couldn't imagine so many different aspects of their health and wellbeing and everyday activities being quantified.
- Up until now, most neurotechnology devices have been silly-looking. They've been headbands that can be in hard plastic, that are uncomfortable to wear and that you wouldn't wear all day, or they might be sensors that are embedded into a hard hat or a baseball cap. The signal that is how much brain activity could be picked up by those devices and the functionality were really limited.
- What's happened over the past few years is that there have been two things that are converging. One is finding a way to embed brain sensors into watches or

into earbuds or into headphones. So they're part of multifunctional devices just like the rest of the sensors are.

The second is moving beyond what have been really niche applications, like meditation or neurofeedback for therapeutic purposes, to using our brain activity as a way to be able to replace peripheral devices like a mouse or a keyboard.

Part of that has also been the growth of AR and VR and the investments into that space, and recognizing that if you're building a new and immersive way for people to interact, you need new sensors, and that the traditional joysticks or handheld controls are very awkward as a way to interface with those. And so, using these sensors that can be embedded into an AR or a VR headset or earbuds or headphones or a watch is what most of the major tech companies are investing in.

And so, the convergence of all those forces puts us literally at the moment before. Every major tech company has a huge investment into bringing brain sensors into their everyday technologies. We have months to a couple of years before those technologies will become quite mainstream.

Aza Raskin:

I remember in 2019 seeing a study out of Harvard where they took a rhesus macaque monkey. They stuck probes into its visual cortex, sat it in front of a screen that was generating images and how to generate images that maximally stimulated those neurons. The images that emerged were psychedelic. They were images of researchers in masks and other monkeys' faces. It was the first time that ... At least I have seen, where memory was being extracted from matter, that you're able to image the contents of a mind. It's terrifying and scary, but also requires somebody sticking probes into your brain, and that feels like it's going to be a long time until that rolls out into the world. How much more information does actual neuroimaging give you than what's already available given all the other sensors that are out there?

Nita Farahany:

Yeah, it's a great question. It's part of why when we get to the idea of cognitive liberty, it's really meant to be an umbrella concept across all of those. But, importantly, let's realize none of this happens in isolation. I mean if a company is able to read your microfacial expressions and they're able to read your heart rate and your movements and your digital traces and activity, and then you have brain sensors on top of it, is there some missing piece that's adding or does it not add anything at all?

The answer is it seems to add something. What you can get resolution-wise from implants in the brain may be much more powerful than wearable sensors. But what wearable sensors get you that these other microfacial changes don't get you are some of the unexpressed inward and deeper feelings, emotions, and reactions.

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

So it's true. You can pick up with some degree of accuracy, for example, if a person is tired just by picking up how they're moving a steering wheel, or using sensors on a car that look at the stripes on a road and try to figure out have they changed and how you react to it.

Earlier then you can pick up with those algorithms, you can pick up fatigue levels in the brain, because there are signals that change, and your pattern of electrical activity in your brain changes as you go from being wide awake to being sleepy. Those inward reflections are things that AI has gotten better and better at, but not precise. It gives you resolution, it gives you additional insights, and it gives you some of the most inward feelings and reactions as well as evoked reactions to information. You can literally mine a person's brain with environmental stimuli to get information that's stored within. You can't do that as well with microfacial changes.

Aza Raskin: So you argue in your book that we need a new definition of cognitive liberty. What defines this concept and why is it especially important now?

Nita Farahany: So cognitive liberty I think is an update to liberty for the digital age. So it's built on classical notions of liberty, of right to self-determination, and a right to self-ownership. But the way I've been defining it is the right to self-determination and a right from interference with our mental privacy and with our freedom of thought.

Those align with human rights concepts. Self-determination is the basic idea of dignity and self-ownership that underlies most other human rights. But it also aligns with this idea of a right to informational self-access and a positive right to really be able to access information about our brains.

Technology can be deeply empowering for people if it is technology on terms that actually align with human values. I think it really needs to cover this idea of freedom from interception, manipulation, and punishment for our thoughts, thought being a robust concept because it's an absolute human right. And so, part of what we're doing right now is trying to read each other's minds, and it can't be that every interaction is off-limits. So it's just certain kinds of things that we think are problematic as interfering with our freedom of thought.

Aza Raskin: Why should any of this matter? Why isn't this just something that's academic?

Nita Farahany: So I think that our interrelationship with both each other and with technology has fundamentally changed what it means to be human. But it's a struggle and a worry without a framing or a naming to help people really understand what is it that's at issue.

And so, I think it's both naming for people what it is that we're all searching for, which is this idea of cognitive freedom, but also taking it from an academic

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

concept to translation into what that means for rights, what it means in context in the employment setting, for example, or in educational institutions, and how those rights translate to specific policies, how that translates to how products should be designed to respect the cognitive liberty of individuals.

To me, it's really how do you enable human flourishing in an age that is far more interdependent with technology. And so, this concept I think helps guide us forward in how to do so.

Aza Raskin:

What I hear you doing is defining a new kind of commons. Listeners to the podcast will be familiar with the three rules of technology that we've been positing, and that the first rule is when you invent a new technology, you uncover a new species of responsibility. It's not always obvious what the responsibilities are. The examples we give are we didn't need the right to be forgotten to be written into law until the internet can remember us forever. It's surprising that what does HTML have to do with the need to be forgotten? Or we didn't need the right to privacy to be written into law until Kodak started producing mass-produced cameras.

What happens is when there's a new technology, it makes what was illegible legible. Suddenly AI is making more and more of the human condition legible, which means it's suddenly able to be exploited. If the technology confers power, it starts a race. If you do not coordinate the race, it ends in tragedy, which is really the story that we've been telling of the attention economy and the engagement economy.

A lot of what you're talking about, of course, is deeply entwined in what CHT's work has been. I know that social media companies are now investing heavily in this space. Of course there's Elon Musk's, Neuralink, and that recently got approved by the FDA. He obviously also owns Twitter. Can you talk listeners through how far advanced these technologies actually are?

Nita Farahany:

Right. So your listeners will be deeply well-read on the fact that most platforms and technologies have been looking at how do you exploit cognitive heuristics and how do you exploit brain mechanisms in order to keep people addicted or keep people engaged in ways that diminish brain health and wellness rather than expand it. Even as it introduces new opportunities for connectedness, it creates greater distance between us and mental health problems and disorders.

So if you take companies that have long held the approach to the brain to be how do you exploit it, how do you diminish it, and how do you addict it, and then you give those same companies the capabilities through the acquisitions that they're making ... So Meta's huge investments by acquiring CTRL-Labs in 2019, and they plan to launch neural interface through EMG Technologies in 2025 that's integrated into a watch, or Apple's new Apple Pro Vision, which is using pupillary response. They had a team of neuroscientists and

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

neurotechnologists who were working on trying to make inferences, or the possibility of Apple acquiring earbud technology that could integrate into its AirPods.

If you take Microsoft's huge investments into the space of neurotechnology, including its research into understanding how the brain reacts to different information in the workplace to create what's called a cognitive ergonomic workspace, a workplace that is designed to be more responsive to the human brain, or Snap's investments of acquiring NextMind out of Paris, an EEG-based company, to integrate into AR and VR.

What we're looking at is major tech companies who have had an approach to the brain, which is use as much information about how the brain operates to exploit it rather than to enable and empower people. And where the business model has been built on commodification of personal data and then you give to those same companies the capability of having much more precise measures of the brain and how it reacts, and then the ability to commodify all of that brain data, whether it's for neuromarketing, for microtargeting, or for manipulation of elections or other processes. I think it paints a really troubling future if we don't reset the terms of service. And so, cognitive liberty really is meant to connect up all of those pieces.

Aza Raskin: I think what most people don't realize is because of the combinatorics of it all, how fast this is moving, the example of what should an image-generating, an art-generating AI have anything to do with the ability to read brains-

Nita Farahany: Right. Everything.

Aza Raskin: ... and it turns out it's everything, right?

Nita Farahany: Yeah.

Aza Raskin: Because what does AI do? It's giving the power to decode, translate, and generate the languages of nature. It turns out the languages of nature are images, are videos, are fMRI, are EEG, are DNA. And so, the ability to go from language to image suddenly means that if you hook that up to the language of fMRI to image, you get state of the art in brain reading, and that happened overnight.

Nita Farahany: I think even more profound than that.

Aza Raskin: Yeah.

Nita Farahany: Let me tell you why. It isn't just what you're seeing but what you're imagining or what you're dreaming that could be decoded. That was all primarily decoding

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

from the visual cortex in the brain, or other studies were looking at the motor cortex to say this is the speech that a person is generating.

In the past year, the studies have shown that there is language representation that is redundant across different regions of the brain, and it's not just auditory, like what you're hearing, motor, the kind of speech you're forming, or visual, which is what you're seeing. It's distributed across the brain. If you connect that up then to wearable sensors, what you're looking at is the ability to have far fewer sensors pick up brain activity, have redundant representation of language, and then the ability to fill in with generative AI and associate text to image.

So you have all of those things happening at once, and this fundamental shift where scientists are able to decode from different regions of the brain with much greater precision and accuracy than before. So the studies with generative AI that have come out in the past year have been startling and the seismic shifts that's been happening.

Aza Raskin: To break in here, one of the core things to understand about what AI is doing is that it is taking all of the signals of the world that we couldn't understand and decode before and making them legible.

We can start to decode wifi signals bouncing around a room and determine who is standing where and in what pose. You can look at brain patterns and understand what somebody is thinking or seeing.

When that hits much cheaper sensor data, suddenly way more about what human beings are doing, thinking, feeling becomes legible to technology, and that opens up brand new ways of being exploited.

I really think just it is so important to just pause and dwell that it is the ability to decode your inner monologue the things that you're thinking, the ability to decode ... Your dreams are not safe ... the things you are imagining. It's really, I think, sobering that what we are learning from the more advanced imaging technologies like fMRI with AI are then able to be backported to these less advanced technologies that may well sit, say, inside of an Apple EarPod.

Nita Farahany: Right.

Aza Raskin: That's something which I think is deeply surprising because when I imagine putting on a helmet, I'm like I'm never really going to do that.

Nita Farahany: Right.

Aza Raskin: Maybe I'll be forced to do that ... And maybe we can talk about authoritarian regimes, both corporate and political. Maybe that's where we go next. But that that same kind of technology, it is unknown what the capabilities will be with

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

consumer-grade hardware that I just put in a set of EarPods, and then I am leaking my thoughts, that's really scary.

Nita Farahany: Let me add one thing to that, which is I think the stark image that you've just created, which is you are aware of the risks, at least to some degree, when you put on a big clunky helmet, or if you're in a fMRI machine that requires you to know what's happening and actively consent to the process. But when the sensors become invisible, because they're embedded in our everyday technology, and those same technologies are multifunctional, you're taking a conference call or you're jamming to music, you're not thinking about the sensors that are embedded and how much data is actually being generated. The nature of the data from brain activity is what we call raw brainwave activity.

So to your point, because it has the capability for being mined for so much more over time as capabilities advance, if that data is stored, it can be returned to over and again to be probed for so much more.

Aza Raskin: This starts to, I think, point at this myth in, well, if I don't want to participate, I'm just not going to use this stuff. You may be unaware that you're using this stuff. But let's go to the two other ways that you might be forced into using it. One is like the corporate route and one is the authoritarian route. I'd love for you to talk about the risks in both.

Nita Farahany: Sure. So it's interesting because you can see the risks already being realized, because the misuse cases are already beginning, even though the technology is still, in many ways, at the earliest stages of dissemination across society.

So one of the earliest companies in the space was a company called SmartCap that had a life band that has EEG, electroencephalography, sensors embedded in it, where workplaces worldwide, more than 5,000 companies, have partnered with SmartCap to have long-range truck drivers or people who are working in mines or pilots be required to have their fatigue levels monitored by monitoring their brain activity. That could give you a more precise interpretation of a person's fatigue levels by being able to see as they transition to those earlier levels of sleep.

But it's not a choice that these employees are having. In fact, there was one group that, based on union activity, was able to prevent the mine from requiring them to wear these SmartCaps. In many ways, I think the way SmartCap is doing it is as privacy preserving as possible. They're keeping all of the raw brainwave data on device. They're overriding the data continuously. They're only providing the extracted interpretation of the score from one to five of fatigue levels. But the fact that there are employees that are already being required to wear it is startling.

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

It's already happening. I think people don't realize that. There are very little protections, at least in the US, for employees, the idea that you can just quit and go elsewhere, not if everybody is using the same technology and not if you don't have the upward mobility to make that easy to move between jobs. It isn't just truck drivers or factory workers. It's knowledge workers where productivity is being tracked already through a suite of different technologies that are put onto their workplace computers.

Now, as you start to have these sensors that are issued by workplaces, it's possible that they could have access to all of the information that it's collecting to mine it for so much more.

A lot of companies in the US have also launched brain wellness programs to bring down stress levels and to address mental health disorders in employees. The problem with most of those wellness programs is they're not subject to the same kinds of privacy rules as HIPAA-governed health insurance plans are. And so, a lot of the data from those wellness programs are also being mined and sold and repackaged by employers.

And so, in those settings, I think as neurotech gets integrated, I think of the risk of discrimination. Authoritarian regimes are even scarier.

Aza Raskin: Yeah. Well, that's interesting. Let's dive in there. Which authoritarian regimes are already making it mandatory, and then what are the implications?

Nita Farahany: So we know already from reports coming out of China, the way it's being used in China. There are also reports of law enforcement using it, which we'll come to in a moment, from places like India and Singapore and the UAE, using brain tech to be able to interrogate a person's brain in a criminal setting.

In China, the earliest reports were factory workers being required to wear hard hats and baseball caps that were embedded with EEG sensors to pick up their attention and their fatigue levels. Students in a classroom in China were reportedly being required to wear headsets that were issued by a US company to track their attention and mind wandering. That information was being sent to a console in the front of the room for the teacher, being sent to parents and being sent to the state. Students reportedly being punished based on what their brain activity revealed.

In the workplace, employees being required to have their brain activity tracked in a setting where you're afraid of how that information can be used and misused.

Then reports of this brain mining that we were talking about in that same setting. So showing people in China political messaging, like communist messaging, and then seeing how their brain reacts to that information to try to

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

get at their inward feelings about the regime and whether or not they're true believers or not. All of that is reportedly happening in China right now.

There are attempts, apparently, at developing brain-controlled weaponry. One of the terms that places like NATO have been talking about is cognitive warfare as the brain being the sixth domain of warfare.

Aza Raskin: Just breaking in here to mention that the five domains of warfare are defined as land, sea, air, space, and information. Nita is saying that cognitive warfare may become the sixth, and I would argue that in fact it already has. As Marshall McLuhan said in 1968, World War III will be a guerrilla information war with no division between military and civilian participation.

Nita Farahany: What we know is that there are a lot of reports coming out of China that there are significant investments in brain-controlled weaponry. The Biden administration, in December of 2021, issued sanctions against a number of Chinese companies for purportedly trying to develop this technology. Whether they have or they have not, the fact that it is something that people are worried about and that there appears to be an investment in is I think of significant concern.

Aza Raskin: Yeah, I think we are absolutely already in cognitive warfare. It used to be that if you wanted to pit Americans against Americans, it took a lot of work to take out the right kind of op-eds and get the right kind of content. Now Facebook or TikTok will give you hand glove treatment to deliver the perfect incendiary statements to the exact fissure lines of society, to inflame them, to pit like fellow citizen against fellow citizen. This just feels like it enhances the ability and efficacy of making those kinds of messages.

Nita Farahany: So the frightening thing is it's possible for that to happen without a person even being aware of it. I mean most people aren't aware of the way in which they're being conditioned or polarized within social media or within the kind of messaging that they have access to. But if you imagine both access to the platform itself, censorship tools together with how a person is reacting to that information and being able to precisely change that without a person even being aware of it.

On the one hand, you can use explicit punishment, but you may not even have to. It's the subtlety with which these changes can be made imperceptible to humans that I worry about a lot.

Aza Raskin: Yeah. Yeah. To go back to where we started in the interview of that image of a rhesus macaque monkey with electrodes in its brain and AI generating images, that's not a sci-fi scenario in terms of right now because that's already happening on social media just without the brain probes-

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

Nita Farahany: That's right.

Aza Raskin: ... plugged into your brain, and it's using fairly unsophisticated signals of what you click on in terms of likes. It's not give people what they want. It's show people what maximally activates their nervous system.

Nita Farahany: Take a company like Intertek, a Chinese-based company that has issued many thousand headsets that people use to do mind-controlled car racing or neurofeedback. Imagine now TikTok has access to the Intertek headsets, picking up brain sensors while you're on platform on TikTok and can have this closed-loop system picking up brain activity. That brain activity in real time is being fed into an algorithm that then generates content and doesn't just give you curated content.

All of that is just a world of sensing where what's changing your brain activity is your environment. There are a whole category of devices that are being developed and that are already in existence that also provide neurostimulation.

The more precise that stimulation becomes ... You were talking about punishment. What if it's a little literally Pavlovian shock that you get in response? That's not so farfetched when there are already devices that exist that can provide neurostimulation in addition to neural sensors.

Aza Raskin: So you're saying it's not just reading the brain, it is writing to the brain.

Nita Farahany: Right. There are soft ways to write to the brain. Everything we do writes to our brains in some sense, and that inputs change our brain activity, how the brain fires, and ultimately what brain signals look like. But that happens through inputs that are not literally providing little jolts of electricity to the brain, or on wrist neurostimulation that stimulates the motor activity in your body in response to whatever's happening with your brain signals.

Aza Raskin: And so, then just to make it, in some sense, more real for the listener, what somebody like a TikTok could do is they could hire 10,000 people to wear these brain caps, to figure out how brains work correlates to signals that they can read, how the phone is being held, in what orientation, how much jitter the sensor has, how often they are moving around. Then the vast majority of people don't actually have to wear the brain caps for them to gain access to the new power that the technology affords.

Nita Farahany: That's right. That's been happening not on TikTok necessarily, but through what's called neuromarketing for a while now, which is people are paid to watch advertisements or to engage in whatever set of activities they're being asked to engage in while the brain sensors are measuring their responses. That's how these devices have been trained to pick up attention levels or fatigue levels, is thousands of people watching inputs, seeing how the brain reacts, and then

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

being able to change products and designs or advertisements to evoke a specific kind of response in individuals. When these sensors become widespread, the ability to do that neuromarketing at scale will become increasingly powerful for companies.

Aza Raskin: And so, really what you're pointing at is a new level of asymmetry of power.

Nita Farahany: That's right, and a very frightening, I think, asymmetry of power, because both of the last fortress of privacy falling, but also the subtle ways in which our brains and mental experiences can be shaped and reshaped without even being aware that it's happening.

Aza Raskin: I think one argument would be like this stuff is scary. Obviously I'm not going to use it. I'm just not going to buy the product. I'm not going to let my kids use it. Will this ever control the means of participation? What do you respond to that?

Nita Farahany: Yeah. So I mean first I would say we've talked mostly about the risks rather than the benefits. It is, I think, important to recognize the benefits. The ability to be able to track things like stress or be able to track your attention and understand when your attention is being hijacked, to be able to track over time things like cognitive fitness levels, or to be able to see the earliest stages of dementia or depression, or, for me, I'm a chronic migrainer, having earlier indications that would allow me to intervene more quickly could usher in a new era.

Aza Raskin: Like a Fitbit for your mind.

Nita Farahany: Yeah. I mean that's what I think Elon Musk called Neuralink at some point, the Fitbit for the brain. I don't think Neuralink will be it. But the idea that we could have valuable information from quantifying our brain activity, I think is something a lot of people will opt into, particularly if there are robust privacy measures and if the right to mental privacy is codified in law, if it's something that truly exists and changes the terms of service.

I think the idea of opting out is a limited generation idea, and that the more capabilities that are developed and the more natural our interaction with the rest of our environment is using neural interface, the more likely it becomes ubiquitous, which is why I think it's so critical that we move quickly to recognize cognitive liberty and codify it in law, because we're not at the inflection point where this is widespread across society. It's still more limited applications. It's still major tech company investments, but not widespread dissemination yet.

It's still possible to buy AirPods without health sensors embedded in them. But as those become obsolete and your only option is to have brain sensors integrated, I think there's virtually no way to claw back rights.

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

Aza Raskin: Well, actually, that leads to a question about how do you start implementing something like cognitive liberty and what are the frameworks? One of the things we've learned as we've started to delve into the AI space is that, honestly, companies don't really care about ethics and responsibility.

Nita Farahany: Yeah.

Aza Raskin: Well, they say they do until market forces start to steamroll it. The one thing that they listen to, the language that companies respond to is liability.

Nita Farahany: Yeah. So I mean I've been thinking about it on five levels. One is to move quickly to update our existing human rights. The reason I've started there is that's global. That means recognizing cognitive liberty as a new human right which directs the updating of three existing human rights, self-determination to be an individual, right to self-determination, mental privacy to be explicitly included within the right to privacy, and freedom of thought to protect more broadly than just religion and belief.

That's great. That sets a legal norm. It creates an enforcement mechanism. But as we know, people violate human rights all the time. And so, you have to move beyond just a human rights regime into what that looks like at a national level. That means context-specific, I think, regulation as well. I don't think it's enough in the US, for example, to say, "Well, the first amendment ought to also include freedom of thought."

I think what we need is national legislation in a context specific way that addresses these issues. What does it mean in the workplace if there's mental privacy? That would mean that there are limited use cases that have to be governed by justifications for gaining access to any brain data. So maybe you can gain access to fatigue levels, but it would mean that the rest of it couldn't be gathered.

Okay, that's the rights level. We can go into depth on that, but I have then been thinking about it on how do you embed cognitive liberty into research design by researchers to answer the empirical questions about how do we create both mechanisms by which people can exercise cognitive liberty and what does that look like? How do we embed it into commercial design from user-level controls to rules that enable people very easily to opt in to brain protecting and mechanisms? What does it mean from aligning incentives in society? How do we make brain health and wellness a national and international priority so that you start to actually try to maximize brain health and wellness? Then what are the tools that we need to enable individuals to cultivate cognitive liberty?

So one example there is one of the big concerns as generative AI explodes is the amount of generated content that we have no way of being able to decipher as real versus fake.

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

Aza Raskin: Correct. This and next year are the years that video and photographic evidence cease to be effective.

Nita Farahany: That's right. And so, how do we help people both recognize you can't trust what you see or read anymore and how they can safeguard themselves against the risks of manipulation? There's a lot of really great research, for example, from marketing where they've looked for a long time at if something's labeled as advertisement or a marketing campaign, your ability to resist it goes up versus if it's unlabeled. And so, there's some content authentication initiatives that major tech companies have signed onto to start to try to create provenance of images. Those are a good start, I think.

Aza Raskin: Is a human rights framework through the UN, is that the way that you bind the race?

Nita Farahany: Yes.

Aza Raskin: Okay.

Nita Farahany: So I mean that's why I think you have to have the human rights framework as part of it, right?

Aza Raskin: Yeah.

Nita Farahany: Because there's trying to figure out the carrots. That's the million-dollar question, right-

Aza Raskin: Yeah.

Nita Farahany: ... is how do you actually incentivize companies to do so? And so, the liability model focuses on the human rights framework to say, "Look, fundamental to human dignity, fundamental to human flourishing is having the right to cognitive liberty as an international human right," which changes the default rule of what you can and cannot do as a company.

Aza Raskin: And so, then when there's a company you do that thing, what happens? Walk me through the actual ramifications.

Nita Farahany: Well, I mean so you have to have enforcement, right? So all three rights that I talk about, self-determination, freedom of thought and privacy, are all codified within the ICCPR, the International Covenant on Civil and Political Rights. It's like a court that oversees a treaty that we and other countries have signed onto that should bind us.

Now the question is how much teeth and enforcement does human rights have? It has, in many ways, more, I think, of a shaming function than it does always an

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

enforcement function. You can't haul the United States into international court and put us in jail somehow. But what you can do is to have a naming and shaming. So you have global norms that develop around it, and you have a court of redress. So you have the ability to actually file a complaint, to have that complaint heard, to have opinions that are issued and recognized.

Then at least, theoretically, that trickles down to what we have agreed to as different countries to align our national policies with. Then it's identifying what those are. And so, I think it starts to have an effect on what national legislation and implementing policies look like as well.

But it's not enough. I wish that human rights were enough. It's not enough because authoritarian governments continue to violate human rights, because people neglect and countries neglect their human rights obligations, unfortunately, all too frequently. That's why I think you have to start to embed it across each of these different dimensions as well.

Aza Raskin: Yeah.

Nita Farahany: I'll make a quick side note here to say I don't think it's enough to just update those three existing human rights. I think we actually need the naming of the thing that we are trying to protect, which is cognitive liberty for that function.

Also, some of the work that I've been doing with NIH, the Neuroethics Working Group of the US Brain Initiative, is to try to figure out if there's different degrees of sensitivity for different kinds of brain data. Not all brain data is going to be equally sensitive, and maybe the technical solutions that are addressing it may be different. You might think that images and thoughts in your mind are far more sensitive, for example, than your fatigue levels.

Privacy by design solutions can and should be implemented. That to me is part of what is taking cognitive liberty from a theoretical concept to a human rights concept, to what does that literally mean technologically and specifically should be embedded into product design.

Aza Raskin: I have two more questions to go. One is there's another framework that I know that you've talked about, we've actually talked about it on the podcast, around asymmetric influence. That's the concept of undue influence.

Nita Farahany: Right.

Aza Raskin: Because it's actually tricky to know when is something like legitimate persuasion and when is something undue influence.

Nita Farahany: I'll tell you, this was the very hardest chapter for me, and I think it's our hardest problem today, is figuring out where does persuasion end and manipulation

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

begins, not just from a philosophical concept but from a legal freedom of thought. If it's right against manipulation, what is that? Because we're trying to persuade each other all the time. When does it become problematic?

So for me, in trying to unpack this, I turned to a lot of different sources and came up with what I think is best described as your freedom of action, and understanding that there are a lot of inputs we don't have control over, that we never will have control over, whether my plane was late, what the weather is outside, all of which affects my mood and our inputs that are beyond our freedom.

And so, free will as a robust concept, I think, is a little outdated. Freedom of action, though, I still believe we have. By that I mean we maintain flexibility of action choices, trying to essentially hack into our automatic reactions to bypass our action choices and to put us into auto mode rather than critical thinking mode in ways that are harmful to us is manipulation. Those two pieces, I think, define for us what manipulation means in the digital age.

Aza Raskin: I wonder if there is a way of saying the degree to which I can know you better than you know yourself, which is to say I can predict you better than you can predict yourself, and, hence, change your actions is the degree to which I need to be in a fiduciary relationship to you. That is to act in your best interests, just like a lawyer knows more about how to exploit your lack of knowledge. They have to be in a fiduciary relationship that any of these technologies, you can get a degree of how much better they can outcompute you and, hence, they should be bound to that fiduciary.

Nita Farahany: Yeah. That's then what brings us back to this question of how do you align incentives to make that happen?

Aza Raskin: Right.

Nita Farahany: I don't want to be on my screen more than 10 minutes today. Then there are technologies and techniques designed to hack into my automatic actions in my brain, my cognitive biases and heuristics that I operate within, overriding whatever my desires are, who I've committed myself to. So how do we change business models and align incentives in society so that people do act in a fiduciary responsibility position where not only have these technologies revealed new insights, but they've put them in a position where the incentives need to be to enable human flourishing, not diminish human beings?

Aza Raskin: Yeah. I guess that gets to one final question. I think there are some people that would argue we have to use this kind of neurotech. There's no choice, because if we're going to compete with AI, we need to do it. We need to augment ourselves.

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

Nita Farahany: Well, I mean that's as if AI is inevitable, as if the race that we have created is inevitable and as if we have no choice. It's this path-dependent. So I'd first question the idea that we have created the need for us to compete with something that we've created and whether or not there are appropriate guardrails we can put into place.

But the second is, look, I'm imagining a more hopeful world, a world in which we use the technology to gain insights about ourselves. We use the decoding to reclaim control over technology that we have allowed to control us.

Aza Raskin: I think smuggled into the ... That we can choose is that it's not individual choice, as we're saying. Like you may be forced to use this technology in the same way that I was very hesitant to ever upgrade my phone to the ones that use the face unlock. Then I stuck with an old phone that used my thumb, and eventually I couldn't buy a phone anymore that did that.

Nita Farahany: Your existence becomes obsolete if at some point you don't assimilate the technology.

Aza Raskin: Exactly. Like technology controls the means of social participation-

Nita Farahany: Right.

Aza Raskin: ... and social participation is absolutely necessary. So when you say it's up to us to choose, it's not individual choice. It's really a reforming of the incentive landscape that our society runs on. That determines which direction the technology goes.

Nita Farahany: Yeah. I think we can start with human rights. That automatically puts pressure onto incentive systems to actually align. But we have a massive power imbalance right now between individuals and tech giants that have set the terms of humanity. We need to reformulate that, and that fundamentally means stopping and saying how do we realign incentives to be human-centered flourishing rather than human-centered diminishment?

If used in that way, rather than the way that some transhumanists talk about it, which is increasing through synergistic brain-computer interface the augmentation of humans, if instead we use it as a way to study ourselves and to understand our brain health, our brain actions, the ways in which our cognitive biases and heuristics are tapped into and reclaim control and cognitive freedom, then I think we actually can compete well with AI or any other technological system because we enable human flourishing. We enable humans to expand rather than diminish.

It depends on how we use the technology. If we use the technology to further addict us and automate us, I don't think we'd compete. I think all we do is give

Center for Humane Technology | *Your Undivided Attention* Podcast
[Protecting Our Freedom of Thought with Nita Farahany](#)

all of our brain activity to AI to be able to improve the systems to outcompete us. If we use it as a way to reclaim what it means to be human, I think the potential for humanity to have a golden age of flourishing is possible.

Aza Raskin:

Nita Farahany is the author of *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology*. She's a distinguished professor of law and philosophy at Duke University, where she teaches in the law school, chairs the Bioethics and Science Policy Program, and she serves as the founding director of the Duke Initiative for Science and Society. From 2010 to 2017, Nita worked on the Presidential Commission for the Study of Bioethical Issues, which she was appointed to by President Obama.

Tristan Harris:

Your Undivided Attention is produced by the Center for Humane Technology, a nonprofit working to catalyze a humane future. Our senior producer is Julia Scott. Kirsten McMurray and Sarah McCrea are our associate producers. Sasha Fegan is our managing editor. Mia Lobel is our consulting producer.

Mixing on this episode by Jeff Sudakin. Original Music and Sound Design by Ryan and Hays Holladay, and a special thanks to the whole Center for Humane Technology Team for making this podcast possible.

Do you have questions for us? You can always drop us a voice note at humanetech.com/askus, and we just might answer them in an upcoming episode.

A very special thanks to our generous supporters who make this entire podcast possible. If you would like to join them, you can visit humanetech.com/donate. You can find show notes, transcripts, and much more at humanetech.com. If you made it all the way here, let me give one more thank you to you for giving us your undivided attention.