Vincent Munoz:

I think what we need to do is explain how our principles of free speech, free inquiry will help serve the cause of justice.

Betty Friendan:

The First Amendment, the constitutional freedom of speech and freedom of conscience that is the bulwark of our democracy.

Bettina Apthekar:

There was a passion in what was being said, affirming what people considered a sacred constitutional right. Freedom of speech and freedom of association.

Michelle Deutchman:

From the UC National Center for Free Speech and Civic engagement, this is Speech Matters, a podcast about expression, engagement, and democratic learning in higher education. I'm Michelle Deutchman, the center's executive director and your host. In the ramp-up to the fast-approaching national election, speech matters has addressed voting and student civic engagement as well as artificial intelligence, AI and speech. But what impact is AI and other technological advances having on our election processes already? And moving forward, how will generative AI change how political messages are spread and how we communicate with one another about the state of our democracy?

This month's guest, Jonathan Mehta Stein, executive director of California Common Cause joins to help us consider these critical questions. But first, let's take a listen to Class Notes, a look at what's making headlines. Returning to a story from last episode, United Auto Workers 4811, the union that represents almost 50,000 graduate students and other academic workers at the 10 UC campuses conducted a standup strike to protest the UC's handling of pro-Palestinian demonstrations, which they argue constituted an unfair labor practice.

The strike began three weeks ago at UC Santa Cruz and spread to six additional campuses. UC responded by filing complaints with the state Labor relations Board, which twice declined to block the work stoppages. This led UC to sue the union for breach of contract, and last week an Orange County Superior court judge handed down an order halting the strikes. UAW says it will continue its efforts calling the ruling as the UC system's latest attempt to avoid accountability.

Harvard and MIT have announced that they're changing their hiring practices, opting to discontinue asking faculty candidates about their commitment to diversity and inclusion, colloquially known as diversity statements. In response to the announcement MIT's, president Sally Kornbluth commented that, "We can build an inclusive environment in many ways, but compelled statements impinge on freedom of expression and they don't work." At Harvard, candidates to the university's largest faculty division are and Sciences will instead be asked to describe how their efforts strengthen academic communities and how they would promote learning environments in which students are encouraged to ask questions and share their ideas.

At universities around the country, faculty have grown concerned over intrusions by administration and governing boards into systems of shared governance, which the AAUP defines as "the joint responsibility of faculty, administrators and governing boards to govern colleges and universities." Recently, the University of Kentucky Board of Trustees took up a proposal, backed by the university's president, which would dissolve the 100-year-old university senate, replacing it with a faculty senate which would have less power to decide on academic programs and courses. Faculty argue that they alone should have control over curriculum considering their disciplinary expertise.

Concerns about shared governance have also been raised in Arizona, where a bill aiming to limit shared governance at public universities is moving through the state legislature as well as at Harvard where faculty have raised concerns about the power of the Harvard Corporation to make core academic decisions for the university.

Now back to today's guest, Jonathan Mehta Stein currently serves as executive director of California Common Cause. Since 1970, Common Cause has been working to hold power accountable through lobbying litigation and grassroots organizing. Its nonpartisan pro-democracy work has helped pass hundreds of reforms at the federal, state and local levels. Before becoming executive director of California Common Cause. Jonathan served as a member of California Common Cause's board of Directors and served four years as board chair. Previous to this role, Jonathan spent four years as the head of the Voting Rights and Census program at Asian-Americans Advancing Justice, Asian Law Caucus, where he worked to increase access to California's democracy for historically disenfranchised communities including immigrant and limited English-speaking voters, communities of color, low income communities, and people with disabilities.

Jonathan also worked as a voting rights staff attorney for the ACLU of California as well as served as a commissioner and chair of the city of Oakland public Ethics Commission. And if that weren't enough, while receiving his master's in public policy in JD from UC Berkeley, Go Bears, Jonathan served as the student regent of the University of California's, board of Regents where he advocated for access, diversity and affordability for the 230,000 students in the UC system. Welcome, Jonathan. We are so happy to have you.

Jonathan Mehta Stein:

Thank you so much, Michelle, and I appreciate the very, very kind and probably too long introduction. I appreciate it so much.

Michelle Deutchman:

All right, so as a former Bear myself, I feel like we must begin our conversation at Cal. Can you tell us a little bit about what brought you to this work and how if at all your time at Berkeley and as a UC student Regent influenced your path?

Jonathan Mehta Stein:

Berkeley was an incredible opportunity to explore my passion for public service because every single issue is present on campus. Every tool in the activist toolkit is present on campus. It's a really easy place to fall in love with public service and to find what you're passionate about. The student regent role was probably the most fascinating, most unique, but also most difficult and most taxing experience of my life. It was a brutally different difficult moment for the UC because it was a moment when the state was cutting hundreds of millions of dollars every year and the administration was contemplating very, very substantial tuition increases. And that was happening concurrent with the Occupy movement, which many of your listeners will remember. And it just led to an explosion of activism on every campus.

And so I was caught in the crosshairs of that slash trying to find a way to mediate between the administration and the student activists, giving voice to student protestors who took over every single regents meeting while also recognizing that I could be... I had access to rooms that they did not, and so I could carry that message in a different way. I learned a ton, I learned what I could endure frankly, but I also learned how to be strategic when I thought that in my first year student regent, I had to be as hard charging as possible.

I couldn't possibly have a relationship with any of the regents because that would represent me being co-opted by the power structure. Over time I realized that there was a more strategic way of going about doing things.

Michelle Deutchman:

Thank you so much. I really appreciate your candor. And it's going to be hard for me not to go off on asking you more questions about that experience, but you know what, we'll save that for another podcast episode. I think I want to obviously turn to Common Cause and the work that the organization accomplishes. And you do that through litigation and lobbying and grassroots organizing and we're barreling towards this election in November. And I'm curious what some of your organization's larger priorities are, as we move closer and closer to that date, and then of course following the election as well.

Jonathan Mehta Stein:

Yeah, so Common Cause is unique among national democracy organizations in that it has state chapters or state organizations in 30 plus states. And in every one of those states, including our team in California, we operate what we call an election protection program, where we are putting trained volunteers in hundreds, maybe thousands of voting sites in each state, in order to ensure disability access, in order to ensure language access, in order to ensure that voter harassment and intimidation is not occurring. In a place like California, we're very focused on language access and making sure that translated ballots mandated by law are actually available in other states. They might be stopping voter challenges or stopping poll worker harassment. It really takes a different shape in every state based on that state's needs, but it's an amazing program and the impacts are really, really meaningful and substantial.

We're also a policy organization and so we're working on a bill right now, AB84 in the state legislature that would dramatically expand language access in California's elections. We've made a lot of progress in California on voter registration access and ballot access, but language access in our elections for our limited English-speaking immigrant voters is an area where the state lags behind other states, which is a real shame, because we have millions of limited English-speaking voters in California, and we've started this new initiative called the California Initiative for Technology and Democracy, which I know we're going to talk quite a bit about. And it has three bills in the state legislature currently that we're trying to pass in order to respond to the danger of AI-amplified disinformation in our democracy.

And that same initiative is trying to level up the public's understanding of the threat posed by AI-powered disinformation, because any bill we pass in 2024 will become effective in 2025, which means that if we're going to protect voters and protect our communities from this sort of accelerated disinformation, this year in this cycle, it is going to be up to all of us. It's going to be up to trusted voices and trusted messengers in the community to make sure everyone knows you need to be a more skeptical, smarter consumer of political information in 2024 and going forward.

Michelle Deutchman:

And I'm sure we're going to get to the fact that higher education is one of those institutions that can help with that. But let's go back and talk about CITED. I know that you launched it about six months ago, and again it's California Initiative for Technology and Democracy, and on the web page when you go there, it says the threat of election disinformation turbocharged by AI is too grave. So why don't you tell us a little bit about the initiative and the goals, and you can even hit on some of that legislation and then we'll go from there?

Jonathan Mehta Stein:

The California Initiative for Technology and Democracy, it was really born out of two realizations. The first was that Congress was not going to act in defense of our democracy, and California as a result, had an opportunity to lead the domestic US response. The second was that Sacramento did not have the policymaking infrastructure it needed in order to do this work well. Sacramento had great expertise on health policy or education policy or environmental policy, but when it came to tech policy, the default was that you went and asked the tech industry trade associations and their lobbyists what to do or how to

act. And the response you got back more often than not was this bill would break our technology or this bill will stifle innovation.

And as a result, very little got passed and we thought it was essential to build an unbiased independent source of tech policy expertise specific to democracy issues so that when the legislature or policymakers or thought leaders in the state wanted to act in defense of our democracy around platform accountability or any number of other things in the tech policy space, they had a source of expertise, interdisciplinary expertise they could go to in order to get real honest advice that wasn't connected to industry.

And we just think that the moment for this is upon us whether we like it or not. We have entered a monumental election year. Our democracy is already weakened, trusted institutions in the media at our all time lows. The truth feels like it's under assault every single day. And now on top of all that we have, we're entering the first ever AI election, and generative AI deepfakes are going to inundate our political discourse. Voters are not going to know what images, audio and video they can trust, and this is not something that you have to be a sophisticated actor in order to use.

The tools we're talking about here are so powerful, they're so easy to access, they're so low cost that basically any candidate, any conspiracy theorist, any foreign state, any online troll now has the tools they need to destabilize our information ecosystems and to deceive voters. And so you put all that together and it was just clear that we needed to do something. And so we founded CITED last year, we're really pleased at the progress we've been able to make in a short amount of time and we're hopeful about our legislative package.

Michelle Deutchman:

So when you paint the landscape, it's pretty terrifying. And I think that obviously the work that CITED is doing is exceptional and so important, but I imagine that a lot of listeners are thinking to themselves, what are some of the things that they can do and think about without the platform of being a legislator? And so I don't know if you have any thoughts about what your person who is engaging with different platforms and social media and thinking about the election can do, and then we'll kind of follow up with a little bit what higher ed might do.

Jonathan Mehta Stein:

Yeah, so let me start by naming for people what they need to be on the lookout for. So there has been examples of generative AI deepfakes targeting politicians or the integrity of elections in a wide range of countries already. So folks may be aware of where it's touched the US. So for example, the "fake Joe Biden." Joe Biden robocall in New Hampshire telling New Hampshire Democrats to not vote in the Democratic primary. That was an example of it touching the United States. There was another example in which Ron DeSantis... Remember Ron DeSantis' presidential campaign? When Ron DeSantis was running for president, he created AI images of Trump hugging Fauci as a way of targeting Trump and undermining his support among Republican primary voters.

But there have been a couple light touches in American politics, but it has been seen in really dramatic ways in Bangladesh, Slovakia, Pakistan, Argentina, India all over the world, to the point in India for example, you can find stories about this, there is just politicians all over the country are now using AI-generated deepfakes of themselves, of their opponents and so forth. Like in Slovakia for example, there was a particular candidate for the head of state was leading, and in the last four days before the election, a audio deepfake of him came out of him purporting to say that he was going to rig the vote and rig the election.

And in the end, a pro-Putin, anti-Ukraine candidate came from behind and won in the final days. And so we're seeing really serious examples like that. And then you're seeing really silly examples like the Trump-Fauci stuff. But I'm really worried about not stuff that just targets politicians, but for example, false audio of an elections official "caught on tape" saying that their voting machines can be hacked. A

fake robocall of Gavin Newsom going to millions of democratic voters on the night before election day, telling them their voting location has changed.

A fake version of an entire county elections website distributing false information to voters, all of what I just described is really easy to make using current technology and that technology is only getting better. So that's what we're worried about and that's what people are seeing in other countries. And I think all of that is really a dry run for what we're going to see in the United States. I think all of that is coming to the US in the next five or six months. And so if that's what you're supposed to be on the lookout for, what should people do about it? Well, the first is don't trust social media to do any content moderation that's going to solve this problem.

YouTube, Meta, X, they've all stopped labeling and removing posts that repeat Trump's false claims of a stolen 2020 election. Facebook has made fact-checking features optional on the platform. Twitter has stopped utilizing a software tool that actually can identify organized disinformation campaigns spreading across the platform. All of the platforms have laid off key members of their trust and safety and integrity teams, right? So the social media platforms are not going to save us. They are walking away from their responsibility to do something about this. The other thing is you can't trust politicians to save us either, because increasingly politicians are operating in an uncertain environment.

They may want to run honest campaigns in which they do not succumb to the allure of this new technology that makes dirty attacks and dirty campaigning so much easier. But if their opponent is doing it, they may see themselves in an arms race where they feel obligated to do it. So don't trust the social media platforms. Don't trust politicians. If you are a community member who wants to make sure you are not deceived, what I would recommend is find trusted news sources and go directly to those news sources for your news. If you trust the AP, go directly to the AP's website.

If you trust the New York Times, go to the New York Times website. Do not get your news from social media if you can help it. And when you do see something on social media that is too good to be true, it's a video that perfectly attacks the person you oppose for president, or it's an image that makes the person you support for mayor look perfect or something, it's very, very flattering to them. Or there's video that perfectly confirms some conspiracy theory, before you believe it, before you retweet it, before you share it, go and Google it, do some additional research. See if what you're seeing reported on Twitter or on Facebook is being reported by actual news sources in other places. So double-check it. If it's a real story, it will be reported in these other news sources that you trust. If it's a fake story, either it won't be reported in those other news sources or those news sources will be actively debunking it.

They'll say, "There is a video circulating of Vice President Kamala Harris saying X, Y and Z. That's a fake, don't trust it." So we're an era of double-checking when it comes to our news content now. We cannot trust our eyes on first glance. We need to do a little bit more research. The way I've put it is that we all need to have more finely calibrated BS meters in 2024 and going forward. And if you are someone who is already doing all these things, then it's your obligation to go talk to the young person in the life, the elderly people in your life. I think those are two populations that are particularly vulnerable. And just to your neighbors, your friends and your community members.

I think everyone needs to be engaged in this process of leveling up our skepticism. And so if you have an uncle or an aunt or a grandma or grandpa who is uncritically reposting what they're seeing on Facebook, have a conversation with them about how our information ecosystems are changing and how we all need to increase our skepticism.

Michelle Deutchman:

That was so terrific. I feel like you were able to crystallize something tangible that really everybody who has access to the internet and social media can do, which I think is so important in your reason, a much more existential question, which is in this moment, who can you trust? And what you're saying is that you need to be really not trusting in this moment, which is sad, but it's the reality. You kind of come at this

sort of tech world from a different vantage point as a voting rights attorney, as an attorney. And I'm curious how your background in the law impacts how you see all of this technology and generative AI and how it's impacting voting in our democracy.

Jonathan Mehta Stein:

Well, I'll say that most people who make tech policy or work on tech policy in one way or another, come out of tech. And I don't. I grew up in Silicon Valley. I grew up on the Cupertino-San Jose border. I went to elementary school that abutted the back parking lot of Apple's World Headquarters. So I grew up right in the heart of Silicon Valley. And instead of being part of it, I sort of watched as tech ate the Bay Area over time. And I think as a result of being around it but not being part of it, I don't have any deference towards it. And so my view is that when we seek platform accountability, when we seek regulation of tech to ensure that we're keeping kids safe and we're ensuring the mental health of our young people, that we're protecting older folks from scams, that we're protecting the integrity of our elections, when we take assertive aggressive action to regulate the tech sector in these ways, we are not doing something anomalous or extreme or overly aggressive.

The reality is if you can take a perspective from outside of tech, the reality is every other industry that impacts our wellness, our family's wellness, the wellness of our communities, every other industry is accustomed to regulation and oversight. So if you're an airline or you are an automobile manufacturer, or you are a food producer, or you're a pharmaceutical company, or if you work in financial services, you are accustomed to really significant regulation, testing of your products, inspection and oversight. And when you fail, there are consequences. When you have negative impacts on the community or on members of the public because of your products, there are significant impacts.

Tech has lived in a regulation-free and consequence-free world for now several decades, and we've seen the impacts of that. I mean, there is more than enough data now that shows that monumental impacts on teen mental health, for example, of social media platforms. And yet as a tech company or a social media platform, you can roll out a new feature or a new platform or a new app and it will be adopted by millions within days with no pre-testing, no sense that you have to check in with any sort of oversight body in order to make sure that this is going to serve people well and not harm people.

And so basically, I just think the rules of the game should be equal for everyone and that tech should be held to the same standard. We hold all other industries, and I think that by virtue of coming from a civil rights and civic engagement career instead of a career in tech, it's easier for me to see that with great clarity.

Michelle Deutchman:

I don't know if you're willing to tell us what that would look like for you. Would that look like sort of an FDA or some kind of oversight body? In your ideal world, if you could spin that or draw that out, what would it be?

Jonathan Mehta Stein:

Well, a lot of the leadership in this space is coming from the EU. And the team at CITED is very grateful to be in deep conversation with the EU regulators who are writing the AI Act and other relevant pieces of law that are coming out of the EU. And one of the approaches they're taking is that they're forcing the AI companies to do self-audits for potential harm to consumers and to report out on the impact of those audits or the findings of those audits, and then to do risk mitigation work. So if you are creating an AI chatbot system that can interact with someone as if they are that someone's friend or even family member or girlfriend, and there's a danger of creating weird parasocial relationships and unhealthy mental health dynamics and increasing isolation, for example, under emerging EU law, you're required to audit your systems to see if that's possible.

And if it is, you have to report out to the public. And then if those dangers are real, you have to do work to mitigate them in one way or another. And that seems like common sense to me. I would argue that that's less extreme than the regulatory infrastructure that pharmaceuticals have to deal with, or food producers have to deal with, or financial services companies in many respects have to deal with. We're just beginning to roll out a regulatory framework around tech, and I don't think anyone is trying to get overly aggressive about it. It's just putting in some common sense basics.

Michelle Deutchman:

And do you think that California is in a position to sort of lead on this?

Jonathan Mehta Stein:

Absolutely. I think that in fact, California has an obligation to lead on this. We can't count on Congress, we can't count on the platforms and the social media companies themselves. California has a legislature willing to be bold, and it has this enormous amount of tech policy expertise, including current and former tech executives who are now grappling with the social, moral and political consequences of the technology they've helped create. We're really lucky that at CITED, we have these advisory councils and these partners who are people who spent years inside Google, Amazon Meta, and they know how these systems work and they can help us find ways to build policy solutions that protect the public while not stifling innovation.

And we can go into a committee hearing in the state legislature and if a tech company says, or a tech trade association says, "This bill is incompatible with technology, this doesn't understand how the technology actually works, or this bill will stifle innovation," or whatever the case may be.

We have tech experts that are working with us, particularly because we're in California, who can say, "I spent years inside the industry. I spent years working for the largest companies. I know for a fact what we're proposing here is not impossible. It may not be great for your bottom line. It may take some work, it may be inconvenient, but it is possible and it will serve the public interest." And so California has all the right ingredients. I think the one challenge is that tech has very, very successfully cultivated allies in powerful positions and they've invested a lot of money into campaign donations and lobbying over time as any smart industry would. And so there is a built up cache and credibility and influence that they have at the highest levels. And so we just need to make sure that it doesn't overwhelm doing what's right.

Michelle Deutchman:

One of the things I think a lot about is the law, especially judicial decisions move much more slowly than what's happening on the ground. And technology seems similar, right? That it's moving way more rapidly than democratic institutions are able to evolve. And that's clearly, like you said, we're not going to catch up by November. What impact do you think that has in terms of voter behavior, that sort of lag time?

Jonathan Mehta Stein:

So that's totally correct. It is the case that these technologies evolve much more quickly, and these companies innovate much more quickly than government or civil society can keep up with. And as a result, historically I think government has been unwilling or unable to act. The courts have been unwilling or unable to act, because they don't feel comfortable or confident in legislating a rulemaking in the space, and they don't want to get it wrong because it feels like the technical expertise you'd have to have in order to get it right is so enormous.

But the reality is that so is flying planes, so is building vaccines and medicines. These are all things that are enormously complicated. And we've just put the work into building oversight mechanisms and regulatory infrastructure and we just have to do the work of doing that. We have to acknowledge that tech is as it shapes and controls our lives to the same extent that cars do, for example.

And cars are heavily regulated and there's all sorts of safety protocols in place. And we just have to build all of that same infrastructure around technology. And the good news is we're starting to do that. CITED is a good example of a new emerging piece of policymaking infrastructure in California. And also the sense that tech can't be regulated is evaporating. The veil is off. I think that people everywhere are understood... But of both parties. This is really bipartisan. People of both parties are recognizing that allowing for a totally unregulated social media universe for 20, 30 years starting with I think Section 230 is one of the original sins in the field. It has led to hugely negative consequences in all areas of our lives. And so we can do better on social media and we can get AI right in the first instance.

Michelle Deutchman:

Well, and you've named one of the very few things that seems like it could be bipartisan, which is amazing, right? So let's dig in a little bit deeper. Last season around this time, we had the director of the Center of Technology Policy at the University of North Carolina, Chapel Hill, Matt Perault on the podcast to talk about AI just as the world was waking up to ChatGPT and its use in the classroom. That was kind of our focus. And recently that center released a report on the new political ad machine policy frameworks for political ads in the age of ai, which laid out a couple of buckets. One is scale, one's authenticity, one is personalization, and one is bias. And we've touched on a couple of these things and I sort of want to let you address any of these four buckets and how CITED is thinking about them as we move forward.

Jonathan Mehta Stein:

Let's talk about scale. So in 2019, California passed the nation's first deepfake law, protecting politicians in particular from malicious deceptive deepfakes. And a number of states around the country are now beginning to replicate that 2019 California law. We think that the existing California law is inadequate and needs to be replaced, and the author of that, Assemblyman Mark Berman, is now part of CITD's legislative package to build something better and more nuanced. But here's why I think it doesn't work. Basically it says that if somebody posts a political deepfake targeting a politician under certain circumstances, that the state or some sort of enforcement entity has the ability to seek an enforcement action against the creator of that political deepfake. And it's narrowly tailored so that it stays on the right side of the First Amendment. And it's really time-limited, it's really cabined so that we're not creating First Amendment issues.

But the reality is, if I were to create, create a political deepfake of President Biden, and I put it on Twitter and it was wildly successful, it might take someone weeks if not months, to pursue some sort of enforcement action against me after which they successfully able to get me to take down the original post. By that time, the post has been seen by millions of people and it has probably been repurposed over and over and over by other users of social media. They've taken it, they've reworked it, they've memed it, they've turned it into all these different things. And by having me take down the original post long after its existence in the news cycle is probably over and long after it has been repurposed a million times and rememed a million times, you're playing a really ineffective game of whack-a-mole, and it just doesn't scale.

And so our view is that we need to put an obligation on the social media companies to be stopping this stuff from entering the information ecosystem on the front end, because backend solutions just don't scale. You can't go chasing one deceptive political deepfake after another, because you will be doing that millions and millions and millions of times. No one has the enforcement capacity or the time to do that, and everything you're trying to stop will have already gone viral before you get to it. And so we think the only way to do this at scale is to put an obligation on the social media companies to not allow the posting or to label the posting a very specific thing. So if you're going to ask the social media companies to take down a certain class of stuff or to label as fake a certain class of fake stuff, you have to give them really actionable direct guidelines that they can follow so it doesn't get messy.

They're not taking down way more stuff than they should be, and it has to be easily identifiable. It has to be stuff that they can readily identify, so they aren't having humans sitting somewhere checking millions and millions of posts in real time that they can automate this. So if you can get that right, then I think you can place the obligation to protect us from this deceptive political garbage, misinformation and disinformation. You can put it on the front end and it can actually work, it can actually clean up our information ecosystems. Do we have the magic bullets? I think we have some pretty good ideas, but we're at the very beginnings of trying to regulate in this space, and so we're doing the best we can and over time things will evolve and I'm sure they will improve.

Michelle Deutchman:

I like that hopeful note. You're a former ACLU attorney and I'm sure you can imagine the arguments that folks like ACLU and others who care very deeply and profoundly about the First Amendment are going to say about the potential for vagueness and ambiguity when you create this legislation. And again, I know there you might not have an exact answer, but I imagine you're anticipating that, and do you think it will be possible to craft something that is nearly tailored enough but also able to be enforced in a way that isn't going to run afoul of the First Amendment?

Jonathan Mehta Stein:

Absolutely. So first and foremost, I think there's lots of room for transparency-based solutions. So for example, I think what you can do is, and what we're looking at is watermarking and labeling requirements. So one of our bills would require that generative AI companies embed digital metadata in the digital media they create. So when they create a fake photo, a fake video, fake audio, it carries embedded metadata that allows us to determine that it is fake. It tells us when it was created, where it was created, how it was created and so forth. And that same bill requires that social media companies read that embedded metadata whenever fake AI content is posted and add a tag or a label under the post or near the post, that indicates that what you're seeing is fake. And there's no First Amendment concerns with simply labeling fake content as fake, so that we have a better understanding of what we're consuming as participants in the social media universe.

But I also think that bands of political defects that seek to undermine our elections and undermine our democracy are possible in a narrowly tailored way that respects the First Amendment. We have some bills that attempt to do this, and we are in regular conversation with the ACLU and the Brennan Center and a number of other friends who care passionately about the First Amendment and they're helping us shape this. Even in some cases through their opposition. We are in conversation with them and taking amendments in order to build the most effective legislation we can while staying on the right side of the First Amendment and respecting free speech concerns.

I think when you're attempting a ban of political speech, it's really important that the ban is not in any way based on the perspective of the speaker. It has to be banned based on the fact that the content is fake, that it is a deepfake of and narrowly tailored this as much as possible deepfake of candidates, elections officials or elected officials that are reasonably likely to deceive voters and to impact the outcome of elections or prevent people from voting. Fake by itself is not enough, but we think if you add these other elements and you time limit it so that this is what we try to do in our legislation, all of these bans, these very specific narrow bans, they're also time-limited so that they apply for a limited amount of time before election day.

And then for with respect to political deepfakes that target elections official or the integrity of elections, they also apply for a narrow period after election day to cover the vote counting period. We have a compelling state interest, which is that we want to protect our democracy in our elections and we have a narrowly tailored solution to protecting that state interest. And so are we trying to be bold? Are we entering some new territory here? Absolutely. But our hope is that following the clear guidelines of First

Amendment jurisprudence and in consultation with the First Amendment advocates, we can build legislation that is both impactful, but on the right side of the wall.

Michelle Deutchman:

I feel very heartened to know that you and CITED are at the helm and leading the way. I want to make sure that we do turn back to higher education a little bit, and I sort of have two questions about that. We've talked about some of the ways that generative AI and deep fakes are being used to ultimately suppress the vote, and I'm wondering if you've seen any evidence of that in particular with student voting or if you expect that we'll see more of that as we get closer to the election?

Jonathan Mehta Stein:

Unfortunately, there's plenty of data now that shows that young people are particularly bad at differentiating between real news and fake news. And I think I'm speculating here, but I suspect it's because if you are under 25 or so, you have grown up in a post-newspaper world. Forget the evening news and some verified stamped source of credible news that everybody believed in together. That's a product of decades ago now, but at least newspapers were seen as trusted institutions in your local community that everybody was... If they didn't read it, they were at least aware of it. And if there was news reported in your local newspaper, everyone knew they could trust it. With the decline of the local news industry, which is a different thing we work on at California Common Cause separate from CITED, but because of the decline of the local news industry and because of the loss newspapers, I think young people live in a world in which they have always gotten news from their social media app of choice at any given time.

Today it might be TikTok, and 10 years ago it might've been Instagram or Facebook. And I think that that leads to a really destabilized information environment and not knowing what you can trust. And universities I think have a really significant role in speaking to that. I think that one way is that they're in conversation with young people all the time, challenging them on where they get their news and how quickly they believe their news is a really useful role for universities to play and forcing people to think more critically, which I know the UC system definitely does, is a really useful enterprise. Questioning what you're seeing and what you're hearing is essential in today's world and a good university system is encouraging their students to do that.

But also in California, our leading universities are a source of a huge amount of brainpower that helps CITED and others in this field act with effectiveness and integrity and sensitivity. That is to say people who understand the technology really, really, really well, but are not working for tech. They are at the university and they are able to make unbiased recommendations. CITD Benefits from Berkeley Professor Hany Farid, who is maybe America's leading expert, an academic on the subject of deepfakes. Professor David Evan Harris at Haas is a close partner of ours who has helped us understand a million parts of this work because of his longtime service doing it from inside the tech industry before he became an academic and a policy person.

Michelle Deutchman:

We always usually end by asking our guests to talk to listeners about tangible things they can do, and you've already done such an exceptional job of covering that. I think what I'll just ask you then before we close is if there's anything else that you want to add about sort of the wide birth of things that we've been able to cover or maybe something we didn't cover.

Jonathan Mehta Stein:

I'm going to flag one additional short-term concern and one additional long-term concern. The short-term concern is that we may be on guard until election day and then put our guard down. And I think actually

the period between the election and the next January 6th is really important. If Trump loses, last time around, he had assertions that the vote was rigged and that votes were stolen and that elections officials in the deep state were conspiring against him. And he took all that to court and it ultimately didn't work. This time around, they will be able to make false but confirmatory evidence of anything they want to assert. So a deepfake of an elections official manipulating votes on a voting machine, they can make that deepfake. A deepfake of an activist, a liberal activist or someone they think looks like a liberal activist stealing vote by mail ballots out of a Dropbox.

They can make that. So in the vote-counting period before our presidential election is certified, it's a really acute period where the dangers are very high, and so we need to be on guard. The longer term concern I have is that in a destabilized information environment in which people don't know what they can trust, they may just simply disengage. They may say, "I thought politics was gross and now it's even grosser. Forget it. I'm just not going to vote, or I'm not going to participate, or I'm not going to pay attention." Or, I think even worse than that, maybe that people begin to believe anything that confirms their biases and beliefs and they disregard as false and fake, anything that challenges them or confronts their beliefs. And we retreat into tribalism. And so the long-term dangers here are disengagement or tribalism. Both are hugely problematic, and in a difficult moment for our democracy already either of those paths or both of those paths simultaneously could mean really dangerous things for our democracy and ultimately for our communities.

And so this is a moment where we need to collectively confront these challenges no matter how tired we are, no matter how cynical we are, no matter how pessimistic we are, no matter how turned off by politics we already are. Essentially, we need to work together in defense of truth. Even when that truth doesn't go our way or doesn't agree with us, we need to act in defense of the idea of collective decision-making. The idea of fact-based decision-making, separate and apart from partisan outcomes and political outcomes. And that's a challenge I don't remember the United States having to face in my lifetime. And so it's new ground for us in many ways, but we have to figure it out because the stakes are too high.

Michelle Deutchman:

That's our call to action, and hopefully all of us individually as members of communities and as members of higher education can take steps to meet that. So obviously, you are very busy trying to solve all of these problems, so we're really grateful that you could give us your time and your insights and your expertise.

Jonathan Mehta Stein:

Thank you so much for having me. I really appreciate it. And thank you so much for this podcast. I love it.

Michelle Deutchman:

Thanks for joining us for another thought-provoking episode. Next month we plan to focus on the supreme Court's outstanding decisions, which are traditionally issued by June 30th. In the meantime, enjoy some of that summertime weather.