

## TRANSCRIPT Episode 13 – Zach Boyd

**Jonathan Bench:** Welcome to another episode of *Lawbalization*. Today I'm joined by Dr. Zach Boyd, the Director of the Office of Artificial Intelligence for the State of Utah, where he leads development of AI regulatory policy and helps clear paths for novel business applications of AI. He's on long-term leave from his role as a professor at BYU, where he teaches applied and computational mathematics. His research lab focuses on artificial intelligence, machine learning, and mathematical modeling in social science applications such as psychology, economics, and social networks. Zach, thanks so much for taking the time to talk with me today.

**Zach Boyd:** Thank you—it's a pleasure.

**Jonathan:** By the way—this will air after the fact—but we got approval to sponsor the AI conference in December. I was already planning to attend, and I'm glad our firm can support it.

**Zach:** Excellent. Thank you for your support.

**Jonathan:** I attended last year as well. It was great to see the Governor involved, and I'm sure we'll talk about how he's led this and continues to be business-minded to keep the state thriving.

Let's start with you. Give us the highlights of your career leading up to where you are now. As someone in mathematics, did you expect AI to roll out the way it has? Were you seeing the signs over the last decade before it hit the mainstream?

**Zach:** My career path has been academic. I was at BYU, then completed a PhD at UCLA. I did work at Los Alamos National Laboratory in the weapons division, and then at the University of North Carolina at Chapel Hill.

My research sits at the intersection of mathematical modeling of complex systems, machine learning, and social science—things like computational psychology and supply-chain management. It's "wet and gooey," but mathematically intense.

We've had machine learning for a long time. Around the time I started grad school, deep learning started showing remarkable progress—especially in computer vision, effectively teaching computers to "see." There were rumblings about scaling laws and the coming of powerful language models. As they improved, coherence increased.

The twist with scaling laws is you don't really know what "one level up" means until you see it. GPT-2 was somewhat coherent; GPT-3 was a level up and good; GPT-4 was fantastic. So even for people watching closely, the leap felt sudden—over just a couple of years.

**Jonathan:** That can still feel fast. I saw something similar in blockchain—thinking I was late, yet everyone kept saying we’re still early. It’s interesting how something can be new, still novel, and yet already have years of traction. We tend to think in days and minutes rather than decades.

**Zach:** Exactly. Right before the reasoning-focused models came out, I was in the Bay Area talking to folks very bullish on AI progress. I asked if it felt like things were slowing down. They said, “Look at the scaling-law trends—we’re right on track.” Six months can feel long without visible progress, but then the next day the reasoning models arrived—boom—up another level. Humans perceive time differently from the actual progression curve.

**Jonathan:** I’m in a field that doesn’t adapt quickly. Lawyers keep pace with legal changes, but not always how technology will impact workflows. Every industry is wondering: Have we adopted AI fast enough to stay competitive? How quickly can we adapt?

**Zach:** It’s hard to predict socio-technical adoption. So much modern tech still isn’t widely adopted. Some financial institutions run code written in COBOL—decades old. It’s tough to know which advances will be critical to stay ahead of and which society will be slow to adopt.

**Jonathan:** Let’s backtrack. You’re the first person I’ve met who spent time in Los Alamos. What happens there, and what can you tell us about your work? Why is the lab world-famous?

**Zach:** Los Alamos is famous as the birthplace of nuclear-weapons science—the Manhattan Project. Since then, it’s been one of the chief scientific centers for nuclear weapons technology. I spent three summers there as a graduate student in the weapons division.

Because of international agreements, the U.S. doesn’t test nuclear weapons, but they remain central to national security. The challenge is like leaving a car in a garage for 30 years and needing to guarantee it will start and run exactly as expected. We worked on physics and simulation to provide assurance to Congress and the defense community that the systems will operate precisely as intended—even with component changes over time.

**Jonathan:** Fascinating. Let’s come back to your current role. Tell us about the Office of AI—Utah tends to be lean but effective. What’s your office like?

**Zach:** Our mission is super ambitious, and our team is small—basically five “superheroes.” We have a mathematician from Dartmouth for technical expertise, a legislative drafter from Congress with 20 years’ experience, a master connector and serial entrepreneur—just

incredibly high-caliber people. It's a tiny team facing massive challenges as every sector adjusts to new technological realities.

**Jonathan:** What is your office specifically tasked with?

**Zach:** Two convergent missions. First, ensure Utah has the best regulatory environment in the world for AI in business—the pro-business angle. Second, be more intentional than we were with earlier tech waves on the consumer-protection side. Many people regret aspects of how social media and smartphones evolved. We want to help the legislature make the best decisions with AI.

**Jonathan:** That resonates with my experience with the Division of Consumer Protection—small but collaborative, business-minded, and ensuring consumers have redress against bad actors.

**Zach:** Consumer Protection has been a great partner. The trick is to catch the bad actors without ensnaring the good ones.

**Jonathan:** How do you strike the balance between supporting the business community with certainty and ensuring guardrails?

**Zach:** We use a few frameworks. One is the difference between regulated and unregulated sectors. In regulated sectors—education, financial services, law—there are lots of rules. AI is disruptive, but the rules aren't. Businesses prefer clarity over relying on enforcement as a proxy for clear regulation. Some rules may be unnecessary now; some should be modified; some should stay.

In unregulated sectors, we think like economists: What will the market equilibrate to in the long run, and will we be happy with that? For example, consumer chatbots might be fine with competition and quality incentives. But with AI “companions,” early data suggests market incentives to isolate people, replacing human relationships with synthetic ones—a potential market failure. Then we ask: What's the lightest, most effective intervention to avoid that failure while letting the market serve people well?

**Jonathan:** How is Utah received in other circles? Do people want to duplicate what you're doing, or say it's not enough?

**Zach:** There's a range of opinions, but generally Utah is seen as a competent, sensible leader. Several states have reached out to learn how to replicate our approach and how we stay on top of the issue without being overbearing. In some places, regulation demands that industry solve deep technical problems without any practical path. Utah is perceived as practical—recognizing the need for solutions but staying common-sense about getting there.

**Jonathan:** With a team of five, what does a good public-private partnership look like? How should the business community use AI responsibly to help the state lead without accidental missteps?

**Zach:** The state's priorities reflect the people's: we want businesses to succeed and feel permission to experiment and innovate—responsibly. One tool is our regulatory sandbox. We recognize existing rules, but some lose relevance as tech advances. Companies come to us with innovative approaches that may conflict with current regulations yet are justified by technological shifts. We can authorize limited-period pilots in exchange for data sharing and adherence to a responsible-AI policy aligned with common-sense principles. The state learns without delaying innovation for years.

**Jonathan:** Which sectors have been early with the sandbox?

**Zach:** Healthcare, because professional licensure and conduct requirements are evolving as AI changes scopes of practice—what to do personally versus delegate to a machine or a lower-skilled professional.

In education, our role is more facilitative—helping stakeholders understand issues so local leaders can implement according to community values.

In financial services, things are slower because of joint state-federal oversight; it takes time to advance together.

**Jonathan:** We're at year-end as we record, when a lot happens legislatively. Your office is only about a year and a half old—still early. In the EU, the approach is different: more taxonomy and ex ante certainty versus the U.S. tendency to let things run and adjust. It's not right or wrong—just different legal philosophies.

**Zach:** Agreed. Regulators are used to categories and bright-line rules, which is hard in a rapidly evolving tech space. We're mindful that private litigation can address many harms, and businesses respond to that.

Another tool: future-proofing regulations. We can't constantly revisit everything, so we often include a safe harbor—a way to “check the box” without binding everyone to a specific method that might be outdated in a year. The safe harbor provides a clear path now while allowing other compliant methods to emerge as the tech evolves.

**Jonathan:** Can you give an example?

**Zach:** Last year we passed a mental-health-related bill. The state licenses mental-health therapists, but an AI that presents as a therapist breaks categories. The law clarified that the Department of Professional Licensing can pursue AIs that harm mental health—paralleling human accountability—because consumer protection matters.

We added a safe harbor: if, prior to deployment, a company creates internal documentation consistent with the state of the art—pre-deployment safety testing, a clinical advisory board, measured risk showing the system is at least as safe as a human therapist for the same task—then the regulator will defer enforcement. Document your safety and we'll leave you be for now, while leaving open how “state of the art” documentation evolves.

**Jonathan:** That makes sense: protect consumers, keep a human ultimately responsible where appropriate, and let innovators move forward if they take reasonable steps. The safe harbor helps avoid the gray-area paralysis.

**Zach:** Exactly. Some states, like Illinois and Nevada in similar areas, have simply blocked activity because they don't know how to regulate it yet. Utah's approach isn't to block—it's to use reasonableness while we learn, then refine when checklists actually work.

**Jonathan:** If you could wave a wand and remove one barrier preventing Utah businesses from deploying AI, what would it be?

**Zach:** It would be sociological: people still don't know how and when to trust AI. For example, AI call-center agents can outperform available human staff in testing, but callers dislike them. We're collectively figuring out how we want these systems in our lives. I'd turn the time dial forward a few years so people are more accustomed to the tech.

**Jonathan:** Curve ball: How do you personally use AI right now—favorite use cases at work or home?

**Zach:** Outside of research, I use the main language models for analysis and drafting. This morning I was debating a legal point with counsel and used AI to gather arguments. I love high-quality voice modes: if I'm driving to a meeting on, say, Medicaid law, I can get a 20-minute briefing. It makes me smarter walking into meetings. I'm not a “fully AI-infused life” person, but I use it where it adds clear value.

**Jonathan:** Same here. For non-billable prep, I won't spend two hours before a call that may not convert. AI lets me digest a slide deck or white paper in ten minutes so I can be responsive without over-investing.

Quick vignette: I run an AI-enabled Dungeons & Dragons group to build deeper connections in Utah's AI community. Someone recorded our first four-hour session (Utah is a one-party consent state), ran it through models, and generated character arcs. It was better than we imagined. We've kept iterating—world-building, characters—mixing casual users and super-users. It's been a blast.

**Zach:** That’s actually been motivating. We hear about AI for D&D in our internal build discussions. On version 80 of some definition, we’ll say, “Make sure the D&D use case still works.” It’s helped push our definitions.

**Jonathan:** We plan to expand it for Utah Tech Week in early 2026—probably up to 50 people. You’re welcome to drop by, even if you’ve never played—grab some food and see what the AI ruckus is about.

**Zach:** Love it.

**Jonathan:** A couple questions left. What measures are most critical to ensure AI strengthens public trust rather than undermines it? How do we help people see AI as an exceptional tool for good—like the internet—if used correctly?

**Zach:** Three factors:

1. Technical reliability. AI will make mistakes; that won’t vanish, but we can reduce frequency and impact.
2. Sociology of use. People need to manage AI like a team—clear instructions, incentives, feedback, and quality control—more like managing a person than buying a product.
3. Business incentives. In B2B contexts, incentives often align better. In consumer contexts, I worry about personalized data sales and engagement mechanics harming certain groups. We’re working to shape markets so incentives align with long-term community interests.

**Jonathan:** Fraud is a concern—robocalls on steroids. I have hundreds of hours of recorded voice out there; I could be deepfaked quickly. Even routine transactions—voice prints, passcodes—can be spoofed. We need to keep up with the latest scams and slow down when needed.

**Zach:** Totally. I’m optimistic we’ll see better usability even without dramatic capability jumps: stronger digital authentication to verify identity against deepfakes, and better agentic integration across systems. The next few years should be exciting.

**Jonathan:** I’m with you—optimistic, with education and awareness. Looking ahead: say you spend a few more years, then hand the baton to someone else. What does Utah’s AI landscape look like in five to ten years, and what do you hope people say about your office’s early days?

**Zach:** I hope people say we laid pioneering foundations for flourishing. I admire the generation that built the internet—open protocols, cryptography, a connected world. I’d like

to be that for AI, not a repeat of early social-media pitfalls. Let's build the groundwork for decades of progress.

**Jonathan:** Last question. If you weren't at BYU and not at the Office of AI Policy, what would you be doing?

**Zach:** I'd likely be at a major AI lab, trying to improve capabilities—or, as a mathematician, digging into why these systems work. We've created working emulations of human behavior. I'd love to be in the guts, understanding mechanisms and what, if anything, they reveal about us.

**Jonathan:** Thank you, Zach. It's been amazing having you on the show. I'm looking forward to the AI conference and hope we can catch up again to see where things have gone since this conversation.

**Zach:** It's been a pleasure. Thanks.

**Jonathan:** Thanks.