

TRANSCRIPT Episode 22 – Medlir Mema

Jonathan Bench: Today I'm joined by Medlir Mema, the Director of AI Ethics and Governance Institute and a Senior Fellow at Faith Matters. He is also a Professor of International Relations at Brigham Young University, Idaho in the United States. Medlir holds a PhD in Political Science from the George Washington University and an MA in European Studies and International Economics from Johns Hopkins University School of Advanced International Studies. His work explores how AI systems shape democratic processes and public trust, as well as the implications of AI-driven automation, algorithmic decision-making, and concentration of technological power for economic systems and public institutions, topics which he covers with invited guests on his Age of AI podcast. Medlir, thank you so much for joining me today.

Medlir Mema: Well, thanks, Jonathan. And I'm not quite sure how I ended up as your guest. I think this started as my idea of having you as a guest on my podcast. And I guess for my listeners on my podcast, this will also double up, right? So this is a collaboration I suppose we'll both have.

Jonathan: That's right. This is called economies of scale.

Medlir: That's exactly right.

Jonathan: I'm very excited. So give us a little bit of your personal and professional background. You're originally from Albania. You're now in the United States. You have advanced degrees in the U.S. and you've leaned very heavily into AI and governance. So maybe give us a quick run through of how a kid from Albania ended up in the United States doing what you're doing and traveling around the world and really leading a lot of the conversations around what AI governance should look like.

Medlir: Yeah, good question. You mentioned my background from Albania, and that actually plays a role in my journey in terms of getting interested in AI. I came here when I was 18 to go to school, as you noted, and I stayed here throughout that education experience all the way through my PhD. About halfway through that, I went back to Europe. So about 15 years ago, I went back to Europe, first Sweden, Belgium, and then Albania. I was involved in political campaigns and conversations and was thinking a lot about development issues. How can countries like Albania, but others that are struggling with catching up with the rest of the developed world, make use of technologies to leapfrog? That led me into the discussion around the fourth industrial revolution, which is now the conversation about AI, big data, and emerging technologies. Initially it was all positive. I was very optimistic, and I still am optimistic about technology. But increasingly I began to think more about issues of labor displacement, education, and the upsides and downsides

of these technologies. As a political scientist coming from a social science perspective, not a technical one, I started asking what these technologies do to societies, especially democratic societies. When you have massive job displacement or educational arrangements being reshaped in this fourth industrial revolution, what does that do to trust in society? What does it do to citizens' relationships with governments?

Once you identify the potential harms and benefits, you have to think about mitigation and alternatives. That's where my work on AI governance comes in. Over the last three or four years, I've worked with colleagues on what AI governance should look like at the global level. More recently, especially since coming back to the United States, I've focused more on what that looks like domestically. That's basically the journey. And if I may, Jonathan, let me ask you the same question.

Jonathan: I consider myself a technology native. I was four or five years old when we got our first computer, so I'm part of the last generation in the United States that grew up pre-technology and then lived through the technological revolution from computers not connected to the internet, to dial-up, to broadband, and now where we are today. When I went into law, it was an interesting path because law school really teaches you how to function in the adversarial legal system. There are classes on business, but they don't really teach you how to be a business lawyer.

I leaned into business law because of my background in China and Chinese. I didn't study coding or computer science in high school or college, so I assumed I had missed the technology boat. But I've always been on the edge of the tech curve without being a true techie. My tech friends call me an honorary techie because I don't have the degree, but I love learning how technology works. I was working in command prompts before true UI and UX experiences, which gives me a little bit of credibility.

As I moved through law school and into practice, I started reading about blockchain and other emerging technologies that had previously seemed like science fiction. I assumed I wasn't the right lawyer for that space because I didn't have a technical degree or an IP background, even though I went to George Washington University Law School where many classmates were deep into IP. But mentors encouraged me to learn blockchain because it was international and technically deep. I realized that one of the core legal skills is not being afraid of new frontiers. So I kept learning. Eventually, like much of the world, I moved from blockchain into AI. We still have blockchain fundamentals and use cases, but AI has been a tsunami that forced all of us to reckon with it.

As a business lawyer, my clients are all using AI to some degree or feel like they should be. It helps that I've leaned into understanding how AI works, how it changes workflows, legal

relationships, and the availability of information. I've spent less time thinking about governance and more time thinking about practical deployment. If I'm candid, I tend to be a realist. I say let the technology roll and see what happens. I don't think about potential harm until I start hearing concrete use cases, and then I think maybe we need guardrails to protect the most vulnerable.

Medlir: I think you're spot on in saying regulation shouldn't exist for regulation's sake. That shouldn't be the objective. But at the same time, AI models are among the least regulated technologies in the United States. Around the world, there are layered regulatory frameworks at global, regional, and national levels. China, which you know well, has multiple layers of regulation. I don't think the United States should imitate China, but leaving trillion-dollar technologies without any regulatory framework is deeply problematic.

As evidence of misuse or harmful outcomes increases, not hypothetical long-term existential threats but real harms like teenagers committing suicide after interactions with AI systems, we have to take it seriously. In the United States, states have taken the lead, and sooner or later the federal government will have to act. The question isn't whether to regulate automatically. It's what to regulate and to what end. There seems to be consensus around child safety and transparency. Those are areas where rules would make sense.

Globally, we've long said China has the data, the United States has the technology, and Europe has regulation. Europe's regulatory model has historically had influence through what's called the Brussels effect. With GDPR, for example, global companies adopted European standards worldwide. The expectation was that the EU AI Act would do the same. But there has been resistance. Even before the current US administration, there was pushback from both American policymakers and big tech. Some European countries like France and Germany began to hesitate.

We've seen retrenchment in Europe. The EU AI Act itself is under scrutiny, with some provisions delayed beyond 2026. There's even talk of digital ombudsman legislation. The language shift from the AI Safety Summit in the UK to the AI Action Summit in Paris signals a change in tone. If Europe feels pressure to soften its regulatory stance, you can imagine the broader global trend.

In the United States, while states experiment with regulation, there is pressure at the federal level for preemption. I don't think we'll see a global AI compact anytime soon. We'll see standards, benchmarks, civil society indexes, and voluntary initiatives, but not a binding global framework. That creates short-term gains for big tech but long-term uncertainty. Large companies can manage uncertainty. Small companies struggle with it.

Jonathan: That uncertainty ties directly into competitiveness. At both the national and company level, everyone is thinking about staying competitive. AI feels like a massive accelerant. Unlike blockchain, which had important but narrower applications, AI has broad applicability across industries and countries. Anyone with a phone can access powerful models. The potential productivity gains are staggering, and that has implications for national growth and organizational structure.

How should governments think about governance in a way that supports adoption without sacrificing competitiveness?

Medlir: There's a strong narrative that AI is a universal accelerant, and I agree to a degree. But we need realism. We risk technological determinism, assuming that technology inevitably improves everything. Many promised breakthroughs, such as dramatic transformations in healthcare, haven't materialized at the scale predicted. Much recent GDP growth tied to AI has come from investment in infrastructure and models, not necessarily transformative productivity gains across the economy.

We could have invested those resources into climate change, cancer research, or other pressing issues. Instead, we're doubling down on AI because we assume it will deliver outsized returns. I'm not saying AI lacks value, but we should question whether everything it touches turns to gold. Some of it may be diluted.

Jonathan: I've heard similar skepticism. An MIT study suggested that only 5 percent of enterprise AI investments translated into real productivity gains. In my practice, I think about billable and non-billable workflows. AI has significantly reduced time on non-billable tasks. For example, I used AI to help design a D and D campaign for my son, shrinking a 10-hour task into two or three hours. That translated into meaningful quality-of-life gains. I've also used AI to build elaborate experiences for my children that would otherwise have been impossible given time constraints.

Those are real efficiencies, but I recognize the broader question: are those marginal gains worth the trade-offs?

Medlir: That's exactly the issue. There's no free lunch. The convenience and time savings are real, but what are we giving up? In the absence of strong regulation, we are often trading privacy and data rights for marginal efficiency gains. Is that a bargain we want? Could we design frameworks that allow efficiency without total surrender of privacy or intellectual property? Governance should ensure that AI benefits individuals and companies, not just big tech. Right now, most financial gains accrue to large technology firms, while users reap incremental productivity improvements.

Jonathan: If I'm cynical, I might say that many of us have already given away our data through years of using big tech platforms. AI may not change the fundamental bargain. But perhaps we can at least build better systems for the next generation.

Medlir: Intergenerationally, we owe it to the next generation. Even selfishly, individuals should have rights to delete or monetize their data. From a business perspective, regulatory ambiguity is dangerous. Uncertainty can be worse than regulation. It creates compliance chaos and deters investment, especially for small companies. States as laboratories of democracy can test models. Ultimately, we need a predictable framework that balances child safety, transparency, whistleblower protections, and innovation.

Jonathan: Should we focus on AI-specific legislation, or adapt existing frameworks like data privacy laws to account for AI?

Medlir: We need both. Many laws were written before 2022 and are not AI-proof. Areas like human trafficking are directly impacted by AI through sextortion and exploitation. Existing frameworks need updating. At the same time, risk-based approaches like the EU AI Act or some state models can differentiate between high-risk and low-risk uses. High-risk applications may warrant bans or close oversight, while lower-risk uses can remain flexible.

Jonathan: As we close, what would be the biggest policy mistakes governments could make in the next five years?

Medlir: Regulatory uncertainty is a major risk. Businesses might even consider voluntary standards or certifications to build trust with consumers. There is a trust gap between producers and users of AI products. Voluntary seals of approval or industry standards could reduce pressure for heavy-handed legislation. Eventually, thoughtful federal legislation focused on child safety, transparency, and predictable guardrails, without overreach, would provide stability.

Jonathan: This has been a fascinating discussion. It's helpful to see these issues from a governance perspective and to think more deeply about the trade-offs behind technologies we adopt so quickly.

Medlir: I've enjoyed it as well. Hearing your practical, ground-level perspective from working with clients is invaluable. Governance debates benefit from understanding how these technologies are actually being used. Thank you for the conversation.