Future Skills Centre | Centre des Compétences futures

Presented by Signal49 RESEARCH

The Future Skills Podcast

Season 6: Episode 4

Building Canada's AI Advantage

In this episode of the Future Skills Podcast, we explore what it will take for Canada to build and sustain a competitive edge in artificial intelligence. Host Jeremy Strachan speaks with three experts who each bring a different perspective on the opportunities and risks shaping Canada's AI future. We begin with Aaron Rodericks, Head of Trust and Security at Bluesky, who offers a candid assessment of the global AI race—highlighting Canada's position relative to major players, the realities of market concentration, and the emerging cybersecurity risks tied to rapidly advancing models. Next, Jaxson Khan, Senior Fellow at the Munk School of Global Affairs joins the podcast to unpack the concept of AI sovereignty, outlining what it means in practice across the AI stack and what structural and cultural shifts are needed for Canada to build, adopt, and retain AI capacity. Finally, CIFAR AI Chair at the University of Alberta Blair Attard-Frost reflects on Canada's newly released AI strategy, emphasizing the importance of public trust, social research, and a broader cultural shift in how we approach innovation, governance, and talent retention. Together, this conversation explores what it will take not only to compete in the global AI landscape—but to do so in a way that is secure, sovereign, and socially grounded.

Guests

Aaron Rodericks, Head of Trust and Security, Bluesky
Jaxson Khan, Senior Fellow, Munk School of Global Affairs
Blair Attard-Frost, CIFAR AI Chair and Assistant Professor, University of Alberta

Host

Jeremy Strachan, Manager, Education & Skills, Signal49 Research

Links

Future Skills Centre Homepage: https://fsc-ccf.ca/

Future Skills Centre LinkedIn: https://www.linkedin.com/company/fsc-ccf

Future Skills Centre Bluesky: https://bsky.app/profile/fsc-ccf.bsky.social

Signal49 Research Homepage: https://www.signal49.ca/

Signal49 Research Facebook: https://www.facebook.com/Signal49Research

Signal49 Research X: https://x.com/S49Research

Conference Board of Canada, The. *Canada's Workforce in Transition: How AI Is Shaping the Future of Work*. Ottawa: The Conference Board of Canada, 2026. English | French

Signal49 Research: *AI and Entry Level Demand*. Ottawa: Signal49 Research, 2026. English | French

Government of Canada. *Canada's National AI Strategy: AI for All.* Ottawa: Government of Canada. English | French

Mullen, Sean, and Jaxson Khan. *Sovereign By Design: Strategic Options for Canadian AI Sovereignty*. Toronto: Munk School of Global Affairs and Public Policy, 2026. English

Loertscher, Oliver. *Going All in on Automation: Economic Growth, Labour Markets, and the Path Toward a Robust Transition.* Ottawa: Signal49 Research, 2026. English | French

Transcript

Jeremy Strachan: Welcome to Season Six of the Future Skills Podcast, brought to you by the Future Skills Centre. I'm Jeremy Strachan, Manager of Education & Skills at Signal49 Research, and your host for the season. On The Future Skills Podcast, we explore what matters most to Canadians when it comes to skills, training, and the ever-changing world of work. Since 2019, the Future Skills Centre has been driving Canada's workforce transformation by funding innovative training solutions, cutting edge research and inclusive partnerships to ensure everyone has the skills to thrive in a changing economy. In this episode, we dig into what it's going to take for Canada to build its AI advantage. From global competition and national security to workforce transformation and public trust, AI is raising urgent questions about how Canada positions itself in a rapidly evolving landscape. Aaron Rodericks of Bluesky starts us off with a global view—what the AI race looks like today and where Canada fits within it. Then we'll zoom in on sovereignty with Jaxson Khan of the Munk School of Global Affairs who tells us what it means to have control over AI systems and infrastructure. Finally, Blair Attard Frost, CIFAR AI Chair at the University of Alberta joins to look at how policy, research, and public trust must evolve if Canada is going to build a sustainable and inclusive AI future.

AI has indelibly transformed how we produce, and how we compete—bringing with it enormous opportunity, but also a high degree of uncertainty. And in Canada, that conversation has taken on a new urgency following the release of the federal government's updated National AI Strategy, AI for All, in June of 2026. The strategy lays out an ambitious vision for expanding adoption, strengthening research, and ensuring AI systems are developed and used responsibly. But as the global pace of AI development accelerates, the big question is how Canada turns that vision into real outcomes.

Despite being an early leader in AI research, Canada now finds itself in a much more competitive and fast-moving landscape. Businesses and public institutions have adopted AI unevenly, with many organizations still figuring out how to integrate AI into their day-to-day operations. At the same time, public trust hasn't kept pace with the technology itself and concerns around job impacts, data use, and the broader social effects of AI continue to shape how Canadians think about—and whether we're ready to embrace—these tools.

My first guest is Aaron Rodericks, Head of Trust and Public Safety at Bluesky. Aaron brings more than a decade of experience working on issues of information integrity, online safety, and platform governance, including previously co-leading Trust and Safety at Twitter. From his vantage point, Aaron has a unique perspective on the intersection of AI, platform power, and global competition—and what all of that means for Canada's place in the rapidly shifting AI landscape. Aaron, welcome to the podcast.

Aaron Rodericks: Thanks, Jeremy. Great to be here.

Jeremy: Help us contextualize where we are in the global AI race—and how Canada compares to major players when it comes to things like market scale, investment, and overall positioning?

Aaron: I looked up a couple of numbers right before this call. We're looking at a global total market for AI that's roughly 4.8 trillion by about 2033. That's an insane amount. Let's look at Anthropic just in the last year. They're now running at a revenue run rate of something like 30 billion. And I took Canada's Cohere and it had an incredible year, we're very proud of it, at 240 million, which is significantly higher than the 100 million they made the year before. But certainly when we look at, Canada-US markets and start to look at the AI one, we start to see that we're really a drop in the bucket. And when people are really charging to grab compute, to grab GPUs, to grab all the available capacity to launch the latest models—and at the end of the day, what you're trying is for a market capture here. And if you don't get market capture, we are far more likely from a Canadian position to end up in that small incremental long tail of per cents of the total addressable market.

Jeremy: You touched on just how concentrated this space is becoming. So let's turn to regulation. We're seeing jurisdictions around the world moving quickly—often borrowing from one another. And I want to ask what would a smart, effective Canadian approach to AI regulation look like in a landscape that's moving this fast?

Aaron: And certainly at the speed that things are going at, you really don't know what you're regulating. And it's very easy to get regulation very, very badly and very wrong. And this happens all the time. I'm in the social media space. I have been there for a while and we've seen that surge of regulation. Everyone's doing copy-paste from "hold my beer" versions to significantly worse versions. And Canada has some incredible failures in that regard. They tried to go after Meta on the news side, and Meta just had so much market power, it said, okay, fine, I'm just not going to serve news. And the sad thing is everyone loses when that happens. I can't send news articles to family in Canada. And when you look at AI, you're going to have the exact same position in that the bigger players dictate the market space from a regulatory perspective. You're seeing that in the UK at the moment. The UK – and I find a lot of countries, want to say, "Hey, we're gonna be a hub for innovation. Everyone's gonna choose us for AI. We are gonna seize a chunk of this market." On the other side, they still want to go, "No no, we're gonna slow it down. We're gonna regulate it. We're gonna be the safest AI market on the planet." You know, and you can't do both at the same time. You can only win one of them, and usually you're gonna win safety by not having anyone in your market. Or you're gonna win the innovation side by lowering the safety side. I wish I could be more pragmatic or tell you that there was a better story to win at both simultaneously, but as long as Canada keeps on trying to do both at the same time, it's not gonna win at either.

Jeremy: Ok, so picking up on the safety piece: From a national security perspective, how should we be thinking about these risks?

Aaron: Yeah, and it's a great time to have this conversation because we've had Anthropic start to talk about its Mythos model that was so powerful, it could basically take apart every operating system, every browser, and find zero-day exploits in them. That's insane. And it was so risky that they basically decided, no, we're going to have to hold off on releasing this. Now, in the US, you have the Fed, Jerome Powell, they're calling emergency meetings with banks saying "You have to get your cybersecurity position in place." Anthropic's trying to do the responsible thing, and it's doing kind of a pre-release to players who need to batten up their security. So let's look at all this happening in the context right now, that we are a couple of

months out from a release that could totally change the shape of cybersecurity as we know it. Where's Canada in that picture? Is Canada arguing for pre-release access? Is the government going to get that access? What are they going to do with that access and capability? So the state of Canada's software space at a federal government level, but also institutional level—look at the big banks, look at this—are they getting access to these models to pressure test their systems inside? And I think the sad answer probably no. So until that happens, Canada is very, very much at risk at not only being left out of the race, but then being left out of the people who get early access to build up their defenses against the adversaries who will then have this tool when it goes live to market.

And look, fundamentally, we're talking every browser—Firefox, Chrome—you know, we're talking your bank accounts, this, that—it can find vulnerabilities in the system because the way that these models are finding vulnerabilities is not necessarily an intuitive way, and then they're just able to go at it endlessly. Now, of course, Anthropic's got significantly more compute than your average actor. But once you get advanced persistent threats, you know, you're talking about the Chinas, the Russias, the paid hacker groups, they have a ton of funds to put towards, "Hey, I want to crack this particular organization. I want to find vulnerabilities in this place," et cetera. And essentially saying that where there's a will, there's a way. As long as you get access to those models, you're the first one with the keys. It's like almost at this point, like having an infinite amount of zero-days.

Jeremy: Ok I just need to jump in: What's a zero-day?

Aaron: So in the cybersecurity world, sometimes they'll find issues like, "Hey, if you press this button 17 times, a glitch happens and then you can put in a code." That's just one example. And what happens in the cybersecurity world is you want to find these vulnerabilities earlier and earlier and earlier so the companies don't catch on. So when you're in the bad guy world, finding a zero day, which basically means no one else has found it, it is at the earliest point of inception, it means you have a long play time to play with that to then get it out there. And Canada, it, is not ready at all for this.

Jeremy: Ok, so let's bring this down to the organizational level, what should Canadian businesses and institutions be doing right now to both protect themselves and take advantage of AI as a productivity and a potential workforce amplifier?

Aaron: Moving into this modern age of AI: I, for example, built an AI agent for our latest AI product in about 48 hours. I did possibly close to a month's work on doing massive amounts of research, then turning it into a harms taxonomy; then distilling it into the framework; then kind of the prompt areas where we want to focus on mitigating the harms, et cetera. I was able to distil all that down into a small, small amount. And I am *terrible* at using AI, relatively speaking. So I think when we're stepping into the future, what Canada needs to do is they need to, at least in these bigger organizations, figure out how do you do your sandbox, how do you do your carve out so that employees—how do you give them credits so they can adopt it? Meta went a little too far where your performance management is linked to your use of AI. So some people are just coming up with ridiculous AI uses just to spend the credits. But in the Canadian example, you are going to shift to a model of: There's no need for a manager to manage other managers to manage other managers. You are going to the manager IC or the "individual

contributor" model. It's basically going to be a fusion of that. You need to figure out how do you manage agentic AI moving forward and at some point managing other humans that then manage agentic AI. And I think this has massive consequences in the Canadian economy and on organizations who have done things—let's look at interns. I'm getting an intern for a summer. They're coming in not knowing anything, being like, "Train me, show me what to do." And then I have to train them how to do a briefing note, how to do this, and then going cycles. Unfortunately, I've been a manager for like 20 years now. So having had that experience, and then compared to an agentic model where I just train it on a skill, add that to memory, boom, it's learned it, keep on refining it, keep going from there—to a certain extent, why would I go through the pain of training up an intern for them only to move on or potentially we get them an offer, et cetera, when I can just have a model that does that for me? And when I actually get a little bit more talented, I can have it automated, running on a schedule, et cetera, et cetera. And these are the types of skills that more businesses need to adopt and more workers need to adopt to have that force amplification.

Jeremy: Aaron, thank you for that. That's a very eye-opening and unvarnished conversation, and I think is refreshing and needed. So I really want to just, again, say thank you for your time in speaking to us from across the pond. And thanks for being on the podcast. Really appreciate it.

Aaron: Thanks, Jeremy. Take care.

Jeremy: That was a clear-eyed look at both the scale of the global AI race and the risks of falling behind—particularly when it comes to security and access. But competing in this space isn't just about scale or speed. It also raises deeper questions about control: who owns the infrastructure, who governs the systems, and how much autonomy do countries really have.

My next guest is Jaxson Khan, Senior Fellow at the Munk School of Global Affairs and Public Policy at the University of Toronto and he's the CEO of Aperture AI. Jaxson spent a lot of his career working at the intersection of technology and public policy, including advising on Canada's AI strategy and sovereign compute investments. He's also the co-author, along with Sean Mullin, of the recent Munk School report *Sovereign by Design: Strategic Options for Canadian AI Sovereignty*, which takes a comprehensive look at Canada's position across the AI stack and lays out some options for strengthening our control and resilience. And he's been thinking deeply about what AI sovereignty actually means in practice—and what Canada needs to do differently if it wants to maintain both competitiveness and control in this space. Jaxson, welcome to the podcast, so glad we get to chat.

Jaxson Khan: Great to chat with you too, Jeremy. Thanks for having me on.

Jeremy: Let's start with the basics. "AI sovereignty." That's a term that gets used a lot, but what does that mean? You unpack that quite a bit in your report – so walk us through some of the major layers, and how that relates to Canada's AI "stack"?

Jaxson: We've been trying to define what that even is, AI sovereignty. It's a term that gets thrown around a lot. And for us, the closest we're able to come to defining it clearly is that we are able to have freedom to operate or at least freedom from coercion in building AI, adopting AI in Canada. There are a lot of different layers of the stack for AI. So we talk about

seven different layers of the stack. That includes the chips and semiconductors that are on the actual graphical processing units or CPUs themselves that power AI products and services. It also includes the data centres where those chips go and are located. And then also the actual AI models you're using. So services like ChatGPT, Claude, Gemini, these are all AI systems that are actually based on an AI model. And so one of the questions is, "Do you have control? Are you able to steer the direction of those models, or are they proprietary?" Sometimes those models are open source or open weight models, which might mean you can actually download them to your system and have more granularity, control. But all that is to say, there's a lot of different levels of AI stack. I think one of those layers is also data. Do you actually even have the data that is critical to these systems? Do you control it? Do you have the ability to fine tune that data?

We also talked about five different dimensions to digital sovereignty, and some of those include jurisdictional sovereignty. So for example, are you able to set and maintain the rules that are actually involved in AI in your country? So one of the things that's come up is, well, if a lot of these products and services, they're stored on foreign-owned infrastructure, or a lot of these models, they're all developed by US or Chinese companies, are we actually able to exert control over them with Canadian law, or are they going to be subject to the law jurisdiction of other countries? So that's one consideration.

Another would be operational control. And so do we actually have the ability to keep going if we're facing, let's say, a trade threat, not letting us have access to the most powerful models or even restricting us from access to frontier models, period? Could also be due to cyber attacks.

And another one that's a bit more societal to think about is just, are we even able to control the algorithms that are shaping our societal conversations. So for example, social media platforms, or the ones that are behind popular platforms. Sometimes there's been criticism, for example, I believe in China that, oh, well, the model's there, they wouldn't let you look up, let's say Tiananmen Square, for example. And this might be an example of epistemic capture sometimes is what it's called—fancy word basically just saying, are we able to control the way in which we receive information? Are we able to form our opinions and express them freely? And I think that's an important context of societal sovereignty, yeah.

But we do talk about what could different layers of sovereignty look like. So for example, again, I talked about jurisdictional sovereignty. So you can have something that we've analyzed to be juridical sovereignty. Try saying that three times fast. There's some examples of this in France and Germany where they actually are using, for example, Microsoft technology. So again, good powerful AI tech and data centre support for Microsoft, but they are routing that through a locally owned company or subsidiary. So at least, again, if this is about, depends on what your outcome is, but if it's about trust, if it's about control, maybe then at least you're able to exert a bit more control over the domestic firm, but then you are still able to use maybe world-class technology from a major provider. So that's kind of one of the examples we offer.

Another one may be that you have more contractual sovereignty. So for example, you may be using directly services from, again, like a Microsoft or Google, Amazon, but perhaps you, in the context of this, new geopolitical environment, you want to ask them to make sure that data stays

in Canada, doesn't get transferred elsewhere, that there are audit rights for the Canadian government, and maybe even that you, and only you, as a service provider or a customer of those services, that only you have access to the encryption keys and not the cloud infrastructure provider themselves.

So those are some of the examples that we talked about in cloud infrastructure specifically. And we focused on that layer because we found that it was more vulnerable or less sovereign in Canada. There's been reports that foreign hyperscalers own something like over 85 per cent of the cloud infrastructure in Canada. And so that's a significant amount. It just means, you know, we have to think critically about what does that mean for our AI stack in Canada.

Jeremy: Ok wow, thanks for that. So beyond those structural and technical layers, you've also pointed to culture as a key factor—and I'm thinking about risk tolerance, adoption, and whether we actually build and buy Canadian technologies. So how big of a barrier is that cultural piece, and what do we need to change for us to make meaningful progress?

Jaxson: Yeah, it's something Canada's been struggling with for, you know, some people have already the better part of a few decades. I think as Canadians, we're very risk averse. We are "small-c" conservative. You'll never hear someone saying, "Yeah, it was super easy to get capital in Canada," whether from private or public sources. The second thing is it's very hard to get contracts here. People are very hesitant to buy Canadian tech, again, whether that's companies here or the government. We just don't really buy and love our own stuff. If those cultural pieces around risk aversion don't change, it doesn't matter what structures or new incentives we try and squeeze in there, but we have to choose to build, and we also have to choose to buy. And we have to do so in Canadian way more and more, doesn't mean all the time. But again, if we just want to go to the lowest common denominator, the pieces with the economies of scale, almost every single time, that's going to be a US solution.

But we also produce some fantastic companies here, and if they produce something that is strong, I think that's a particularly good reason to at least take a look at it. And sometimes forming those strategic partnerships for companies, if they want to develop an in-house capacity, another option is they can have a strategic partnership with a Canadian partner. And that way, they might even have more ownership of that solution. And also for governments—governments have been *desperate* to catalyze our innovation ecosystem for years, often through a major combination of grants and loans. But if there's no ultimate strategy for economic value capture from those loans; if you don't have suitable IT protections; if we do not focus on ownership and long-term having stakes in those companies and incentivizing them and maybe having conditions on massive grants so they stay in Canada, or at least in a greater sense, those are all things that we can try and do.

And right now, I'd say we do a great job at subsidizing talent, subsidizing startups that then move or get bought up. And if we want to change that, think about business: It's way easier to retain a customer than to acquire a new one. And so rather than continue to fund new business development in Canada—not that we should stop—but for every marginal effort, perhaps we do a better job at retaining and ensuring that the business environment is very good for them to exist. But again, I will say a lot of these structural changes will help, but a lot of, in my opinion, still comes down to culture.

Jeremy: Let's pivot to skills and workforce impacts. So we all know AI is reshaping jobs, it's creating new ones, it's taking others, but: What does a strong national response look like when it comes to retraining and helping workers transition effectively?

Jaxson: Skills is hard when it feels like the ground is shifting beneath you. I think the government's goal of, in this new strategy, trying to create new AI-related opportunities is a good one, because there are going to be new jobs that are created by AI. The only question is, if someone loses their job in an area which might be more exposed to AI, or maybe there's some automation involved, how do you help retrain those folks? I do think that supporting people early and fast through a strategic response is very important. Like you don't want people to get in a position of insecurity or jaded or demoralized, which can certainly happen when one loses their job. And so capturing those people fast and early, I think developing a really, really good predictive data platform as a country to understand where and when proactively those changes are going to happen—those shifts—is really, really important. It would probably help us to do a much better quick retraining resource for support, and then also reallocate those people very fast.

I think there are tens of thousands of open jobs in Canada. People often simultaneously, even though there's a job crunch, unemployment at 6.6 per cent, employers will complain not being able to find the right talent. And so I think that's a market failure, but there isn't always necessarily a market solution for this. So some combination of public and private engagement on this front, of skills retraining, rapid response in the AI era, I think is very overdue. It sounds like AI strategy will speak to this to some extent, but I don't know if we're going to quite see exactly what we need there, certainly not from government action alone. So there's definitely an opportunity for both private sector entrepreneurs as well as more public interest government entrepreneurs as well to try and solve this gap.

Jeremy: Ok, I think we'll leave it there. That's a great point to end on. There's been a lot in this conversation, Jaxson. Really, my thanks to you for joining us on the Future Skills Podcast. This was a great chat.

Jaxson: Thanks, Jeremy. This was a lot of fun.

Jeremy: Those are big structural and cultural questions—and they point to the complexity of actually implementing an AI strategy that works for Canada. But even the best strategy depends on something else: Public trust, strong research foundations, and a clear understanding of how AI is actually affecting people's lives.

My last guest is Blair Attard-Frost, a Canada CIFAR AI Chair, Fellow at the Alberta Machine Intelligence Institute, and Assistant Professor of Political Science at the University of Alberta. Blair's work sits at the intersection of AI governance, power, and participation. And her research focuses on the social and political dimensions of AI—she looks at how trust, equity, and decisions about technology are made and who gets to be part of that process. Blair's been closely tracking how Canadians feel about AI, how that trust is evolving over time, and what that means for the success of policies like Canada's new AI strategy. Blair, welcome. I'm looking forward to our conversation.

Blair: Yeah, me too. Thanks very much for having me on.

Jeremy: Let's start with your reaction to the federal government's AI strategy. From your perspective, what stands out, and especially when it comes to how it addresses public trust?

Blair: It's important to contextualize this by recognizing that the whole strategy exists against a backdrop of really significant public mistrust in AI and public disinterest in using AI. It's very much the elephant in the room. The strategy refers to trust as the North Star. So they're constantly dancing around this issue without really addressing it head on. And there's survey data on this that is remarkably clear. There was an Ipsos poll, the Ipsos AI monitor. The 2026 edition of it showed that 26 per cent of Canadians feel nervous about AI and only 26 per cent of Canadians feel excited about AI. And that's not a one-off survey. This follows from 2025 when there was a survey by KPMG, University of Melbourne, showed that only 34 per cent of Canadians are willing to trust AI. 75 per cent of Canadians expect our government to regulate it. I've been following these public opinion polls on AI for many years now, for as long as they make this to like five, six years. There's been all kinds of them. And the trend line in terms of public trust, it kind of keeps going down year after year.

So I think the federal government is quite aware, you know, that this is a really significant challenge that they have to contend with, but they don't really fully acknowledge that in this strategy. I think we have a really strong ecosystem for research and talent established in Canada. I think it's good to continue investing in it if we want to build public trust. I think a really crucial investment we're going to have to make more of if we want to address the trust and enthusiasm piece is stronger investment in social sciences and humanities research on the social impacts; the labour impacts of AI. I was pleased to see that in the strategy, they're also allocating—and I think this follows an announcement from budget 2025 too—15 million dollars from existing departmental funds to a new StatsCan program to actually track social economic labour impacts of AI, which is a really important development, but it's sort of tacitly acknowledging that we have this really serious gap in our evidence base and our knowledge of the social impacts of the technology. And we're making all of these investment decisions and policy decisions without a strong evidence base. So there is, it's great to have that StatsCan program, but I think there's a lot more we can do, not only to strengthen our knowledge of social impacts and labour impacts of AI, but to build people's trust in showing that we have a government that's doing evidence-based policies around these issues.

So I think there's a social science and humanities piece that we really need to address. We need to take that—there's a lot more of a strategic priority than just building AI, which is important too. But we need to kind of elevate the SSH side a little bit more.

Jeremy: Ok, I'd like to ask you now about Canada's dependence on foreign infrastructure—particularly in research environments. How does that shape conversations that we're having about sovereignty, around openness, and around research security?

Blair: I think there's two pieces to that. For the first piece of it, I mean, our research ecosystem is extremely dependent on US platforms. I might shift the frame just a little bit here to talk about old-fashioned institutional IT governance instead of AI. Because I think one piece that gets lost in all of the conversations about the shiny new AI toy is just boring old enterprise IT systems, IT security procurement, which is really, really foundational to all the other AI stuff we're building on top of it, right? Many of the large post-secondary institutions in Canada operate on US cloud

platforms. So Microsoft, Google, AWS are really like the productivity drivers of Canadian research institutions. If we're going to have a conversation about US dependence in Canada's research ecosystem, we should really take a close look at the IT backbone that all these AI systems are being built on top of.

In my own work, I actually don't build machine learning systems, but I do a lot of qualitative research with human participants who sometimes tell me very sensitive information about their personal lives or the businesses or government departments they work for. And my own research security practice for this is I don't put primary data on our enterprise Google Cloud—interview recordings, transcripts, files that link participants to interview data. I always keep that on an encrypted local hard drive. But some universities across the country are trying to implement IT policies that would require that kind of data to be put on a Microsoft, Google, or AWS cloud environment. So I think there is a need for a national conversation right now in post-secondary and research institutions about IT infrastructure, procurement, research security that extends beyond purely AI issues. I think, you know, the digital infrastructure backbone of that is a piece that we really need to be looking at.

Then the question around academic openness, academic freedom, censorship, political interference, and science is really pronounced for social sciences, humanities, even health researchers in the US, as well as lots of folks in computer science engineering too, who are seeing really significant federal funding cuts if their research isn't perfectly aligned with the Trump administration's political priorities. So that imposes some barriers on collaboration with US researchers. And I think it also opens some opportunities for the Canadian research ecosystem to support collaboration with US researchers.

And also, in the AI research space, this could be everything from supporting collaborative research addressing bias and discrimination in LLMs; automated decisions; agentic systems; health equity issues for AI in healthcare; human rights impacts of AI; policy research for inclusive methods of engaging marginalized communities on AI policy. I think there's a really huge opportunity here to leverage Canadian values of pluralism, justice, democracy, and supporting international research collaborations with not just US researchers, but researchers in other jurisdictions as well, who are facing different kinds of restrictions related to censorship and political interference in their work.

Jeremy: Great, so I'd like to wrap up our conversation again with the question of talent, adoption, and how we can build our AI advantage: Canada has long struggled to retain and commercialize its AI expertise. So what needs to change?

Blair: Yeah, I mean, it's a great question. I think, though, that the strategy should really prompt us to pause and have a moment of national self-reflection, where we ask ourselves, "Do we want to keep doubling down on some of the same old innovation policy levers that we've desperately been trying to pull for decades, that we've convinced ourselves are the right levers to be pulling?" Or do we want to maybe take a moment to do a pulse check here and see what people across this country are actually thinking and feeling in this moment where AI is upending so many aspects of how we live and how we work? I think we really need to bring more people into these kinds of policy decisions around—that are historically centred around things like innovation and talent. And we need to rethink them through the lens of social and cultural

policy. So I think a big change we need is a cultural shift away from the sort of technocratic, expert-driven mentality where we have these groups of institutes and experts on a particular topic who are saying what needs to be done.

I think we really need to start valuing the perspectives of everyday Canadians: Indigenous communities; vulnerable populations; precarious workers; pre-seed founders; entrepreneurs who are struggling to scale; again, researchers who study and understand the social and cultural dynamics surrounding the technology. I think we need to be letting those people lead these kinds of national conversations and allow thinking about policy action to emerge from that bottom-up approach a lot more often, because those people are the ones with their boots on the ground who understand the cultural context from which all this mistrust, anxiety, disillusionment, and ultimately migration away from Canada all emerge. It's all rooted in culture and context. And I think unless we change our mindset to think about AI governance and tech governance more broadly as primarily an issue of social and cultural policy, rather than as an issue of "innovation policy" with all of the different sort of ideological trappings that conventionally entails, I'm really skeptical that we're ever going to see much improvement with the whole talent retention problem. So I think this is a social, cultural policy issue. Let's treat it as that. Let's understand what's driving all of this migration away from Canada a little bit better. And to do that, I think there's a role for publicly funded research, for social science, humanities, scholars, who are doing work on AI to play in all that.

Jeremy: I think that's a great place to end. I really appreciate that, Blair. Thanks a lot, it's been a great conversation.

Blair: Yeah, thanks a lot, Jeremy.

Jeremy: In this episode, we've explored what it will take for Canada to build its AI advantage—from competing in a global market dominated by major players, to strengthening sovereignty across the AI stack, to rebuilding public trust and rethinking how we approach innovation to curb the outbound migration of tech and intellectual property.

What's clear is that there's no single lever to pull. Success is going to depend on how well we align policy, infrastructure, talent, and culture—and how quickly we adapt to a landscape that's evolving in real time.

My thanks to Aaron Rodericks of Bluesky, Jaxson Khan of the Munk School of Global Affairs, and Blair Attard-Frost from the University of Alberta for sharing their thoughts.

You can hear all six seasons of the Future Skills Podcast on your favourite podcast app. Give us a follow if you haven't, and stay tuned for more conversations this season. This episode was produced, edited, and hosted by me, Jeremy Strachan. Sound design also by yours truly. Thanks for listening.

In partnership with: Toronto Metropolitan University, Signal49 RESEARCH, Blueprint

Funded by the Government of Canada's Future Skills Program | Canada

This research was prepared with financial support provided through the Government of Canada's Future Skills Program.
The responsibility for the findings and conclusions of this research rests entirely with Signal49 Research.