11 May 2026

AI Governance: Sovereignty, Trust and Sustainability

Conference Summary | Ivey Business School, Toronto

Prepared by Noor Us Sahar with review by Erik Bohlin, Romel Mostafa, and Zsofia Agoston Villalba.[1]

Overview

The conference brought together industry, government, and academia to examine three core challenges in AI governance: sovereignty over data and infrastructure, trust built through accountability rather than technical performance alone, and the sustainability of long-term governance frameworks. Romel Mostafa (Ivey Business School) opened by noting that the rules governing the AI economy are being written right now, under geopolitical fragmentation and significant uncertainty, and that the window to shape them remains open but may not for long.

The Lawrence National Centre works directly with Ontario's Ministry of AI, and the day's discussions were intended to inform ongoing policy deliberations. The conference was co-organized with Erik Bohlin, Ivey Chair in Telecommunication Economics, Policy and Regulation, and Romel Mostafa, Professor and Director of the Lawrence National Centre, whose research and public engagement work contributes actively to discussions surrounding digital policy and governance.

Public Sector AI: ESDC

Ima Okonny (Assistant Deputy Minister, Employment and Social Development Canada) grounded AI governance in ESDC's mandate: the department holds data on virtually every Canadian and administers Employment Insurance, Old Age Security, the Canada Pension Plan, and student loans. Her central thesis was that public sector AI governance is ultimately about people, not systems. During COVID-19, ESDC deployed a custom natural language processing model to triage millions of Employment Insurance forms, eventually reaching 99 percent accuracy. The deployment succeeded because policy, legal, privacy, and security teams were involved from the outset rather than as late-stage reviewers. ESDC also built an AI system that proactively identifies Canadians entitled to benefits who had never applied, often the country's most vulnerable people, earning international recognition from the Social Benefits Association in Geneva. Okonny's governing principles: centre everything on people and real-world impact, treat data quality as fitness for purpose, maintain transparency at every level, and actively engage communities to surface gaps in data representation.

International Governance

Marc Rotenberg (Center for AI and Digital Policy, Georgetown Law) argued that internationalization and national sovereignty are not in conflict and that meaningful governance architecture has been built over the past six years: the OECD AI Principles, the UNESCO Recommendation on AI Ethics endorsed by all 193 member nations, the EU AI Act, and the Council of Europe AI Treaty signed by Canada in February 2025 and already endorsed by more than 40 countries. A Brussels Effect is emerging, with countries including Peru and Vietnam adopting AI laws that closely mirror the EU AI Act. The United States has largely withdrawn from normative AI leadership, rescinding Biden-era executive orders and retreating from multilateral commitments, creating a vacuum that middle powers including Canada, Australia, and EU member states are positioned to fill. Rotenberg's priorities for Canada: enact comprehensive AI legislation, establish independent AI oversight with genuine technical competence, ratify the Council of Europe treaty through Parliament, and deepen multilateral engagement.

Regulating AI

Nathalie Smuha (University of Toronto Faculty of Law; former drafter of the EU AI Act at the European Commission) argued that AI governance cannot be achieved by law alone because technology is itself a form of architecture that regulates behaviour through design. She identified several core difficulties: there is no consensus definition of AI; context shapes risk in ways that make horizontal regulation across all sectors difficult to calibrate; and law is inherently over-inclusive and under-inclusive. The EU's own drafting process illustrated the moving-target problem most vividly: generative AI provisions had to be inserted into an almost-finalised text

[1] This summary report includes material developed with the support of AI-assisted drafting and summarization tools, alongside editorial contributions from the authors.

after ChatGPT launched midway through negotiations. Her closing message was that choosing not to regulate is itself a choice whose consequences accumulate, and Canada should not wait too long to establish its own sovereign position on how AI should be developed and used.

Economics of AI

Johannes Bauer (Quello Center, Michigan State University) argued that AI is an infrastructural epistemic technology that does not merely process information but accelerates the discovery process itself, as seen in protein mapping and AI-driven materials science. This raises governance questions about who controls the direction of innovation and who captures the benefits. He warned of a Hayekian risk where AI concentrates information in ways that undermine the decentralising function of markets, and of a social paradox in insurance where AI-enabled perfect risk pricing destroys the pooling function that makes insurance socially valuable. On employment, reliable net forecasts are not yet possible; early data suggests effects are heterogeneous, with younger workers in coding and customer support disproportionately affected. Agentic collusion, where AI systems coordinate pricing without explicit programming, is already generating legal cases in rental and gasoline markets. Good governance, he concluded, requires continuous monitoring and adaptive policy rather than a one-time legislative settlement.

Trust and Sovereignty

The panel on trust and governance limits produced pointed disagreements. Kevin Chan (Meta) argued trust is built through demonstrated positive impact, pointing to AI-assisted medical diagnosis, cancer trial matching, and Indigenous language preservation tools. Meta's open-source model strategy was presented as a step toward democratising access for countries that cannot build foundation models independently. Philippe Lefebvre (FIPRA, Brussels) countered that AI follows the same dynamics that produced platform monopolies, and that data rather than model architecture is the lasting competitive differentiator; Europe has given away much of its data without capturing its value. Matthew da Mota (Canadian SHIELD Institute) defined sovereignty as the ability to exert governance and law over digital infrastructure, and argued Canada does not currently have this in any meaningful sense. Jonathan Obar (York University) challenged the framing of trust itself, distinguishing between trust as mere compliance, the product of digital resignation and privacy fatigue, and trust as meaningful consent backed by real accountability. He argued governance frameworks should start from context, caution, and genuine consent.

Enterprise Governance

Jennifer Curtiss (Chief Data Officer, Scotiabank) described a layered governance framework with reporting lines to the board, a responsible AI function examining bias and transparency, and an automation effort to keep pace with deployment volume without slowing legitimate innovation. She was frank that requiring user consent as a condition of service access is not a workable regulatory model and that a fundamental rethink of digital consent is needed. Michael Page (Unity Health Toronto) described building a governance architecture from scratch in 2017, requiring every AI use case to pass review by legal, ethics, IT security, and privacy officers before deployment. Unity Health's on-premise infrastructure, initially a budget constraint, has proven advantageous: under the U.S. CLOUD Act, data held by a U.S.-incorporated cloud provider can be compelled by American authorities regardless of physical location in Canada, a serious risk for health data covering pregnancies, mental health conditions, and cancer diagnoses. Page also noted that Canada ranked 38th of 40 advanced economies on AI literacy, and that low literacy drives low trust and low clinical adoption regardless of technical performance.

Sovereign Infrastructure

Mark Graham (Bell) and Alexandre Guilbault (TELUS) both argued that the critical AI governance risk is no longer primarily about consumer privacy but about the possibility that a foreign entity could turn off systems running critical national operations. Meaningful sovereignty requires owning the full stack: Canadian-owned data centers, compute, storage, and a domestic telecommunications network that does not route data through the United States, plus Canadian-controlled encryption. The U.S. CLOUD Act means physical location alone is insufficient if the provider is U.S.-incorporated. Bell has committed to more than 300 megawatts of sovereign Canadian data center capacity. TELUS opened a sovereign AI factory in Rimouski, Quebec, and announced three new facilities in British Columbia totalling 60,000 next-generation GPUs powered by BC Hydro renewable energy, with waste heat warming approximately 150,000 homes. Both panelists called on government to clarify which workload types must remain within Canadian-owned infrastructure and to act as an anchor customer by procuring sovereign AI solutions for public operations.

AI Diplomacy for Canada

Mark Daley (Chief AI Officer, Western University; NSERC Scholar in Residence in AI) closed the conference by arguing that AI must be understood as a full physical and logical stack, from critical minerals and energy through chips, platforms, foundation models, and applications, and that Canada's strategic position differs across each layer. Canada's genuine advantages include abundant renewable hydroelectric capacity, which he described as a geopolitical asset since the binding bottleneck on AI infrastructure expansion globally is the inability to build power substations fast enough, as well as world-class AI research talent and stable democratic institutions. Because Canada lacks a decisive choke point equivalent to the Netherlands' control of extreme ultraviolet lithography equipment for semiconductor fabrication, its primary lever is aggregation: forming coalitions of like-minded nations large enough that dominant AI powers cannot ignore their collective market access conditions. Daley introduced the concept of variable geometry, noting that coalitions will form and dissolve issue by issue, and gave the example of a potential arrangement in which a Dutch-designed supercomputer is built in northern Quebec on hydroelectric power in exchange for Canadian access to Dutch semiconductor technology, advancing sovereignty for both countries across different layers simultaneously. He identified standards as Canada's most underused diplomatic tool, called for placing technically credible researchers onto international AI standards committees far more aggressively, and warned against subsidy programs that generate economic value flowing primarily to foreign shareholders. He closed by reframing sovereignty not as autarky but as optionality: preserving real choices across the AI stack so that no foreign actor can unilaterally change the terms of access to something critical. Canada's goal, he concluded, is not to win a race it cannot win, but to remain a country with genuine agency in the AI era.

Prepared from conference proceedings | Ivey Business School | May 2026

The workshop was adjourned shortly after 5PM.

Annex

Agenda of the Day – 11 May 2026

Speaker Bios

Summary Report

Monday, May 11, 2026
12 – 6 p.m.
New Donald K. Johnson Centre, 100 King St W, Suite 129, First Canadian Place

Host
Ivey Business School, Western University, Ontario

Organizers
Erik Bohlin, Ivey Chair in Telecommunication Economics, Regulation and Policy, and Romel Mostafa, Director, Lawrence National Centre for Policy and Management, Ivey Business School

Schedule **subject to change*

Time	Location
12:00 – 1:00 p.m.	**Buffet Lunch**
1:00 – 1:10 p.m.	**Welcome and Introduction** Erik Bohlin, Professor and Chair in Telecommunication Economics, Policy and Regulation, Ivey Business School; and Romel Mostafa, Director, Lawrence National Centre for Policy and Management, Ivey Business School
1:10 – 1:30 p.m.	**Opening Keynote** Ima Okonny, Assistant Deputy Minister and Chief Data Officer, Employment and Social Development Canada
1:30 – 1:50 p.m.	**Internationalization vs. Sovereignty in AI Governance** Marc Rotenberg, Executive Director and Founder, Center for AI and Digital Policy, and Adjunct Professor, Georgetown Law, Washington D.C.
1:50 – 2:10 p.m.	**AI Governance: The Challenges in Regulating AI** Nathalie Smuha, Professor, University of Toronto, Faculty of Law
2:10 – 2:30 p.m.	**The Economics of the Emerging AI Ecosystem** Johannes Bauer, Professor and Director, Quello Center, Michigan State University; and former Chief Economist, Federal Communications Commission
2:30 – 2:55 p.m.	**Coffee Break**
2:55 – 3:40 p.m.	**Panel: The Limits of AI Governance: Innovation, New Big Tech and Trust** Moderated by Erik Bohlin, Ivey Business School Kevin Chan, Public Policy Director, Meta Matthew da Mota, Research Director, Emerging Technology and National Security, Canadian SHIELD Institute Jonathan Obar, Associate Professor, York University Philippe Lefebvre, Senior Advisor, FIPRA, Brussels
3:40 – 4:10 p.m.	**Panel: Perspectives by Users** Moderated by Romel Mostafa, Ivey Business School Jennifer Curtiss, Chief Data Officer, Scotiabank Michael Page, Interim Senior Director, Data Science and Advanced Analytics, Unity Health Toronto

4:10 – 4:40 p.m.	**Panel: Perspectives by Service Providers** Moderated by Romel Mostafa, Ivey Business School Mark Graham, Senior Vice President, Legal and Regulatory Affairs, BCE Alexandre Guilbault, Vice President, AI Enablement, TELUS
4:40 – 5:00 p.m.	**Closing Keynote: AI Diplomacy for Mid-size Powers: Implications for Canada** Mark Daley, Chief AI Officer, Western University
5:00 – 6:00 p.m.	**Networking Reception with Cocktails and Hors d'oeuvres**

Speakers

Johannes Bauer is a researcher, writer, teacher, and academic entrepreneur. He is interested in the digital economy next-generation media. His position as Quello Chair for Media and Information Policy facilitates the pursuit of rigorous and actionable research. From September 2023 through December 2024, he was on leave from MSU to serve as the Chief Economist in the Office of Economics and Analytics (OEA) of the U.S. Federal Communications Commission in Washington, DC. Educated as an engineer and social scientist, he obtained advanced degrees in economics from the Vienna University of Economics and Business, Austria. Michigan State University has been his home institution since 1990. He had the privilege to spend extended times affiliated with Delft University of Technology, The Netherlands (2000-2001), the University of Constance, Germany (Summer 2010), and the University of Zurich, Switzerland (2012).

Erik Bohlin is Professor and Chair in Telecommunication Economics, Policy and Regulation at the Ivey Business School. He is an expert in telecommunications policy, an inter-disciplinary topic concerned with the impact of digitalization in the economy and society. He is Editor-in-Chief of *Telecommunications Policy*, a premier journal in the field. He is on leave as Professor at Chalmers University of Technology, Sweden. His graduate degree is in Business Administration and Economics at the Stockholm School of Economics (1987) and his Ph.D. is from Chalmers University of Technology (1995). He is a Member of the Swedish Royal Academy of Engineering, and Past Chair of the International Telecommunications Society, an inter-disciplinary professional society convening conferences on the evolving digital society and policy needs.

Kevin Chan is a Public Policy Director at Meta, focusing on the future of the internet including artificial intelligence and wearables. In 2024, he was elected Chair of the Board of MediaSmarts, Canada's digital literacy centre. A former government executive and university administrator, Kevin launched Facebook's Canadian policy function and spent 7 years as its Head and then Director of Policy. He has driven impact in areas as diverse as platform integrity and the promotion of Indigenous languages online. His work fighting white nationalists made NOW Magazine's 2019 Year In Review, and his efforts leading Facebook's Canadian Election Integrity Initiative were recognized with a Harvard Technology and Democracy Fellowship. In 2023, UNESCO named his effort translating Facebook into Inuktitut an initiative of "digital empowerment driving the International Decade of Indigenous Languages". A Fellow of the Royal Canadian Geographical Society, Kevin was awarded the Quest Medal in 2025 for contributions to the discovery of Quest, Shackleton's last ship, and laying the groundwork for digital immersive experiences that explore this discovery. He is also a recipient of Canada's Meritorious Service Medal, the Queen's Diamond Jubilee Medal, and the King's Coronation Medal.

Jennifer Curtiss has been the Chief Data Officer at Scotiabank since 2025. In this role, she is responsible for enterprise data strategy, data governance, and the advancement of data-driven decision-making across the organization. Prior to joining the bank, Jen held senior data roles at American Express, Freddie Mac, Citi and Moodys, leading large-scale data initiatives, modernized data platforms, and built high-performing teams. Jen is recognized for a pragmatic and collaborative leadership approach, with a strong track record of translating complex data challenges into clear, actionable outcomes. She is a committed advocate for data quality, responsible data use, and building a strong data culture that empowers teams and enables sustainable growth.

Mark Daley was appointed to the as Western's first-ever Chief AI Officer for a five-year term in October 2023. A respected researcher in the field of neural computation, Mark's career includes a tenure as Vice-President Research at CIFAR, a world-renowned institute supporting AI research and leading Canada's national AI strategy. Additionally, Mark is a multidisciplinary scholar and has held cross-appointments in several departments across campus, including Computer Science, Mathematics, Statistics & Actuarial Sciences, Biology, Electrical & Computer Engineering, and Epidemiology & Biostatistics. Western is the first university in Canada to create an AI leadership role within its senior executive, and Mark is uniquely qualified for this exciting new role that will help propel Western to the forefront of AI research and application. Most recently Mark served as Western's Chief Digital Information Officer leading Western Technology Services (WTS). In this capacity he brought together the distributed IT community with a collaborative, respectful and federated approach through the

creation of the Strategic Technology forum, drafted the Agile IT Governance framework, and co-created a service-oriented set of strategic objectives guiding the work of WTS.

Matthew da Mota is Research Director, Emerging Technology and National Security at the Canadian Shield Institute. His research examines how emerging technologies and shifting geopolitical dynamics reshape governance, knowledge, and sovereignty — with a particular focus on the national security implications of AI and digital infrastructure. Matt previously held a post-doctoral fellowship jointly between the Centre for International Governance Innovation (CIGI) and the University of Toronto's Media Ethics Lab, where he studied AI's transformation of research and knowledge production. Prior to that, he worked at CIGI's Digital Policy Hub on international AI governance. His Ph.D. research examined the relationship between historical narrative, written evidence, and political power. He currently serves on multiple international standardization technical committees and leads research projects on trust, governance, and epistemology.

Mark Graham is Senior Vice President, Legal & Regulatory at Bell and oversees the teams responsible for all regulatory, competition, commercial, corporate development and M&A, litigation, and privacy / cybersecurity legal matters. Prior to joining Bell, Mark was a member of the Competition, Antitrust, & Foreign Investment and Communications groups at Blake, Cassels & Graydon LLP obtained an honours joint specialist degree in Economics & Philosophy from the University of Toronto (2006) and J.D. at the University of Toronto Faculty of Law (2009).

Alexandre Guilbault is Vice President of AI Business Enablement at TELUS. He also leads the TELUS AI Accelerator, a team of over 120 professionals within the Chief Information Office, responsible for guiding the company's AI transformation strategy. The team takes a strategic consulting approach to AI adoption, working closely with leaders across TELUS Business, Consumer, Digital, Health, Agriculture & Consumer Goods (TAC) and the Office of the CTO to modernize business operations and deliver practical, measurable value using applied, generative and agentic AI. Under Alexandre's leadership, the AI Accelerator has grown from 15 people to over 120 and generates more than $100M in annual business value, validated through A/B testing and signed off by the business. The team maintains under 10% attrition in a market where the industry average exceeds 28% and achieved a 97% engagement score. With 15 years of experience in the technology industry, Alexandre brings a deep background in software engineering, machine learning and entrepreneurship. He holds a Master of Business Administration from HEC Montréal and a Bachelor of Engineering from Polytechnique Montréal.

Philippe Lefebvre is Senior Advisor at FIPRA, Brussels. He has built a career combining strategic leadership and engineering roles spanning the telecommunications and financial sectors, with experience in both the private sector and public policy at the EU level. Since July 2024, he has been working as an independent consultant, specialising in digital and telecommunications strategy, as well as advising clients in the fields of payments and decentralized finance (DeFi). From 1995 to 2024, Philippe held various middle management positions at the European Commission. In his most recent role, he was responsible for the Commission's 5G deployment strategy, including policy development and investment incentive actions to achieve the connectivity targets of the EU Digital Decade Programme. Earlier in his career at the Commission, he held positions with responsibilities covering broadcasting regulation, radio spectrum policy, electronic commerce, and EU strategic research in the field of financial technologies and digital euro. Prior to joining the European Commission, Philippe began his career at Belgium's Générale Bank, now Fortis Group, from (1983-1986). He then held management roles in the international payments industry, serving as Vice President, Technology Programmes at Eurocard, then Europe's arm of Mastercard (1987-1991) and later as Director of Debit Products at Visa International in the U.S. (1992-1995). Born in Luxembourg, Philippe Lefebvre holds a MSc degree in Electrical and Mechanical Engineering and an BA in Economics, both from the Brussels University.

Romel Mostafa is an Assistant Professor of Business, Economics and Public Policy at the Ivey Business School. Romel's areas of research and expertise include strategy & capability development in new firms, innovation & competitive dynamics, industrial evolution & policy, as well as behavioural decision-making. He has published in several leading academic journals, including Academy of Management Journal, Journal of Behavioral Decision Making, Journal of Risk & Uncertainty, Organization Science and Management Science. His research and commentaries have been featured in global media outlets such as CNN, NPR and the New York Times. Romel has taught both at graduate and undergraduate levels, and received several teaching awards. He obtained his PhD and MSc from Carnegie Mellon University, and BA from Lawrence University. As the Director of Ivey's Lawrence National Centre for Policy and Management, Romel spearheads the Centre's research, outreach and teaching initiatives. The Centre advocates for sound policy and corporate action towards unlocking national competitive advantage, by focusing on critical challenges and opportunities around digital, trade and social infrastructural pillars.

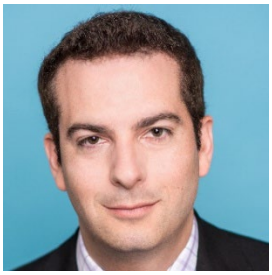

Jonathan Obar is an Associate Professor in the Department of Communication and Media Studies at York University. His research focuses on communication policy and the relationship between digital technologies, civil liberties, and the inclusiveness of public cultures. Recent research addresses corporate AI transparency, consumer behaviours, and online consent protections (visit www.ainfocus.net and www.biggestlieonline.com for more information). He is currently co-editing a special issue of *Telecommunications Policy* entitled "Policy Responses to Generative AI". He has served as Associate Director of the Quello Center at Michigan State University, Research Fellow with the New America Foundation and Free Press, and as Senior Advisor to the Wikimedia Foundation's Wikipedia Education Program.

Ima Okonny is Assistant Deputy Minister and the Chief Data Officer at Employment and Social Development Canada (ESDC), where she leads enterprise-wide initiatives advancing Data Science, Data Management, and Data Sharing. Her work directly supports the department's policy development, service delivery, and results reporting, strengthening the use of evidence across the organization. Over a 26-year career in the data and analytics domain, Ima has driven significant enhancements to the federal evidence base. She has successfully designed and implemented Data Literacy programs, Data Strategies, and enterprise policies, and has led the development of Artificial Intelligence and Data Ethics Frameworks, along with core tools for assessing and managing data-related risks in complex operational environments. Prior to joining ESDC, Ima served in an executive role at the Office of the Superintendent of Financial Institutions (OSFI), where she led a team responsible for the effective collection, governance, and management of data from federally regulated financial institutions and pension plans.

Michael Page is the interim Senior Director of Data Science and Advanced Analytics (DSAA) at Unity Health Toronto. In this capacity, he leads a high-impact team dedicated to deploying AI solutions and optimizing hospital workflows to enhance patient outcomes and health system efficiency. Mike also serves as the Director of AI Commercialization at Unity Health, where he bridges the gap between technical AI development and scalable, real-world healthcare applications. A committed educator and strategist, Mike is an Instructor for Ivey Executive Education and an Advisor and Lecturer for WatSPEED at the University of Waterloo. His academic background includes serving as a Sessional Lecturer for the Executive MBA program at the Ivey Business School. Additionally, he provides strategic mentorship to emerging startups through the Lab2Market program at Toronto Metropolitan University and McMaster University. With over 15 years of experience, Mike has held senior leadership tenures at the Vector Institute for Artificial Intelligence and the University of Toronto. His expertise lies at the intersection of innovation, R&D, and corporate strategy, with a focus on translating

complex technology into social and economic impact. Mike holds an MBA from the Ivey Business School and a BA from the University of Toronto. He remains an active mentor to students and professionals navigating the fields of business, innovation, and social impact.

Marc Rotenberg is Executive Director and Founder of the Center for AI and Digital Policy. He is a leading expert in data protection, open government, and AI policy. He has served on many international advisory panels, including the OECD AI Group of Experts. Marc helped draft the Universal Guidelines for AI, a widely endorsed human rights framework for the regulation of Artificial Intelligence. Marc is the author of several textbooks including the *2020 AI Policy Sourcebook* and *Privacy and Society* (West Academic 2016). He teaches privacy law and the GDPR at Georgetown Law. Marc has spoken frequently before the US Congress, the European Parliament, the OECD, UNESCO, judicial conferences, and international organizations. Marc has directed international comparative law studies on *Privacy and Human Rights, Cryptography and Liberty,* and *Artificial Intelligence and Democratic Values.* Marc publishes widely in academic journals. Marc is a graduate of Harvard College, Stanford Law School, and Georgetown Law. He is a Life Member of the American Bar Foundation, the Council on Foreign Relations, and the European Law Institute.

Nathalie Smuha is an Assistant Professor of Law at the University of Toronto. Her research and teaching focus on the intersection of law, philosophy and technology, with particular attention to the impact of digital technologies on human rights, democracy and the rule of law. She is the author of *Algorithmic Rule By Law: How Algorithmic Regulation in the Public Sector Erodes the Rule of Law*, and the editor of the *Cambridge Handbook on the Law, Ethics and Policy of Artificial Intelligence* (both published with Cambridge University Press, 2025). Previously, she was Assistant Professor and FWO Postdoctoral Fellow at the KU Leuven Faculty of Law, and the academic coordinator of the KU Leuven Summer School on the Law, Ethics, and Policy of AI. She has also taken up Adjunct Professorships at NYU School of Law and Columbia Law School, and she held visiting positions at the University of Oxford, the University of Chicago and the University of Birmingham. Besides her academic activities, Professor Smuha regularly advises governments and international organizations on AI policy and regulation. She coordinated the work of the European Commission's High-Level Expert Group on AI and acted as a scientific expert in the Council of Europe's (Ad Hoc) Committee on AI. She also assisted the OECD and UNESCO with the development of AI policy.

AI Governance: New Tradeoffs for Sovereignty, Trust and Sustainability

Conference Summary

May 11, 2026 | Ivey Business School, Toronto

Session 1: Welcome and Introduction

Speaker: Romel Mostafa, Director, Lawrence National Centre for Policy and Management, Ivey Business School

Setting the Stage

Romel Mostafa opened the conference by pointing to the wide range of participants across industry, government, academia, and other sectors as proof that AI governance is a cross-cutting issue that requires collaboration beyond any single domain. He organized the discussion around three core themes: sovereignty, trust, and sustainability. Sovereignty focused on questions of control over data, infrastructure, and standards in a context of deep global interdependence, with particular implications for countries like Canada balancing autonomy and reliance on dominant technology ecosystems. Trust was framed as something that must be built through governance, accountability, and public engagement, not just technical performance. Sustainability emphasized the need for long term, adaptable governance frameworks that prioritize future generations.

He also used the land acknowledgement to introduce a broader philosophical perspective, drawing a parallel between Indigenous stewardship and AI governance as a responsibility grounded in accountability, reciprocity, and long term care rather than simple control. Mostafa underscored that the conference was closely tied to real policy work in Ontario, positioning participants as active contributors to shaping AI governance, and concluded by recognizing the complexity of the issues and the importance of thoughtful, forward looking dialogue.

The Urgency of the Moment

Mostafa was direct about the historical stakes of the moment: governance frameworks for AI are being written right now, in real time, under conditions of significant uncertainty and geopolitical fragmentation. Governments and firms around the world are experimenting with radically different models, from the relatively prescriptive EU AI Act to the more permissive, market-led approaches favored in the U.S. under the current administration, to the heavily state-directed approaches of China. There is no settled consensus, and the window to shape the architecture of the AI economy remains open — but may not remain open for long.

He emphasized that the Ivey conference was not merely an academic exercise. The Lawrence National Centre works directly with Ontario's Ministry of AI (the office of Minister Mark Schaan), and the insights generated from the day's discussions would be passed on to inform real policy deliberations. In that sense, the participants were not just observers of AI governance — they were, as Mostafa put it, "building the institutional architecture of the AI economy while simultaneously living in it."

He closed the introduction with a note of intellectual humility: the issues are genuinely complex, there are no easy answers, and the goal of the conference was not to arrive at tidy conclusions but to advance practical governance thinking through rigorous, constructive, and forward-looking dialogue. He also offered a warm personal thank-you to co-organizer Erik Bohlin for his tireless and graceful work in bringing the event together.

Session 2: Opening Keynote

Speaker: Ima Okonny, Assistant Deputy Minister and Chief Data Officer, Employment and Social Development Canada (ESDC)

ESDC's Mandate and the Stakes of AI

Ima Okonny opened by grounding her remarks in the unique scale and sensitivity of ESDC's mandate. The department touches virtually every Canadian: if you have a Social Insurance Number, ESDC has your data. Student loans, Employment Insurance, Old Age Security, the Canada Pension Plan -- all of these programs flow through the department, making it arguably the most data-rich ministry in the federal government. This is not a point of pride so much as a point of obligation. The sheer volume and sensitivity of data ESDC holds means that every AI decision carries real consequences for real people's lives.

Her keynote was built around a central thesis: AI governance in the public sector is ultimately about people, not systems, not infrastructure, not efficiency metrics. Every AI deployment at ESDC is evaluated first and foremost against the question of whether it serves Canadians equitably, safely, and transparently. ESDC serves Canadians ranging from children and families to seniors, newcomers, persons with disabilities, and Indigenous peoples across vastly different geographies. Leaving any group out, Okonny stated plainly, is not acceptable.

COVID-19 and the EI Case Study

The most detailed example Okonny presented came from the COVID-19 pandemic. As millions of Canadians lost jobs during the lockdowns, ESDC was inundated with Employment Insurance applications. Officers were overwhelmed by Record of Employment forms, such as paper and digital documents with extensive freeform comment fields, while Canadians urgently needed their checks processed to feed their families.

Okonny's data team developed a custom AI model, trained entirely on data held within the department, to triage these forms automatically. Using natural language processing, the model scanned comment fields, flagged issues, and routed forms to the appropriate processing path, allowing officers to concentrate their attention where it was most needed. The model eventually reached 99% accuracy. Forms it was uncertain about were escalated to human officers for review.

What made this notable was not just the result but the context: the EI system is one of the oldest legacy systems in the federal government, originally built simply to issue cheques. Getting a custom machine learning model into production there was widely considered impossible. ESDC succeeded by bringing all relevant parties, including policy, legal, privacy, cybersecurity, and program officers, into the process from the outset, not as late-stage reviewers. This cross-functional approach allowed the team to navigate the Privacy Act, ESDC's own enabling legislation, and various security protocols without hitting legal obstacles mid-deployment.

The Broader AI Portfolio

Beyond the EI example, Okonny organized ESDC's AI work around three goals: productivity, efficiency, and effectiveness. On productivity, the department uses natural language processing and generative AI to process large volumes of incoming information more quickly, with generative tools applied only to non-personal data. On efficiency, optical character recognition is now used to pre-screen documents in the SIN application process, catching problems before an officer reviews the file and reducing unnecessary in-person visits. On effectiveness, ESDC developed an AI system to proactively identify Canadians entitled to benefits who had not applied, often the most vulnerable people in the country. This initiative earned international recognition from the Social Benefits Association in Geneva.

Guiding Principles

Okonny closed by articulating the principles that underpin all of ESDC's AI work. First, center everything around people and their real-world impact. Second, treat data quality as fitness for purpose, not just statistical validity; a lesson learned acutely during COVID, when contextual differences between regional data sources required sitting with program officers to understand what variables actually meant in practice. Third, learn, iterate, and scale continuously, with ongoing attention to demographic shifts and model drift. Fourth, maintain transparency at every level, from public disclosure of AI projects to Deputy Minister-level governance oversight. Fifth, engage communities actively, bringing in external voices to challenge internal assumptions and surface gaps in data representation, particularly for Indigenous peoples and women.

Session 3: Internationalization vs. Sovereignty in AI Governance

Speaker: Marc Rotenberg, Executive Director and Founder, Center for AI and Digital Policy; Adjunct Professor, Georgetown Law

Framing: Not a Trade-Off

Marc Rotenberg opened by establishing a key conceptual position: he does not regard internationalization and national sovereignty as inherently in conflict, and he resists the language of "balancing" or "trade-offs" that often dominates AI governance discussions. In his view, good public policy should always seek non-zero-sum solutions, and his goal was to show how international frameworks and national sovereignty can coexist and reinforce each other rather than compete. This framing shaped everything that followed.

Rotenberg's organization, the Center for AI and Digital Policy, produces the AI and Democratic Values Index -- a comprehensive annual assessment scoring 90 countries across 12 criteria, generating both a detailed narrative report and a quantitative ranking. The project began in 2019 as a modest effort to curate responsible AI guidelines and has grown substantially as the governance landscape has developed. The report is freely available online.

Progress on International Governance Frameworks

Rotenberg argued that progress on international AI governance is a genuinely positive story that tends to be underappreciated. While public discourse focuses heavily on the pace of AI development and the inadequacy of regulatory responses, a meaningful body of governance

architecture has in fact been built over the past six years, and the Center has been an active participant in constructing it.

He traced the key milestones. The OECD AI Principles (2019), developed with U.S. participation, established shared baseline expectations around transparency, accountability, and human-centeredness. The UNESCO Recommendation on AI Ethics followed, endorsed by all 193 member nations. The EU AI Act created a binding, risk-based framework for AI systems across the European market. Most recently, the Council of Europe AI Treaty, which was signed by Canada in Paris in February 2025, represents what Rotenberg considers the most promising development: unlike the EU AI Act, the treaty is open to both member and non-member states of the Council of Europe, creating a genuine pathway to a global governance instrument. More than 40 countries have already endorsed it. Rotenberg also noted the United Nations' establishment of an independent scientific panel of 40 leading AI experts as an important institutional step forward.

He observed early signs of a "Brussels Effect" in AI: the well-documented tendency for EU regulations to set de facto global standards because multinational companies prefer a single compliance framework. Already, new AI laws in Peru and Vietnam appear to closely mirror the structure of the EU AI Act, suggesting the Act's reach may extend well beyond Europe's jurisdiction.

The U.S. Withdrawal from AI Governance Leadership

The most pointed section of Rotenberg's remarks addressed the significant shift in U.S. engagement with global AI governance under the current administration. He drew a careful distinction between two different kinds of U.S. presence: market dominance and commercial activity on one hand, and normative leadership in shaping international governance frameworks on the other. The U.S. remains dominant on the first dimension but has largely withdrawn from the second.

Rotenberg traced this shift through specific decisions: the rescinding of the Biden administration's executive order on AI safety; the withdrawal from UNESCO, which the U.S. had recently rejoined specifically because of the quality of its AI ethics work; and a broad retreat from multilateral commitments. He also noted a subtle but telling rhetorical shift in U.S. official language: where previous administrations spoke of American "leadership", a term that implies other countries have agency and can choose to follow, the current framing uses the word "dominance," which carries a different connotation entirely. He noted with some irony that the first Trump administration had actually been a constructive participant in the 2019 OECD process, and had issued a solid executive order in 2020 establishing AI guardrails across federal agencies. That progress, along with the Biden-era contributions to the Council of Europe treaty process, is now effectively being unwound.

The consequence, Rotenberg argued, is a real vacuum in global AI governance leadership. In earlier years, a U.S. withdrawal of this kind would have been read primarily as an opportunity for China to expand its influence in international standards bodies. But the AI governance landscape today is different. Middle-power countries, such as Canada, Australia, EU member states, the UAE, and others, have developed genuine governance competencies and credibility. Rather than a U.S.-China binary filling the space, Rotenberg expects to see increasing coalitions among middle powers that are cautious about both Washington and Beijing, finding common ground where cooperation is possible.

Sovereignty: Legitimate Interests and Real Risks

Rotenberg then addressed the sovereignty side of his title directly. He identified five legitimate reasons nations pursue AI sovereignty: protecting national infrastructure from external disruption; preserving cultural identity in the face of AI systems built on foreign data reflecting foreign values; advancing economic development; ensuring democratic accountability over automated systems; and safeguarding data protection frameworks rooted in domestic legal traditions. All of these, he argued, are genuine and defensible interests that international frameworks should accommodate.

The challenge arises when sovereignty claims are used to shield authoritarian AI practices from oversight, block legitimate international accountability mechanisms, or fragment global standards in ways that allow repression to operate without scrutiny. He gave the example of AI-powered surveillance: a country that builds pervasive citizen monitoring infrastructure and then invokes sovereignty to resist international review is not protecting legitimate national interests -- it is using sovereignty as cover for impunity. The goal of good governance design, he suggested, is to construct frameworks that honour the former while constraining the latter.

Canada's Standing and Priorities

Canada has ranked in the top tier of the AI and Democratic Values Index for all six years the report has been published, alongside Japan, the Netherlands, Norway, Switzerland, and the UK. Rotenberg praised Canada's strong data protection infrastructure, active participation in international governance processes, and serious culture of responsible AI in the public sector. He did note one deduction: the failure of Canada's Artificial Intelligence and Data Act (AIDA) to advance through Parliament cost the country half a point in the most recent assessment.

His to-do list for Canada had four priorities. First, enact comprehensive AI legislation. The absence of a coherent national framework is a real gap, and other top-performing countries are demonstrating that it is achievable without stifling innovation. Second, establish independent AI oversight with the technical competence to evaluate risk; many governments are creating dedicated AI offices precisely because data protection regulators alone lack the specialized expertise. Third, ratify the Council of Europe AI Treaty through Parliament. Canada has signed; parliamentary ratification is the essential next step. Fourth, continue and deepen Canada's tradition of constructive multilateral engagement, particularly given the leadership vacuum left by the current U.S. posture. Progress, as Rotenberg concluded, moves forward.

Session 4: AI Governance -- The Challenges in Regulating AI

Speaker: Nathalie Smuha, Professor, University of Toronto, Faculty of Law

Framing the Regulatory Task

Nathalie Smuha opened by noting how well her presentation complemented Marc Rotenberg's preceding remarks. Where Rotenberg had surveyed the international governance landscape, Smuha brought a practitioner's perspective: she had worked inside the European Commission during the drafting of the EU AI Act, and her talk was designed to answer a more fundamental question. Assuming we agree that AI needs to be regulated, what choices does a government actually face, and what are the genuine difficulties of making those choices well?

She began by broadening the definition of regulation itself, drawing on Lawrence Lessig's influential framework from Harvard. Regulation is not only law. It encompasses four distinct modes: formal law, social and ethical norms, market forces, and architecture, meaning the design

of physical or technical systems that shape behavior. Her example of the Robert Moses bridges on Long Island, deliberately built too low for public buses so that only car-owners could access the beach, illustrated how design encodes values and constrains choices without any explicit legal command. Technology, she argued, is architecture in exactly this sense. Smartphones, social platforms, and AI systems all regulate us through their design, not just through the terms and conditions governing their use. For policymakers, this means that AI governance cannot be achieved by law alone; it requires attending to all four modes simultaneously.

She was also careful to note that regulation is not inherently restrictive. It can be enabling as well as protective. Attracting AI talent from abroad, granting subsidies for AI research, creating fast-track processes for AI companies -- all of these require regulation. The policy choice is not between regulation and no regulation, but between different kinds and combinations of regulatory tools.

The Specific Challenges of Regulating AI

Smuha identified several challenges that make AI particularly difficult to regulate through law, each of which requires deliberate choices from legislators.

The first is definitional. There is no consensus, even among AI experts, on what AI actually is. The OECD's work on a legally workable definition, which Rotenberg had helped shape, was influential precisely because it tried to answer this question for regulatory rather than scientific purposes. But the challenge is compounded by the AI effect: the tendency for technologies to stop being perceived as AI once they become familiar. Spam filters, autocomplete, and navigation software were all once considered AI. If a regulation defines AI too narrowly, companies can argue their way out of scope. If it defines it too broadly, it captures activities that pose no meaningful risk.

The second challenge is that context matters enormously. The risks of a rule-based system differ from those of a machine learning system, which differ again from those of a generative AI system. Privacy concerns cut across all of them, since they all run on data, but hallucinations are specific to generative AI. Regulatory design must account for this heterogeneity. Should a country adopt a horizontal AI Act covering all sectors, as the EU chose to do, or sector-specific rules for healthcare, transport, finance, and so on? Each approach has tradeoffs. The EU's horizontal approach risks applying the same requirements to very different contexts, while sector-specific approaches may miss cross-cutting risks or create inconsistencies.

The purpose for which a technology is used also matters, sometimes more than the technology itself. Smuha used the example of deepfake technology: the same AI capability that can generate synthetic advertising models can also produce a realistic video of a political figure saying something they never said, timed to drop days before an election. By the time the target can prove the video is fake, the election may already be over. The EU has moved to ban AI systems capable of producing such content for certain purposes -- but writing laws that prohibit specific harmful applications without also banning beneficial uses of the same underlying technology is genuinely difficult.

A third challenge is that law is both over-inclusive and under-inclusive. Any regulatory definition will capture things legislators did not intend to regulate while missing things they did. Legislators can choose between exhaustive lists of covered applications, which provide certainty but may quickly become outdated, and broad principles-based definitions, which are more flexible but harder to enforce. Equally, they must choose between prescriptive rules (you must produce technical documentation and go through this specific certification process) and open standards

(you must take reasonable care not to discriminate). Rules provide clarity but may not age well; standards are more adaptable but create uncertainty for those trying to comply.

Perhaps the most vivid illustration of the moving-target problem came from the EU's own experience. When the first draft of the AI Act was written, generative AI was theoretically possible but practically obscure. Then ChatGPT arrived, midway through legislative negotiations, and regulators scrambled to insert new language into a text that was almost finalized. Smuha described this as a useful lesson: no matter how carefully a regulation is drafted, technological developments will outpace it. Future-proofing legislation means writing it with enough flexibility to absorb developments that cannot yet be anticipated, while still providing sufficient clarity to be enforceable.

She also addressed the ex ante versus ex post question: should regulation prevent harm before a system reaches the market, or respond to harm after it occurs? Ex ante requirements, such as mandatory conformity assessments or market authorizations, catch problems before they affect people but can slow deployment and impose costs on smaller innovators. Ex post approaches, such as liability regimes and compensation schemes, allow more rapid deployment but may leave people harmed before the legal system catches up.

Objections and Takeaways

Smuha addressed three common objections to AI regulation directly. The first is that regulators do not understand the technology well enough to regulate it. She dismissed this efficiently: legislators do not need to be microbiologists to enact food safety law, and they do not need to be machine learning engineers to enact AI law. The second objection is that law can never keep pace with technology. She pointed to Belgium's civil code, parts of which trace back to Roman law, as evidence that well-drafted legislation can remain applicable across centuries when written at a sufficiently principled level. The third and most politically potent objection is that regulation stifles innovation. Her response was that this depends entirely on how regulation is designed. Poorly designed regulation does create friction and overhead. Well-designed regulation can actively stimulate innovation toward socially desirable directions, as the EU's investment in SME compliance tools attempts to do.

Her closing message was that there is no ideal solution. Every regulatory choice involves tradeoffs, and each jurisdiction must make its own judgments about what those tradeoffs should be. She encouraged Canada not to wait too long. Watching the EU's experience and learning from it is valuable, but establishing a sovereign position on how AI should be developed and used in Canada is itself a form of sovereignty. Choosing not to regulate is also a choice, and its consequences accumulate over time.

Session 5: The Economics of the Emerging AI Ecosystem

Speaker: Johannes Bauer, Professor and Director, Quello Center, Michigan State University; former Chief Economist, Federal Communications Commission

A New Kind of Technology

Johannes Bauer opened with a provocation: he asked the room how many people trusted government to regulate AI effectively. The response was sparse. He then asked how many trusted economists. Somewhat to his own surprise, more hands went up for the second question, which

he noted was genuinely unusual in his experience. He accepted the mandate with good humor and proceeded to offer an economist's analysis of why AI is so difficult to govern well, and why the stakes of getting it wrong are so high.

His central thesis was that AI is not simply a faster or smarter version of existing digital technologies. It is what he called an infrastructural epistemic technology: a technology that not only stores, moves, and processes data, but also generates knowledge about its own functioning and about the world it models. Earlier epistemic technologies -- search engines, Wikipedia, the internet itself -- changed how humans access and aggregate information. AI does all of that, but also accelerates the discovery process itself. It is, in a phrase he borrowed from economic history, an invention of a method of inventing. The protein-mapping breakthroughs achieved by AI in recent years, and the AI-driven materials science being done in robotic laboratories, are examples of a technology that does not just answer questions but generates new questions worth asking.

This distinction matters for governance. Technologies that process information are governed differently from technologies that generate knowledge and accelerate discovery. The former can be assessed on accuracy, speed, and privacy. The latter raises deeper questions about who controls the direction of innovation, what kinds of discovery are incentivized, and who benefits from the results.

How AI Disrupts Economic Theory

Bauer organized much of his talk around three intellectual traditions in economics and what AI does to each of them.

The first is Friedrich Hayek's insight that knowledge in society is decentralized: no single actor can aggregate it all, and competitive markets are the most efficient known mechanism for doing so. Price signals coordinate the dispersed knowledge of millions of individuals. AI challenges this framework in two ways. On one hand, it can aggregate information at a scale and speed that was previously impossible, potentially enabling new efficiencies. On the other hand, it may allow information to become more concentrated rather than less, reducing the decentralizing function that gives markets their value. Worse, when AI formalizes tacit knowledge -- the contextual, local, relational knowledge that program officers at ESDC understand better than any algorithm, as Okonny had illustrated -- it risks encoding that knowledge in systems controlled by a small number of actors rather than leaving it distributed across communities. Bauer called the worst-case version of this a Hayekian nightmare: a world of extreme information concentration in large organizations and governments, with the decentralizing function of markets undermined.

The second tradition is the information economics of Stiglitz, Akerlof, and Spence, who showed that markets often fail when information is asymmetric or hidden. Insurance markets are a canonical example: adverse selection arises because individuals know more about their own risk than insurers do. AI can reduce some of these asymmetries by enabling hyper-personalized risk assessment. But Bauer pointed out the paradox this creates: if health insurers can perfectly price individual risk, the social pooling function of insurance collapses. The market becomes technically efficient while losing its social purpose entirely.

The third tradition concerns organizations and innovation. Oliver Williamson showed that organizations arise when markets fail because transaction costs are too high. AI can reduce those costs, potentially reorganizing firm boundaries in unpredictable ways. And for innovation specifically -- which Bauer noted accounts for 60 to 70 percent of economic growth -- AI both accelerates and distorts the discovery process. It can explore the space of possible innovations at unprecedented speed, but it does not inherently decide which innovations are desirable. That requires human judgment. The risk is that AI accelerates innovation in directions shaped by

whoever controls the models and training data, not by broader social deliberation about what problems are worth solving.

Jobs, Inequality, and the Limits of Forecasting

On the question of employment, Bauer was honest about the limits of economic prediction. Estimates of AI's net effect on jobs range from massive losses to substantial gains, depending on the model used and the assumptions built in. What early data does suggest is that effects are heterogeneous and industry-specific. In coding and customer support, for example, younger workers appear to have been disproportionately affected following the introduction of capable AI tools, while more senior workers have been less impacted. Some industries with strong demographic demand drivers, such as nursing, have seen job growth that masks AI-related displacement. The honest answer, Bauer argued, is that the net outcome cannot be forecast reliably in advance because it depends on how a complex dynamic system evolves over time -- including how consumer spending, industry structure, and policy all interact.

He was similarly nuanced on market concentration. Some degree of market power in AI is probably necessary and beneficial, because the costs and risks of frontier AI development require the ability to fund large, long-term projects. But if that power becomes entrenched rather than contestable, it creates the familiar problems of monopoly: reduced innovation incentives, suppression of competition, and extraction of rents. Agentic collusion, where AI systems coordinate pricing strategies in ways that no human explicitly programmed them to do, is already producing legal cases in rental and gasoline markets.

His closing argument was that AI's effects on human flourishing are not automatic or stable. They depend entirely on the institutional framework within which the technology operates. Too little policy leaves harms unaddressed. Too much poorly designed policy creates friction that prevents beneficial uses from developing. The path to a good outcome requires continuous monitoring, rapid learning, and genuinely adaptive governance, not a one-time legislative settlement. The ESDC model Okonny had described, he noted, was a practical example of exactly this kind of agile, iterative approach.

Session 6: Panel -- The Limits of AI Governance: Innovation, New Big Tech and Trust

Moderator: Erik Bohlin, Ivey Business School

Panelists: Kevin Chan (Meta), Matthew da Mota (Canadian SHIELD Institute), Jonathan Obar (York University), Philippe Lefebvre (FIPRA, Brussels)

Setting the Panel's Agenda

Moderator Erik Bohlin framed the panel around the multidimensional nature of AI trust, a concept that had surfaced repeatedly in earlier sessions but had not yet been examined from the inside. He identified several distinct dimensions: privacy and data use, safety and accountability for AI-related harms, and the labor market anxieties that AI is generating across the economy. A recent ruling by Canada's Privacy Commissioner against OpenAI's data practices in Canada was fresh context. With that backdrop, he invited each panelist to offer a perspective.

Kevin Chan: Trust Through Demonstrated Value

Kevin Chan, Public Policy Director at Meta, argued that trust in AI is most effectively built through concrete demonstrations of positive impact rather than through abstract regulatory assurances. He offered several examples from his team's work facilitating AI adoption for social benefit.

In healthcare, Meta has worked with hospitals where AI is providing clinicians with a second opinion on medical diagnoses. The goal is not to replace physicians but to reduce diagnostic error rates; measurable outcomes confirm that the AI-assisted process produces safer and more accurate results. In a separate initiative, AI is being used to match cancer patients with clinical trials in ways that human review would likely miss, identifying patterns in complex data that point toward more effective treatment candidates.

On language preservation, Chan described Meta's partnership with UNESCO to develop online translation tools supporting low-resource Indigenous languages. These are languages with sparse training data -- many have not developed vocabulary for modern concepts -- and conventional approaches to natural language processing struggle with them. AI's zero-shot capabilities, where a model trained on language generally can make reasonable inferences about a language it has never specifically encountered, offer a pathway to preservation that would otherwise be practically impossible.

Chan also discussed Meta's open-source AI strategy. Rather than keeping foundation models proprietary, Meta publishes model weights, allowing developers, researchers, and governments worldwide to use and adapt them without API subscriptions or commercial agreements. He argued that this democratizes access to frontier AI capabilities, particularly for middle-power and developing countries that cannot afford to build foundation models from scratch. It enables those countries to build applications on top of state-of-the-art models while maintaining some degree of sovereign control over their own deployments. He acknowledged that this is not a complete answer to sovereignty concerns, but suggested it is a meaningful step toward broader access.

On the question of AI wearables, Chan's team is running an AI Glasses Impact Grant program, inviting applicants to propose beneficial uses for Meta's AI glasses technology, with seed funding for the most promising ideas. He described overwhelming demand for the program and noted that many of the most compelling proposals came from organizations with no commercial relationship to Meta.

Philippe Lefebvre: European Lessons on Governance and Data

Philippe Lefebvre, drawing on his long career at the European Commission and his experience watching the EU AI Act take shape, offered a more cautionary perspective. His overarching message was that the EU's experience illustrates both the importance of governance and the difficulty of executing it well.

On market concentration, Lefebvre argued that AI follows the same dynamics of increasing returns and network effects that have produced platform monopolies in digital markets over the past two decades. Europe, he said, has been effectively colonized by U.S. platforms for a decade, and AI risks deepening that dependence. The EU's Digital Markets Act attempts to recreate competitive conditions, and AI regulation will likely mandate that platforms not give preference to their own AI systems over third-party services. But he was candid that the EU's track record on actually reversing concentration is unproven.

On trust specifically, Lefebvre noted that even as the AI Act represents an ambitious governance framework, it is already facing implementation pressures. The EU AI Office has roughly 100 staff, paid at public-sector rates that are a fraction of what Silicon Valley engineers earn. The practical

result is that what looks like independent regulation may function more like structured self-regulation, with the regulator observing rather than genuinely auditing. Germany recently secured a complete exemption of its machinery sector from AI regulation under direct pressure from the Chancellor. A broader simplification initiative called the AI Omnibus is modifying more than 20 pieces of existing legislation simultaneously, and many AI-related obligations are already being delayed.

His most distinctive contribution was on data as the strategic differentiator. Lefebvre argued that AI model architectures will likely converge asymptotically in performance over time. The lasting competitive advantage will come from data, not algorithms. Europe, he argued, is poorly positioned because it has given much of its data away without monetizing it. The EU's Data Governance Act is an attempt to address this, making public data accessible by default and requiring companies to share user data with third parties under consent conditions, trying to unlock the latent value in data held across European institutions and industries. Whether this is sufficient to close the data gap with the U.S. and China remains an open question.

Matthew da Mota: Trust, Governance, and Sovereignty as a Triangle

Matthew da Mota from the Canadian SHIELD Institute presented trust, governance, and sovereignty as three mutually reinforcing concepts rather than separate concerns. His framing drew on research he had led across disciplines on what trust in human-machine interaction actually means.

The research found that trust is earned through a relationship between a more vulnerable party and a less vulnerable institution, that it must be continuously verified rather than assumed, and that the metrics engineers use to measure trustworthiness do not necessarily map onto the experience of trust that actual users have. A system can perform excellently on benchmarks and still command no trust; a system can perform poorly and still attract strong user confidence. High performance does not automatically produce legitimate trust.

Da Mota was pointed in his skepticism of claims that AI systems are democratic simply because they are accessible or open source. Democratic governance involves checks and balances, accountability to affected communities, and the genuine ability of those communities to say no to a technology or practice they find unacceptable. Private companies, however well-intentioned, cannot provide these structures. He raised the example of AI glasses as a specific concern: the device's camera can identify and record people who are in the vicinity of the wearer without their awareness or consent. The question of whether governance frameworks will ever reach a point where someone with authority says that this particular capability is not acceptable -- and whether that authority exists outside of the companies themselves -- remains unresolved.

On sovereignty, da Mota offered a direct definition: the ability of a country to exert governance and law over its markets, systems, and digital infrastructure. By this definition, Canada does not currently have digital sovereignty over its AI ecosystem in any meaningful sense. Laws exist on the books, but the ability to actually enforce governance over the systems that Canadians interact with daily remains severely limited. He argued that law, iterative community engagement, and third-party verification of standards are all necessary components of the trust architecture Canada needs to build. Standards bodies, he noted, are also heavily captured by large technology companies, which makes third-party verification harder than it sounds.

Jonathan Obar: Trust as a Concept Worth Questioning

Jonathan Obar of York University offered the panel's most skeptical intervention, challenging the framing of trust itself as a governance goal. His argument was that the word trust is being used in two very different ways in AI policy discourse, and conflating them creates serious problems.

In one usage, trust means something like compliance or non-resistance: users accept terms of service, proceed with transactions, and do not complain. Obar described this as the practical outcome of what researchers call digital resignation, online apathy, and privacy fatigue -- states where people have decided that reading privacy policies is not worth their time, that the services are too convenient to abandon, and that complaining accomplishes nothing. His empirical research at York, funded by Canada's Privacy Commissioner and documented at BiggestLieOnline.com and AIinfocus.net, found that people want to enjoy digital services without being impeded by the conditions attached to them. They click through consent screens, not because they trust the company, but because they have stopped trying to resist.

In the second usage, trust means something much more demanding: meaningful consent, genuine transparency about risks rather than just data routing practices, accountability structures that give users real recourse, and the ability to contest how their data is used. This is the vision of the Office of the Privacy Commissioner of Canada, and it bears little resemblance to what currently exists.

Obar argued that if policymakers are pursuing the first kind of trust, i.e. a population that does not object, engages without friction, and keeps using services, then what they are really pursuing is hegemony by a dominant group of AI stakeholders proceeding, as Microsoft researcher Kate Crawford has argued, without context, caution, or consent. His policy prescription was deliberately simple: governance frameworks should start from context, caution, and consent; understanding the specific situations in which AI operates, applying genuine caution rather than defaulting to permissiveness, and ensuring that consent is meaningful rather than a checkbox. He also noted that transparency efforts tend to communicate infrastructure details rather than risks, and that his research suggests video-based disclosures are significantly more effective than the text-based privacy policies nobody reads.

Session 7: Panel -- Perspectives by Users

Moderator: Romel Mostafa, Director, Lawrence National Centre for Policy and Management, Ivey Business School

Panelists: Jennifer Curtiss, Chief Data Officer, Scotiabank; Michael Page, Interim Senior Director, Data Science and Advanced Analytics, Unity Health Toronto

Opening: What Does Governance Mean in Practice?

Romel Mostafa opened the panel by carrying forward a question that had already surfaced in the preceding discussion: what does AI governance actually mean when you have to live with it every day? He noted that both panelists operate in sectors defined by trust, accountability, privacy, and security, and asked them to ground the concept operationally before moving to the strategic level.

Jennifer Curtiss, Chief Data Officer at Scotiabank, was direct. Governance, to her, means having the right guardrails in place to understand the risks that AI introduces and to ensure those risks remain within the bank's tolerance levels. Scotiabank operates a layered framework of controls, risk assessments, key risk indicators, and key performance indicators, with steering committees whose reporting lines run all the way to the board. The bank has a responsible AI function,

formerly called data ethics, that looks specifically at bias and transparency. It has separate controls for security, for privacy, and, as agentic AI becomes more prevalent, for the new categories of risk that agentic systems introduce. Curtiss acknowledged that much of this work has historically been manual, and the bank is now in the process of automating its governance platform to keep pace with the volume of AI deployments without creating friction that slows legitimate innovation. Her goal is to automate the assessment of low-risk use cases so that teams can move quickly through a sandbox and proof-of-concept stage, while reserving full human review for high-risk deployments, particularly those involving third parties.

Michael Page offered a parallel framing from healthcare. Unity Health Toronto, the network that includes St. Michael's Hospital, operates in what Page described as an almost complete regulatory vacuum. There is no overarching standard that each hospital or care provider must follow for AI. His team has therefore built its own governance architecture from the ground up, assembling a review panel that includes a legal expert, an ethicist, an IT security specialist, and a privacy officer to assess every AI use case before deployment. Page was candid that this approach was novel when Unity Health adopted it in 2017. At the time, the idea that data scientists should not be the only people making AI risk decisions in a healthcare setting was not widely shared. Most institutions were looking only at technical benchmarks. Page's framing was deliberately broader: governance for him means ensuring that an AI system is safe, robust, and treats the community the way he would want to be treated himself.

Sovereignty in Practice: On-Premise Infrastructure and the Cloud Act Problem

Mostafa turned the conversation to data sovereignty and infrastructure. He noted that both finance and healthcare handle some of the most intimate and consequential data that exists, and asked how each institution navigates the tension between sovereign control and access to cutting-edge capabilities.

Page's answer was striking in its bluntness. Unity Health's entire data infrastructure is on-premise, not because of a philosophical commitment to sovereignty, but because the budget never allowed cloud migration. As of 2017, moving patient data to cloud platforms felt too risky and too uncertain, so the hospital made capital investments in local compute instead. The institution now runs over a dozen H100 GPUs on-site and has deployed more than 60 applications into clinical practice, with over 90 percent of that work funded by philanthropy rather than government. Page acknowledged that this path creates real constraints. The hospital cannot easily access the frontier models and platforms available to better-resourced institutions. But the upside has become clearer in recent years: Unity Health still owns its data physically, something that many healthcare organizations that moved earlier to cloud providers now struggle to say with confidence.

The U.S. CLOUD Act was mentioned as the reason that confidence is fragile for those who did migrate. Even if a cloud provider stores data in a Canadian data center, if the company is incorporated in the United States, U.S. authorities can compel disclosure of the data under the CLOUD Act. For health data, which includes information about pregnancies, mental health conditions, HIV status, and cancer diagnoses, Page argued this is not a technical risk but a fundamental breach of the relationship of trust between a hospital and its patients. He expressed genuine relief that Unity Health's budget constraints had, inadvertently, been the right call.

Curtiss approached the sovereignty question from a different angle. Scotiabank operates in Mexico, Canada, and the United States, and has direct experience with what data localization requirements mean in practice. When Mexico introduced data localization rules, the bank built an on-premise deployment in Mexico to comply, building in the same controls it would apply

anywhere else. Her view is that a fragmented, country-by-country patchwork of data rules is operationally difficult and produces inconsistent customer experiences across geographies. She would prefer a coherent global standard for how customer data should be treated and protected, applied consistently, over mandating physical data infrastructure in each jurisdiction. She acknowledged that this is a hard political sell, but argued it would ultimately produce better outcomes for customers.

AI Literacy, Consent, and the Limits of Governance Without Education

A recurring theme across both the formal discussion and audience questions was the relationship between AI governance and AI literacy. Page was the most direct in naming the scale of the problem. A survey conducted by KPMG found that Canada ranked 38th out of 40 advanced economies on AI literacy, and that the survey drew an explicit connection between literacy and trust: Canadians do not trust AI, and they trust it significantly less than their counterparts in the United States, Asia, and Southeast Asia. Page argued that this is not a peripheral issue. An AI system that clinicians do not understand, do not trust, and do not integrate into their workflow delivers no benefit regardless of how well it performs on technical benchmarks. He described a bottom-up approach to AI adoption at Unity Health: rather than pushing AI tools down to clinical staff, his team goes to the floor, identifies the specific problems that nurses and physicians lose sleep over, and brings AI to solve those problems. This is, he noted, the only adoption strategy that actually works.

In response to an audience member who pressed on the consent problem, Curtiss was frank. Requiring users to consent to data use as a condition of accessing a service, and then framing that consent as voluntary, is not a workable regulatory model. The cookie consent experience has demonstrated this conclusively. The problem is not that people are uninformed; it is that they have no real choice. Requiring consent in exchange for access is, as she put it, simply not going to produce what regulators want. She argued that a complete rethink of how consent functions in digital services is needed, and that incremental adjustments to existing frameworks are unlikely to close the gap. For its part, Scotiabank is attempting to automate much of its own governance layer so that the burden of compliance does not fall on customers or on individual business units, and so that the bank can maintain consistent standards even as the volume and variety of AI deployments grows.

Innovation, Risk Thresholds, and the Role of Change Management

On the balance between governance and innovation, both panelists arrived at similar conclusions from different directions. Curtiss described a staged model where low-risk use cases can move quickly through a sandbox to proof-of-concept without full governance assessment, while anything moving into production triggers a complete review. The goal is to ensure that governance does not become an obstacle to experimentation, while ensuring that what reaches customers has been properly scrutinized. She used a concrete example: her supervisor had fed a spreadsheet containing 142 governance items into Microsoft Copilot and received an output claiming there were 379 no-go cases. The mismatch illustrated exactly the kind of uncritical AI use that governance frameworks are designed to prevent. Page noted that in healthcare, the motivating force for AI adoption is not cost savings but crisis. The healthcare system is visibly struggling, and clinicians know it. That sense of urgency, rather than abstract efficiency arguments, is what makes AI adoption feel both necessary and worth the effort of doing carefully.

An audience member asked whether either panelist had considered involving change management specialists with expertise in adoption architecture, noting that most major consulting firms now cite change management as accounting for 70 percent of the success or failure of AI

transformations. Guilbault, joining from the adjacent panel, seconded the point emphatically. Page agreed that building the best model in the world produces no benefit if no one uses it, and described several cases where Unity Health had developed technically strong solutions that achieved minimal adoption simply because the rollout had not been designed around the end user. The lesson both panelists drew is that governance and adoption are not separable problems. A system that is technically safe but organizationally rejected has failed just as surely as one that is adopted but causes harm.

Session 8: Panel -- Perspectives by Service Providers

Moderator: Romel Mostafa, Professor and LNC Director

Panelists: Mark Graham, Senior Vice President, Legal and Regulatory Affairs, BCE; Alexandre Guilbault, Vice President, AI Enablement, TELUS

Framing: From Consumer Privacy to Infrastructure Risk

Romel Mostafa opened the panel by noting that the conversation was shifting registers: from how enterprises use AI to how Canadian service providers are building the infrastructure that will make AI possible at national scale. He invited both panelists to address what AI governance and data sovereignty mean from an infrastructure provider's perspective.

Mark Graham, Senior Vice President for Legal and Regulatory Affairs at Bell, opened with a reframing of the risk model. He argued that the mental model Canadians have used for the past 15 years to think about AI risk, which is primarily a model built around privacy and the practices of large U.S. consumer platforms, is no longer the right lens for enterprise AI. Privacy remains important, but when organizations begin running critical operational workflows on AI systems, a new and more serious category of risk emerges: the risk that a foreign entity can turn the system off. He offered a series of escalating examples. A government agency running its payroll through an AI-enabled workflow that depends on a U.S. platform could find itself unable to pay employees. An enterprise could lose its customer service capability. Trains could stop running. Planes could be grounded. These are not hypothetical risks confined to adversarial scenarios; they are the logical consequence of allowing critical national infrastructure to depend entirely on systems whose ultimate control lies outside Canada's borders.

Alexandre Guilbault, Vice President of AI Enablement at TELUS, made a parallel argument from the perspective of compute dependency. AI is following the same trajectory as electricity and telecommunications: once pervasive enough, it becomes infrastructure rather than technology. Just as no one would accept a scenario in which Canada's electricity grid or telecom backbone was wholly controlled by a foreign actor, the same logic must eventually apply to AI compute. The practical concern Guilbault identified is that even developers who are highly skilled today are increasingly using AI tools for coding and other core tasks. Research already shows that if AI tools were removed, many developers would struggle to maintain the codebases they have built with AI assistance. Dependency of this kind will only deepen over the next two to three years.

The Full-Stack Argument for Sovereign AI

Both Graham and Guilbault converged on the same structural argument: meaningful AI sovereignty requires owning the full stack, not just isolated pieces of it. Graham described that stack in concrete terms: a sovereign data center physically located in Canada; compute

infrastructure in the form of owned or long-term-leased GPUs; a platform layer that abstracts the hardware for practical use; Canadian-controlled storage; a domestic telecommunications network that transmits data without routing through the United States; and a governance layer controlling access, identity management, and encryption keys. Any gap in that stack, he argued, undermines the sovereignty of the whole. This is not primarily a privacy argument, though privacy is one dimension; it is an economic and national security argument about retaining control over systems that will become as fundamental to daily life as power or water.

Guilbault added an important legal point. Even if a company stores data in a Canadian data center, if that company is incorporated in the United States, the U.S. CLOUD Act allows American authorities to compel disclosure regardless of where the data physically resides. True sovereignty therefore requires not just physical location in Canada but Canadian ownership throughout the stack. He noted that TELUS has built diverse transport paths across Canada that do not route through U.S. networks, meaning that data moving between TELUS facilities never crosses into U.S. jurisdiction on the transport layer. This was presented not as a theoretical position but as an operational reality already in place.

Bell has committed to building more than 300 megawatts of sovereign Canadian data center capacity, funded entirely without government subsidy. TELUS recently opened its first sovereign AI factory in Rimouski, Quebec, with 2,000 H200 GPUs running on 99 percent renewable energy, and the facility sold out quickly. A subsequent announcement at the conference itself covered three new AI factories in British Columbia, totaling 60,000 next-generation GPUs, also powered by BC Hydro renewable energy, with waste heat being used to warm approximately 150,000 homes in the Vancouver area. Guilbault described these as carbon-neutral facilities and positioned Canada's combination of cold climate, abundant renewable energy, stable democracy, and established AI talent base as a genuinely competitive advantage in the global market for sovereign AI infrastructure.

The Role of Government: Clarity and Anchor Customers

The panel addressed what role government policy should play in accelerating the buildout of sovereign AI infrastructure. Both panelists were careful to distinguish between the kinds of government action that would be helpful and the kinds that would create friction.

Guilbault argued that regulation needs to provide clarity rather than comprehensiveness. Technology moves faster than legislation, and attempting to regulate every technical dimension of AI in real time will produce rules that are outdated before they come into force. What enterprises need most is a clear definition of what sovereign AI means in the Canadian context, specifically which workload types must remain within Canadian-owned infrastructure. Defense data, critical IP systems, and health data were offered as obvious candidates. Once enterprises know which of their workloads are subject to sovereignty requirements, they can make long-term infrastructure commitments with confidence. Without that clarity, they remain in testing mode rather than committing to full deployment.

Graham identified a second role for government that goes beyond regulation: anchor customer. The economic viability of building out a full-stack sovereign AI ecosystem in Canada depends, in part, on there being sufficient demand to justify the investment. Government is uniquely positioned to create that demand by procuring Canadian sovereign AI solutions for its own operations, particularly in defense-adjacent and critical public services. This would give domestic providers the commercial anchor they need to scale, which in turn makes the solutions available to private-sector enterprises at competitive cost.

On the question of Canadian providers competing globally, both panelists expressed genuine optimism. Guilbault framed it as a coalition-building opportunity: middle powers working together, including Germany, other EU member states, and Asian democracies, can collectively offer an alternative to dependence on U.S. or Chinese infrastructure. Canada's stable democratic institutions, clean energy infrastructure, and internationally recognized AI research community give it meaningful standing in that coalition. Bell's ecosystem approach, described by Graham, involves partnerships with Canadian hardware companies including HyperTech, Celestica, and Buz for compute and storage, and a partnership with Cohere for the model and agentic platform layer. The goal is a fully integrated offering that an enterprise can adopt without having to stitch together individual sovereign components on its own.

On the question of jobs, Guilbault pushed back gently on anxious framing. TELUS alone announced a thousand construction jobs from the new facilities on the day of the conference. Running and maintaining the data centers will require hundreds more. On the broader labor market question, his view, drawn from ten years of AI deployment experience, is that AI has consistently transformed work rather than eliminated it. No one at TELUS has been replaced one-for-one by an AI system; AI has made people more capable of doing more and better work. He acknowledged the counterargument and offered a clean formulation: it is not AI that replaces workers, but workers using AI who will replace those who do not. The analogy he used was computing. Refusing to use a computer did not preserve a job; it simply made someone less competitive than a colleague who embraced the tool.

Closing Keynote: AI Diplomacy for Mid-size Powers -- Implications for Canada

Speaker: Mark Daley, Chief AI Officer, Western University

Opening: A Computer Scientist in a Realist World

Mark Daley introduced himself as a theoretical computer scientist who was pulled into science diplomacy somewhat against his will, and who has since found it to be among the most consequential work he has done. He is currently NSERC Scholar in Residence in AI, has written on international relations for policy outlets, and was a driving force behind the R7+ research leadership summit that Canada convened in Ottawa in November 2025 as part of its G7 chair year.

He opened with a blunt diagnosis of the geopolitical context. Canada's Prime Minister had recently spoken at Davos, reiterating a message from a piece published in The Economist: the rules-based liberal order is over, or at minimum under serious threat, and institution-building in the traditional sense is no longer a reliable path forward. Daley accepted this premise and built his talk around a single question: in a realist world, what does effective AI diplomacy for a middle power actually look like?

He also offered a counterpoint to a claim made by an earlier panelist, specifically the formulation that AI will not replace workers but that workers using AI will replace those who do not. Daley disagreed. He argued that humanity has done something with AI that has never been done before: commoditized intelligence itself. Not replaced human beings, but made intelligence cheap, scalable, and rentable from the cloud. Every previous source of intelligence, he noted, has been another human or, charitably, a reasonably clever cat. Getting access to large amounts of intelligence has always required hiring, managing, and communicating with people, at significant

cost and overhead. AI changes that equation fundamentally, and he predicted that the full consequences of that change would be visible within three years.

AI Is Not a Ghost: Understanding the Full Stack

Daley pushed back on what he described as a tendency in popular discourse to think of AI as a set of magic weights, a ghost in the machine. AI is not a model. It is an entire physical and logical stack, and governance must address every layer of that stack. He enumerated those layers: critical minerals, energy and physical infrastructure, chips and semiconductor hardware, system software and platforms, foundation models, and applications. Canada's strategic position differs substantially across each of these layers, and the country's diplomacy must be calibrated accordingly.

One unexpected constraint Daley flagged from his own experience as a researcher in neural computation: the binding bottleneck on progress in AI infrastructure is not clever graduate students or algorithmic breakthroughs. It is the inability to build power substations fast enough to energize data centers. Canada's abundant hydroelectric capacity, already discussed by the telecom panelists, is therefore not merely an environmental asset but a geopolitical one. It is one of the few layers of the AI stack where Canada has a genuine, hard-to-replicate advantage in the near term.

A Framework for Middle Power Strategy: Dual Use and Choke Points

The analytical core of Daley's talk was a two-axis framework he has developed and published, which he uses to map different types of innovation and determine the appropriate openness posture for each.

The first axis measures how dual-use a given technology or area of innovation is: that is, how readily it can be converted from civilian to military or strategic application. The second axis measures whether a country has a choke point: a capability so technically difficult to replicate that others must come to you for it. Fundamental astrophysics sits in the low-dual-use, no-choke-point quadrant. Anyone with a telescope can contribute, there is no strategic advantage to keeping discoveries secret, and international scientific cooperation is the appropriate posture. This is the domain of the traditional scientific commons.

AI chip fabrication equipment sits at the opposite extreme: extremely high dual-use salience and a massive choke point. The Netherlands, through ASML, controls the only technology capable of producing the extreme ultraviolet lithography machines that TSMC uses to fabricate the most advanced GPU chips in the world. No other company in the world can replicate what ASML does at the required scale and precision. This gives the Netherlands extraordinary leverage, and it must govern the use of that leverage carefully.

Canada does not have a choke point equivalent to ASML anywhere in the AI stack. This matters because it shapes what strategic options are actually available. Daley argued that for a middle power without a dominant choke point, the primary lever is aggregation: forming coalitions with like-minded nations to create collective market access conditions that even dominant AI powers cannot afford to ignore. If Canada, Japan, South Korea, and the EU jointly set standards for AI systems that must be met before those systems can be sold into their markets, they represent a combined market large enough to compel compliance. No single country in that group has the leverage individually; together, they do.

Variable Geometry and Coalition Building

Daley introduced the concept of variable geometry, which he attributed to work by Kearney and which Canada's Prime Minister had recently invoked at Davos. Variable geometry is the idea that in a multipolar world without functioning universal institutions, coalitions will form and dissolve issue by issue rather than being fixed alliances. Different countries will cooperate on different layers of the AI stack, and Canada's partners at the energy layer may not be its partners at the standards layer, which may differ again from its partners at the talent layer.

He illustrated this with a concrete example. The Netherlands has an extraordinary chip fabrication capability but limited green energy infrastructure. Canada has the opposite profile. A deal in which a Dutch-designed supercomputer is built and operated in northern Quebec, powered by hydroelectricity, in exchange for Canadian access to Dutch semiconductor technology, would be mutually beneficial and would advance sovereignty goals for both countries across different layers of the stack simultaneously. These kinds of cross-layer trades are, Daley argued, exactly the form that practical AI diplomacy should take.

Canada convened the R7+ summit, bringing together research leaders from the G7 nations plus Spain and the EU, and the group reached three agreements, one of which established a shared talent mobility hub allowing researchers to move among participating countries more easily. This is a working example of the coalition-building model Daley advocates: a small, targeted agreement among like-minded partners on a specific layer of the stack, verified and enforced within the club, and designed to accumulate value over time.

Standards Statecraft and the Risk of Passive Subsidy

Daley identified standards as the most underused tool in Canada's AI diplomacy toolkit and the one with the highest return on investment given Canada's existing reputation. Canada is recognized globally as a first mover in AI research. Geoffrey Hinton and Yoshua Bengio, two of the three recipients of the Turing Award for deep learning, did their foundational work in Canada funded by NSERC and CIFAR. This history confers genuine credibility and prestige in international technical forums. Daley argued that Canada should be placing people of that caliber onto international AI standards committees far more aggressively than it currently does. Standards are, as he put it, genuinely boring but actually quite powerful. They are the mechanism by which technical requirements get embedded into procurement processes, regulatory frameworks, and supply chain expectations across dozens of countries simultaneously.

He also flagged two risks that he argued deserve more attention from policymakers. The first is subsidy without value capture: the danger that Canada invests public funds in attracting AI infrastructure or research, only for the resulting economic value to flow primarily to foreign companies or foreign shareholders rather than back into the Canadian economy. He was careful not to be overly specific, but the implication was that the design of incentive programs matters as much as the amount of the incentive. The second risk is the capture of public infrastructure for private ends: allowing publicly funded research or publicly owned physical infrastructure to become the foundation for private AI capabilities without appropriate conditions on that use.

Sovereignty as Optionality

Daley closed with a reframing of sovereignty that he described as the most important conceptual move in his talk. Sovereignty, he argued, is not a binary condition and it is not about autarky. Canada will not fabricate its own advanced chips in the next five years; that would require a generational industrial program competing directly against the most heavily capitalized companies in the world, and even then success would not be guaranteed. Treating sovereignty as an all-or-nothing proposition leads either to paralysis or to poorly designed programs that waste public resources chasing objectives that are not achievable.

The right frame, Daley proposed, is optionality. Sovereignty means preserving as many choices as possible across as many layers of the AI stack as possible, so that Canada is never forced into a position where a foreign actor can unilaterally change the terms of access to something critical. Where Canada has genuine advantages, such as energy, talent, and democratic stability, it should invest aggressively to deepen and extend those advantages. Where it does not, it should seek trusted partners and negotiate arrangements that preserve at least some leverage rather than accepting full dependence.

His closing prescription for Canada as a middle power was a five-point summary: understand the full AI stack and update that understanding as the technology evolves; aggregate with like-minded and like-valued partners but verify their behavior rather than simply trusting; engage actively in standards statecraft and put Canada's best technical minds onto the relevant international committees; use procurement as a diplomatic and industrial policy tool; and protect the scientific commons, ensuring that concerns about dual-use technologies do not inadvertently push Canada away from the fundamental AI research that produced its international standing in the first place. Canada's goal, he concluded, is not to win a race it cannot win, but to remain a country that has real choices about how it participates in the AI era.