Canada Doesn't Need to Win the AI Race. It Needs to Run a Different One.

Canada's new AI strategy gets the destination right. Whether we arrive depends on a transition inside our firms that the strategy never mentions — and that too few leaders can see.

The number that matters

On June 4, 2026, Prime Minister Mark Carney and Evan Solomon, Canada's first Minister of Artificial Intelligence and Digital Innovation, launched *AI for All*, the federal government's renewed national AI strategy.[1] It arrives with serious money — between $2.5 to 3.5 billion in committed and near-term investment — and one target that matters more than all the others: raising business AI adoption from roughly 12 per cent today to 60 per cent by 2034.[2]

That is the right target. Artificial intelligence is a general-purpose technology, and the history of such technologies — electrification, computing — is that the countries that benefit most are not necessarily those that invent them, but those that adopt them most effectively. Development economists call this infusion: the unglamorous process by which new technology spreads through firms and value chains, raising productivity along the way.[3] Canada's overarching economic problem is not invention, per se; it is productivity. On this much, the strategy and we agree.

What the national AI strategy leaves underdeveloped in our view are the mechanisms of adoption: How, concretely, does a mid-sized Canadian firm get from "our staff have AI training" to "our workflows run on AI"? We write as academics who study technology diffusion and spend much of our time working with executives wrestling with real AI implementations. From both vantage points, the missing middle — the bridge between the strategy's Pillar 2 (skills) and Pillar 3 (adoption) — is where Canada's productivity dividend will actually be decided.

Realism about the race

First, a word about the pillars we would de-emphasize somewhat. At Davos this January, the Prime Minister argued that middle powers are "not powerless," and described his approach as values-based realism: principled about ends, unsentimental about means.[4] Canada's AI debate needs the same honesty. Frontier model training now costs billions per run, and the relevant competition is between trillion-dollar American platforms and the Chinese state. The "bitter lesson" of modern AI, in Richard Sutton's phrase, is that capability tracks compute and capital.[5] That precursors to modern AI were in a real sense born in Canadian universities is a fact about history, not about who can afford to push the frontier of the technology today.

This is an argument about emphasis. Parts of Pillar 4 — notably the Compute Access Fund's $700 million in affordable sovereign compute for SMEs — are adoption infrastructure by another name, and we welcome them. But sovereignty investments only pay off where

absorption capacity exists. The sequencing question matters, and in our view, the strategy's centre of gravity should sit squarely on Pillars 2 and 3.

Why the bridge matters more in Canada

There is a distinctly Canadian reason underpinning our argument. Unlike fellow middle powers such as Korea and Japan whose economies are anchored by technically deep industrial conglomerates, Canada's is over-weighted toward small and mid-sized firms that few would describe as "tech." Organizations with fewer than 250 employees account for 69 per cent of Canadian employment — higher than Germany, France, or the United States.[6] Most such firms, even mid-sized companies that are sophisticated in their domains and have learned how to slot into global supply chains, have never hired a data scientist. Many are data-rich but lacking in the modern technical capabilities needed to leverage the data they have. If productivity does not rise in their back offices and on their shop floors, it will not rise nationally.

From chat to automation: the bridge nobody talks about

In our work with Canadian firms seeking pathways to AI adoption, we have noticed a frequent stumbling block: the transition from chat interfaces to true automation. Chat is where organizations learn: giving every employee a license to an AI chat-interface creates a sandbox in which non-technical staff experiment on real problems. We encourage this approach, at least early on. Getting employees comfortable with AI, and having them find ways to leverage it for personal productivity, are positive developments.

But big productivity gains do not come through chat windows. They come from workflow automation and more deeply transformational approaches that motivate firms to change existing processes. Some of these are on the revenue side of the business — finding new growth opportunities, winning competitive bids, and pricing more intelligently. Many are on the cost side — driving workplace efficiencies, improving logistics, and modernizing procurement. Some of these gains are born from prompts, of course, polished and stress-tested in daily use. But the big wins come from migrating successful prompts into automated pipelines through API calls or locally run open-source models. The chat window is where you learn; the API is where you scale.

What strikes us most in our work with companies is how few senior leaders know this gulf exists. Many equate "adopting AI" with buying licenses and mandating training — Pillar 2 thinking — and sincerely believe they have "done AI." The predictable result is pockets of enthusiastic prompting that never aggregate into redesigned workflows or measurable cost curves. This is the difference between using AI and changing because of it; it is the difference between adoption and infusion.

What this means for the strategy

If the bridge from individual-level prompting to scaled, API-driven automation is where adoption succeeds or fails, federal programs should be built around it. Adoption funding — the

$500 million LIFT program, and the Regional AI Initiative — should reward workflow redesign, not license counts or training completions. And the National AI Literacy Initiative needs an executive tier: not prompting classes, but API economics, model selection, and data classification for the CFOs, CIOs, and general counsel who often stand, mostly unknowingly, as gatekeepers on the bridge.

None of this is glamorous. Sovereign compute and national champions offer prestige; the bridge offers productivity. A middle power whose economy runs on mid-sized, non-technical firms, and whose prime minister has made realism respectable again, should know which race it is running. Canada doesn't need to win the AI race. It needs to run a different one — and this time, actually finish it.

[AI tools were used in the research preparation for this piece.]

Footnotes

1. Prime Minister of Canada, "Prime Minister Carney launches AI for All: Canada's new national artificial intelligence strategy," news release, Toronto, June 4, 2026. https://www.pm.gc.ca/en/news/news-releases/2026/06/04/prime-minister-carney-launches-ai-all-canadas-new-national-artificial ↩
2. Innovation, Science and Economic Development Canada, *Canada's National Artificial Intelligence Strategy: AI for All* (Ottawa: Government of Canada, 2026). https://ised-isde.canada.ca/site/ised/en/canadas-national-artificial-intelligence-strategy-ai-all ↩
3. World Bank, *World Development Report 2024: The Middle-Income Trap* (Washington, DC: World Bank, 2024). ↩
4. Mark Carney, "Principled and Pragmatic: Canada's Path," Special Address, World Economic Forum, Davos, January 20, 2026. Full text: CBC News, https://www.cbc.ca/news/politics/mark-carney-speech-davos-rules-based-order-9.7053350 ↩
5. Richard S. Sutton, "The Bitter Lesson," March 13, 2019. https://www.cs.utexas.edu/~eunsol/courses/data/bitter_lesson.pdf
6. OECD, *OECD SME and Entrepreneurship Outlook 2023* (Paris: OECD Publishing, 2023). https://doi.org/10.1787/342b8564-en ↩