

JUNE 2026

# AI INFRASTRUCTURE AND DATA CENTRE RISKS FOR CANADIAN FINANCIAL INSTITUTIONS

Author

**Chris Collins**, Cascade Institute



GLOBAL  
RISK  
INSTITUTE



CASCADE  
INSTITUTE

# AI Infrastructure and Data Centre Risks for Canadian Financial Institutions

June 2026

Prepared by the Cascade Institute for the Global Risk Institute

© 2026 Chris Collins. This “AI Infrastructure and Data Centre Risks for Canadian Financial Institutions” is published under license by the Global Risk Institute in Financial Services(GRI). The views and opinions expressed by the author are not necessarily the views of GRI. “AI Infrastructure and Data Centre Risks for Canadian Financial Institutions” is available at [www.globalriskinstitute.org](http://www.globalriskinstitute.org). Permission is hereby granted to reprint the “AI Infrastructure and Data Centre Risks for Canadian Financial Institutions” on the following conditions: the content is not altered or edited in any way and proper attribution of the author, GRI is displayed in any reproduction. ***All other rights reserved.***

# Contents

- Executive summary ..... 2
- Introduction ..... 4
- Background on AI ..... 5
- AI-System Related Risks for Financial Institutions ..... 8
- 1. Operational risks ..... 8
  - AI use and dependency across financial institution operations ..... 8
  - Operational risks from AI agents ..... 10
  - Cybersecurity risks ..... 10
  - Cloud infrastructure concentration ..... 11
  - Grid stress and backup power limitations ..... 12
- 2. Investment risks ..... 14
  - Direct and indirect AI investment exposure ..... 14
  - Market concentration: The “Magnificent Seven” ..... 15
  - AI capital expenditure and crowding-out effects ..... 16
  - The AI bubble question ..... 17
  - Data centre energy demand and grid stress ..... 19
  - Water consumption and cooling infrastructure ..... 20
- 3. Downstream risks from AI ..... 23
  - Regulatory and political risks from data centres ..... 23
  - Shifting AI governance and regulatory frameworks ..... 24
  - Geopolitical risks ..... 26
  - Economic and social risks ..... 28
  - Misinformation and disinformation risks ..... 31
  - AGI, ASI, and AI safety concerns ..... 32
  - Risk interaction and amplification ..... 34
- Conclusion ..... 36
- References ..... 37

## Executive summary

Artificial intelligence (AI) is advancing faster than any technology in modern history, and this has major implications for Canada’s financial sector. AI is reshaping how financial institutions operate, where they invest, and the broader economic and geopolitical environment in which they compete.

Critical to this shift is the physical infrastructure that enables AI. AI systems depend heavily on electricity, water, land, and specialized materials. Data centres are particularly energy-intensive, and their rapid expansion is placing increasing strain on power grids in North America and globally. In Canada, data centres could account for a significant share of electricity demand within the next decade, raising concerns about grid reliability, energy price volatility, and regulatory intervention.

While much of the focus has been on AI’s capabilities, its underlying infrastructure is emerging as a critical constraint and source of risk. In Canada, recent industry discussions through the Financial Industry Forum on Artificial Intelligence (FIFAI II) have highlighted that AI-driven risks are increasingly systemic in nature—cutting across operational, market, and consumer domains simultaneously and challenging traditional siloed risk management approaches.<sup>1</sup>

Against this backdrop, this paper explores three primary categories of AI-related risks facing Canadian financial institutions. These include:

**1. Operational risks** that arise from the growing dependency of Canadian financial institutions on AI systems and the physical infrastructure that supports this technology. These risks are compounded by the extraordinary concentration of AI infrastructure among a small number of American-owned cloud providers, and by the mounting strain that AI data centres are placing on electrical grids and the impact a disruption in these energy supplies could have for AI operations.

**2. Investment risks** that arise from the exposure of Canadian financial institutions to both AI-concentrated markets and to the physical resource constraints that may limit AI’s growth trajectory. Currently, the global AI investment landscape is defined by two features. The first is an extraordinary level of capital expenditure, and the second is significant market concentration in a small number of mega-cap technology firms such as the so-called “Magnificent Seven,” which collectively account for over 30% of the S&P 500. Investment portfolios of Canadian financial institutions carry correspondingly high exposure to AI and the broader technology sector, through both direct holdings and passive index strategies. This portfolio concentration creates a vulnerability to correlated losses in the event of an AI

downturn. Underpinning these financial risks are physical resource constraints on AI infrastructure, particularly around energy and water, that may limit the growth trajectory currently embedded in AI valuations.

**3. Downstream risks** are the second- and third-order effects of AI's proliferation across the economy, society, and the world's geopolitical order. These are the least visible and potentially most consequential dimensions of AI risk for Canadian financial institutions, and they are the risks for which institutions are least prepared. A defining feature of downstream risks is their ability to interact and amplify one another. For example, resource constraints affecting data centres could trigger regulatory action, shift market sentiment, and lead to asset revaluations, all while disrupting operational systems. Similarly, geopolitical shocks could simultaneously affect supply chains, market valuations, and regulatory environments. These interconnected risks challenge traditional risk management approaches that treat risks as isolated and independent.

In this context, managing AI-related risk is not solely an operational challenge; it is a strategic imperative. Institutions that develop a comprehensive understanding of AI's investment and downstream risks, particularly those linked to infrastructure and resource constraints, will be better equipped to sustain resilience and capture long-term value in an increasingly AI-driven economy.

## Introduction

Artificial intelligence (AI) is advancing faster than any technology in modern history. The scope and rapid pace of change from AI is unprecedented, as is the growth in users, which has surpassed the pace of the Internet's adoption.<sup>2</sup>

The financial industry is at the forefront of this change. AI is reshaping how financial institutions operate, where they invest, and the broader economic and geopolitical environment in which they function. This presents both significant opportunities and new risks.

For Canadian financial institutions, AI risks can be thought of in terms of **operational risks**, arising from growing institutional dependency on AI systems and the physical infrastructure that supports them; **investment risks**, arising from concentrated financial exposure to AI-driven markets and the physical resource constraints that may limit their growth; and **downstream risks**, arising from the second- and third-order effects of AI's proliferation across the economy, society, and the geopolitical order.

This paper examines these AI risks. At the operational level, the discussion covers risks arising from AI agents, the concentration of cloud infrastructure, and mounting pressure on energy grids. The investment level addresses market concentration, the unprecedented scale of AI capital expenditure, the possibility of a speculative bubble, and the physical resource constraints facing data centres, and how these may impact valuations of AI firms. At the downstream level, the paper considers regulatory and political risk, geopolitical disruption, systemic risks from AI agents, labour displacement, disinformation, and AI safety concerns. It also explores how risks at each level interact with and amplify one another.

Throughout, the paper also points to considerations relevant to assessing and managing these risks, such as risk interaction analysis, scenario planning, and stress testing of AI-linked exposures.

Given the rapid pace of AI development, the risk landscape will continue to evolve quickly. Accordingly, this paper does not attempt to predict precisely how AI will transform the financial sector—no one can do so with confidence. Instead, it argues that institutions that invest now in the necessary analytical capabilities, governance structures, and monitoring systems for integrated AI risk assessment will be far better positioned to navigate future challenges than those that wait for risks to materialize before acting.

## Background on AI

Artificial intelligence (AI) is a branch of computer science that enables machines and computer systems to perform tasks that typically require human intelligence, such as reasoning, problem-solving, or perception. Powered by algorithms, machine learning, and deep learning, AI technologies analyze large amounts of data to identify patterns, make predictions, and act autonomously.

Generative AI (Gen AI) is a type of AI that creates new, original content (such as text, images, code, music, and videos) by learning patterns from existing data. As IBM writes, “AI has been a hot technology topic for the past decade, but generative AI, and specifically the arrival of ChatGPT in 2022, has thrust AI into worldwide headlines and launched an unprecedented surge of AI innovation and adoption.”<sup>3</sup>

The capabilities of AI systems have advanced significantly in recent years, and AI performance across a wide range of benchmarks continues to improve.<sup>4</sup> AI models are being trained to break down problems and use step-by-step “reasoning,” and this is improving AI performance.<sup>5</sup> Most notably, over the past year, AI systems have made significant improvements in math, science, and software engineering.<sup>6</sup> AI tools are becoming especially commonplace in software development, with over 65% of developers reporting they use AI tools at least weekly.<sup>7</sup>

At present, both globally and in Canada, many organizations are using AI to improve processes, analyze data, and enhance customer experience.<sup>8</sup> AI has the potential to transform multiple aspects of business, society, and financial markets, but it also carries with it significant risks that remain poorly understood.<sup>9</sup>

Financial services are one of the sectors most suited to AI adoption, and, as AI develops and its capabilities improve, the technology is rapidly reshaping the sector’s strategic landscape. Currently, financial institutions are exploring a range of AI use cases to increase efficiency, reduce risk, and manage fraud.<sup>10</sup> Showing the scope of the potential impact of AI, according to new research from Toronto Metropolitan University, 98% of Canada’s approximately 800,000 financial sector workers are in occupations that are highly exposed to AI technologies.<sup>11</sup>

One particularly relevant development for the financial sector is the emergence of AI agents: autonomous systems that can plan, execute, and adapt multi-step tasks with minimal human oversight.<sup>12</sup> Major technology companies and AI laboratories are now deploying agents that can navigate complex workflows, interact with external systems, and make decisions in real time. AI agents are rapidly gaining traction across a wide range of business use cases, and the AI agent market is forecast to grow substantially in the coming years.<sup>13</sup> Globally, financial institutions are deploying AI agents across a wide range of use cases.<sup>14</sup>

Data centres are foundational to the development and use of AI technology. Data centres are facilities housing the technological infrastructure required to build and run both AI and non-AI digital applications.<sup>15</sup> Data centre facilities may provide AI or other digital infrastructure to multiple clients, or they can be operated on behalf of a specific firm.

As AI grows, the demand for computing power, commonly called “compute,” has also grown. This has led to the rise of specialized AI data centres, high-performance facilities designed specifically to train, deploy, and run AI models.<sup>16</sup> Unlike traditional data centres, which could be used for both AI and general-purpose computing, these dedicated “AI factories” leverage specialized hardware to process the vast datasets required for the training and operation of AI models and machine learning.<sup>17</sup>

Demand for data centres is primarily measured in megawatts (MW) or gigawatts (GW) of power consumption, which reflects the total electrical capacity available to power the servers hosted by the data centre.<sup>18</sup> The largest data centres, frequently used for AI and other compute-intensive applications, are called “hyperscale” data centres: massive facilities containing over 5,000 servers inside a large physical footprint and drawing over 100MW of electricity.<sup>19</sup> Meta plans to open a 1GW AI data centre in Ohio, which should come online by the end of 2026, and the company has also announced a future 5GW AI data centre will be constructed in Louisiana.<sup>20</sup> For comparison, 1GW of electricity would power roughly 750,000 homes.<sup>21</sup>

AI data centres require significant amounts of electrical power and water, and as demand for data centres expands, this will need to be managed.<sup>22</sup> For example, global data centre electricity demand is forecast to more than double in coming years, from 82 GW in 2025 to 219 GW in 2030, with the majority of this increase coming from the growth of AI.<sup>23</sup> Hyperscale data centres specifically will be a significant component of this growth.<sup>24</sup> In many geographies, data centre construction is facing political opposition from local communities, often because of these resource demands.<sup>25</sup>

Economically, the overall impact of AI remains uncertain, and economists and investment strategists are still debating how AI will affect global productivity, growth, and the labour market.<sup>26</sup> Despite the progress with AI’s capabilities, the technology still has some distance to cover before it is widespread in the workplace; as one recent report noted, AI “success rates on more realistic workplace tasks remain low, highlighting a gap between benchmark performance and real-world effectiveness.”<sup>27</sup>

Thus, as recently highlighted by the Washington Post, as AI starts “to reshape the economy, [the impact] may be challenging to detect and clearly measure. That may leave political and corporate leaders to choose the numbers that fit their preferred narratives on how AI is

changing... life and work.”<sup>28</sup> But even the more conservative estimates predict AI technology will have a “nontrivial” impact on the economy of the future.<sup>29</sup>

However, in addition to its potential for unlocking significant economic opportunities in the years ahead, AI will also bring a set of novel risks. Therefore, while seeking to capture value from AI, Canadian financial services firms will need to develop new ways to manage AI-related risks.

# AI-System Related Risks for Financial Institutions

## 1. Operational risks

Across Canada’s financial services sector, AI has rapidly transitioned from experimental use to deep operational integration. While these tools promise significant benefits, a growing dependency on AI, and the physical infrastructure that supports it, is also a source of systemic stress in the financial system. As Peter Routledge, the Superintendent of Financial Institutions, has noted, “AI is a transformative force—both awe-inspiring and potentially perilous.”<sup>30</sup>

To mitigate risks from AI, starting in May 2027, federally regulated Canadian financial institutions will need to follow the Office of the Superintendent of Financial Institutions (OSFI) Guideline E-23 on Model Risk Management.

Guideline E-23 states that financial institutions should be “cognizant of how the use of [AI] models in their business can impact their risk profile and should have effective risk management practices to mitigate the risks. Model risk management should be conducted with integrity, at all times, particularly in a world where newer use cases, including those powered by AI, play a greater role in day-to-day operations.”<sup>31</sup> AI-related operational risks are also covered in the existing OSFI E-21 Guideline on Operational Risk and Resilience, which includes technology, cyber, and data risk management.<sup>32</sup>

To manage and mitigate AI-related risks to their operations, Canadian financial institutions will need to develop policies and procedures that fit within and complement their broader risk governance frameworks.<sup>33</sup> Below, some of the main types of operational risks arising from AI are explored in greater detail.

### AI use and dependency across financial institution operations

Globally, the adoption of AI tools and AI agents by large corporations continues to accelerate. This trend is also occurring in the Canadian financial services sector, where AI is now embedded across the major operational domains of Canadian financial institutions. This increasing use of AI is driven by more reliable and secure AI models, productivity gains seen by early adopters, and increasing client expectations for digital and personalized services.<sup>34</sup> Indeed, according to a survey by KPMG, over 90% of Canadian financial services leaders now say they view generative AI as “critical” to competitive advantage.<sup>35</sup>

Within Canadian financial services, AI is being used across multiple domains. In capital markets, AI is being used to analyze data and conduct investment research, while algorithmic and high-frequency trading systems are increasingly relying on machine learning models for signal generation, execution optimization, and risk management.<sup>36</sup> In retail and commercial banking,

AI drives credit scoring, loan origination, and collections optimization.<sup>37</sup> Insurance companies are deploying AI for actuarial modelling, claims processing, and fraud detection.<sup>38</sup> And, across all sectors of the Canadian financial services industry, AI powers automated client engagement, customer service chatbots, anti-money-laundering surveillance, and regulatory reporting systems.

While AI delivers substantial benefits, widespread adoption of, and dependency on, the technology by financial institutions also creates heightened operational risks. As OSFI noted in a recent report, “Financial institutions are now using AI for more critical use cases, such as pricing, underwriting, claims management, trading, investment decisions, and credit adjudication. The use of AI may amplify risks around data governance, modelling, operations, and cybersecurity. Third-party risks increase as external vendors are relied upon to provide AI solutions. There are also new legal and reputational risks from the consumer impacts of using this technology that may affect financial institutions without appropriate safeguards and accountability.”<sup>39</sup>

Furthermore, as financial institutions become more dependent on AI, this dependency itself will become a source of systemic stress. For example, a significant disruption to AI systems, whether caused by software failure, a cyberattack, a cloud infrastructure outage, or a regulatory intervention, could simultaneously affect multiple business lines within a single financial institution and multiple institutions across the sector. This common exposure is a hallmark of systemic risk.<sup>40</sup>

The fact that both the capabilities of AI technology, and the use of this technology by financial institutions, are growing faster than regulation and risk management practices is a concern for OSFI.<sup>41</sup> This prompted OSFI to collaborate with the Global Risk Institute (GRI) to convene a group of stakeholders from across the AI sector to create the Financial Industry Forum on Artificial Intelligence (FIFAI).<sup>42</sup> In 2023, FIFAI discussions led to the development of the “EDGE” (Explainability, Data, Governance, and Ethics) principles to guide risk management approaches for AI technologies.<sup>43</sup>

Building on this foundation, GRI and OSFI extended their FIFAI work to new partners spanning multiple workshops held throughout 2025 to deepen understanding of how AI technologies are reshaping opportunities and risks. In the spring of 2026, this work introduced the “AGILE” framework (Awareness, Guardrails, Innovation, Learning, and Ecosystem Resilience).<sup>44</sup> This work highlights how the rapid scaling of AI is creating new operational dependencies, concentration risks, and systemic vulnerabilities—reinforcing the need for coordinated, ecosystem-level resilience.

## Operational risks from AI agents

As mentioned above, AI agents are rapidly gaining traction across a wide range of business use cases, and the AI agent market is forecast to grow substantially in the coming years.<sup>45</sup> Financial institutions are beginning to integrate AI agents into trading operations, compliance monitoring, and customer-facing services, and are reporting a positive return on these investments.<sup>46</sup>

The capabilities of AI agents are powerful, but these tools also introduce new and, in some cases, poorly understood risks. AI agents that make decisions without human oversight can negatively impact risk management by propagating errors, interacting with other automated systems in unanticipated ways, pursuing hidden objectives, and generating outcomes that are difficult to audit or reverse.<sup>47</sup> According to one study, 80% of companies report that their AI agents have taken unintended actions.<sup>48</sup> This adds complexity of risk management for firms using these agents in their operations.

AI agents also present cybersecurity risks. As a recent report from McKinsey & Company noted, “in cybersecurity terms, you might think of AI agents as ‘digital insiders’—entities that operate within systems with varying levels of privilege and authority. Just like their human counterparts, these digital insiders can cause harm unintentionally, through poor alignment, or deliberately if they become compromised.”<sup>49</sup> AI agents can access unauthorized data or be “tricked” into revealing access credentials.<sup>50</sup> As a recent high-profile example, in March 2026, an AI agent used at Meta caused a large leak of sensitive user data.<sup>51</sup>

AI agents can also be used by hackers to mount more sophisticated cyberattacks. For example, in late February 2026, cybersecurity researchers were able to use an AI agent to hack into McKinsey & Company’s own in-house AI platform, gaining access to 46.5 million chat messages and almost 750,000 files.<sup>52</sup> The CEO of Codewall, the firm that conducted the test, warned that “we used a specific AI research agent to autonomously select the target... Hackers will be using the same technology and strategies to attack indiscriminately.”<sup>53</sup>

As firms become increasingly dependent on AI agents, the potential for cascading operational failures will also grow. This makes robust risk management critical.

## Cybersecurity risks

As AI’s capabilities grow, the technology will pose greater cybersecurity risks. The case of Anthropic’s new Claude “Mythos” AI model provides a compelling illustration of these risks.

In April 2026, Anthropic announced that its new general-purpose model, named Mythos, had discovered thousands of cybersecurity vulnerabilities across every major operating system.<sup>54</sup> The company described Mythos as a “cybersecurity reckoning” and stated that it would delay

the public release of the model, deeming it too powerful for unrestricted access.<sup>55</sup> Highlighting the gravity of this risk for financial services firms, U.S. Treasury Secretary Scott Bessent and Federal Reserve Chair Jerome Powell held an urgent meeting with bank CEOs in early April 2026. The goal was to ensure that banks were aware of the potential dangers posed by Mythos and similar AI models—and to confirm that they were taking appropriate precautions to defend their systems.<sup>56</sup>

These cybersecurity risks will increase as AI continues to improve. As one prominent cybersecurity business leader recently wrote, “These are not incremental improvements. Imagine a horde of agents methodically cataloging every weakness in your technology infrastructure, constantly. Over the next six months, the barrier to entry for sophisticated attacks will continue to diminish. A hacker’s dream weapon will be available to anyone with a credit card and compute. What makes this moment different is not just capability. It is the asymmetry, and for now, it favors the attacker. A single bad actor will now be able to run campaigns that once required entire teams.”<sup>57</sup>

## Cloud infrastructure concentration

Beyond the AI tools themselves, the infrastructure that powers AI also presents operational risks for Canadian financial institutions. The reliance on cloud infrastructure is one example.

Cloud infrastructure comprises the “building blocks” that support cloud computing services, including AI tools.<sup>58</sup> AI systems deployed in the cloud rely on this infrastructure to store data, train models, and deliver services over the internet.<sup>59</sup> The majority of AI systems are hosted in “hyperscale” data centres, which are “large-scale data centers that provide a wide range of cloud computing and data solutions for businesses that need vast digital infrastructure, processing, and storage.”<sup>60</sup> As AI grows, hyperscale data centres are estimated to capture about 70% of the forecast capacity in the U.S. data centre market.<sup>61</sup>

Currently, the AI systems used by Canadian financial institutions are overwhelmingly hosted on cloud infrastructure provided by a small number of American-owned hyperscalers, including Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), who collectively control approximately two-thirds of the global cloud infrastructure market.<sup>62</sup> These firms provide the majority of the world’s cloud services, including in Canada.<sup>63</sup>

This concentration of cloud infrastructure providers creates a critical vulnerability, as Canadian financial institutions have limited practical alternatives for the scale of compute required to run advanced AI workloads. This vulnerability aligns with the FIFAI II report’s findings that “growing dependence on a small number of AI providers and opaque AI supply-chain dependencies heighten systemic fragility,” pointing to structural concentration risk in the AI ecosystem.<sup>64</sup> And

a prolonged outage at a single hyperscale provider could simultaneously disrupt the operations of multiple major Canadian financial institutions.<sup>65</sup>

Additionally, given that these hyperscalers are American-owned, and are thus subject to U.S. law, Canadian firms that use their services are vulnerable to geopolitical tensions that may arise in the future and may face increasing concerns about data sovereignty.<sup>66</sup> As a recent report from the University of Toronto's Munk School noted, "U.S.-based technology giants have become central instruments of American geopolitical power.... As their technologies become global standards, they bring American law with them, allowing the United States to hardcode its technical choices into geopolitical leverage."<sup>67</sup> To mitigate risks from this dependency on foreign cloud infrastructure, the Canadian federal government is planning to build a "Canadian Sovereign Cloud."<sup>68</sup>

However, even a sovereign Canadian cloud would face issues with data centre concentration that could be a risk for Canadian financial institutions. As the FIFAI II report highlighted, "Data that is inconsistent or incomplete, and fragmented across platforms (including those of third parties and offshore storage) can lead to increased data sovereignty concerns, privacy risks for consumers, and challenges for regulatory oversight."<sup>69</sup> And while financial institutions do maintain business continuity plans, the depth of AI integration into core operations means that cloud infrastructure failures increasingly threaten functions that cannot be easily performed through manual fallback procedures.

## Grid stress and backup power limitations

The energy needed for the physical infrastructure that supports AI operations, including data centres, power grids, and cooling systems, constitutes another layer of operational risk.

AI data centres are extremely energy-intensive, requiring continuous and dense electrical power. As a report from McKinsey & Company noted, "all data centers consume significant amounts of energy, but AI-ready ones are especially demanding because of their high average power densities—the energy consumption of servers in the racks. Average power densities have more than doubled in just two years, to 17 kilowatts (kW) per rack, from eight kW, and are expected to rise to as high as 30 kW by 2027 as AI workloads increase."<sup>70</sup>

As AI growth drives the expansion of data centres, energy demand from these facilities is increasing significantly. For example, total data centre energy requirements in the U.S. are forecast to be between 325 and 580 terawatt-hours by 2028, representing as much as 12% of total U.S. electricity consumption.<sup>71</sup> Similarly, in Canada, by 2030, data centres could account for 14% of total electricity consumption.<sup>72</sup> And, according to a paper by the law firm Osler, "in some provinces, the amount of power requested by data centre developers alone exceeds the total amount of power otherwise required by the entire province."<sup>73</sup>

This is placing unprecedented burdens on electrical grids. In some jurisdictions, grid capacity is struggling to keep pace with data centre energy demand, which may affect grid reliability.<sup>74</sup> Across North America, strain on electrical grids could be particularly acute during severe weather events, and reliability issues (such as blackouts or brownouts) may arise.<sup>75</sup>

Blackouts and brownouts are a risk for data centre operators. Backup power systems such as diesel generators provide a limited buffer. These backups can sustain data centre operations for hours, but not for the days or weeks that a major grid disruption could last. And while data centre operators are expanding use of co-located energy resources, few data centre operators have the expertise required to co-locate batteries or natural gas and solar electricity production on the massive scale required for hyperscale data centres.<sup>76</sup> Currently, some hyperscalers are building their own natural gas plants to power data centres, but these plants are also vulnerable to weather shocks that impact natural gas production.<sup>77</sup>

Therefore, grid stress impacting data centres is a risk to operational continuity for Canadian financial institutions using AI. This risk will increase as data centre energy demand grows in the coming years, as discussed in the section “Data centre energy demand and grid stress.”

## 2. Investment risks

Globally, financial institutions hold substantial investment exposure to the AI sector and its supporting infrastructure.<sup>78</sup> This exposure includes:

- Direct equity holdings in AI and semiconductor companies,
- Indirect positions through debt markets and supply-chain investments,
- Indirect exposure to a small number of mega-cap technology stocks that make up a large concentration of the overall market,
- Physical asset exposures related to data centre real estate and energy infrastructure.

The investment risks embedded in these positions are amplified and compounded by several factors: the extraordinary pace of capital expenditure flowing into AI infrastructure, physical resource constraints that may limit the sector's growth trajectory, and the possibility that current AI valuations contain speculative excess that could unwind sharply.

This section provides a brief overview of these direct and indirect investment risks facing Canadian financial institutions because of their exposure to the AI sector.

### Direct and indirect AI investment exposure

The most visible dimension of AI investment exposure for Canadian financial institutions is direct equity positions in AI-focused companies, including chipmakers, cloud providers, and AI software firms. Additionally, Canadian financial institutions have significant positions in AI infrastructure, such as semiconductor manufacturers and data centres.<sup>79</sup> These positions have generated substantial returns in recent years, driven by the rapid growth of AI adoption and the increase in capital expenditure this has catalyzed.

Canadian financial institutions also have significant indirect exposure to AI. Canadian financial institutions hold positions in corporate bonds, leveraged loans, and structured credit instruments linked to data centre developers, energy companies supplying AI-related power demand, and the broader AI technology supply chain. For example, in July 2025, one major Canadian pension fund provided a \$225 million loan for the construction of a hyperscale data centre in Ontario.<sup>80</sup>

In the coming years, it is estimated that between three and five trillion U.S. dollars could be spent to build AI infrastructure, and a significant amount of these funds will come from global debt markets.<sup>81</sup> According to a recent report from the Bank for International Settlements, "as the need for AI-related investment grows, firms are increasingly turning to external sources of financing. Debt financing, through corporate bonds, leasing arrangements or loans, allows

investors to spread costs over time and align financing maturities with the long economic life of data centre assets... A particularly fast-growing source of external financing is private credit.”<sup>82</sup>

These debt instruments carry credit risk that is correlated with the overall trajectory of AI investment and, ultimately, with the realization of the revenue growth embedded in current AI valuations. This poses a risk should these valuations unwind, which is a concern for central banks and financial regulators.<sup>83</sup> As Bank of Canada Governor Tiff Macklem noted in a recent speech at the Global Risk Institute (GRI), “Private credit is expected to be an important source of the debt funding needed to grow AI infrastructure. The issue is not private credit itself. It’s how private credit will behave under stress—and the risks it poses to the broader financial system.”<sup>84</sup>

Finally, AI comprises an “extreme” share of the total stock market.<sup>85</sup> This market concentration, which is best exemplified by the so-called “Magnificent Seven” large mega-cap technology companies, means that even investors passively exposed to broad indexes have significant exposure to AI. This risk is explored further below.

## Market concentration: The “Magnificent Seven”

The concentration of AI-related market capitalization in a small number of large American technology companies presents a distinctive form of investment risk, with both direct and indirect investment implications.

Collectively, Alphabet, Amazon, Apple, Tesla, Meta Platforms, Microsoft, and Nvidia are commonly referred to as the “Magnificent Seven.”<sup>86</sup> These seven high-performing, dominant, and influential U.S. mega-cap technology companies have largely driven the stock market’s growth in recent years, particularly during the current AI-led bull run.<sup>87</sup> As of this writing, these seven companies collectively account for over 30% of the value of the S&P 500’s total market capitalization, a level of concentration not seen since the late 1990s.<sup>88</sup> Canadian institutional portfolios, whether through direct holdings or through passive index-tracking strategies, are correspondingly concentrated in these firms.

This concentration creates a vulnerability to correlated losses. For example, in February 2026, Nvidia’s shares fell after its earnings release, dragging down the broader index.<sup>89</sup> In the future, a downturn in AI-related sentiment (triggered by disappointing earnings, a technological setback, a regulatory intervention, a shift in the macroeconomic environment, or some other factor) could simultaneously depress the valuations of all major AI-linked equities. Because these stocks dominate major indices, such a correction would propagate well beyond the technology sector, affecting balanced portfolios, pension obligations, and the wealth effect supporting consumer spending and credit quality.

## AI capital expenditure and crowding-out effects

The scale of capital investment flowing into AI infrastructure is historically extraordinary. As RBC noted in a recent report, “The launch of ChatGPT in late 2022 unleashed one of the fastest and largest capital expenditure (capex) cycles in decades... Capex among [big tech] firms has more than doubled in the last two years, reaching \$427 billion in 2025. Momentum shows few signs of fading heading into 2026, with projections pointing to a further 30 percent year-over-year increase to roughly \$562 billion.”<sup>90</sup>

Multiple forecasts suggest a continued acceleration in this AI-related capital expenditure. For example, four of the biggest U.S. technology companies (Alphabet, Amazon, Meta, and Microsoft) alone are planning to spend \$650 billion on AI infrastructure in 2026.<sup>91</sup> In one striking example, Alphabet has said its capital expenditure could double in 2026.<sup>92</sup> This investment is financing new data centres, advanced semiconductor fabrication, power generation and grid infrastructure, and the extensive cooling systems that AI hardware requires.

While this capital expenditure drives economic activity, it may constrain investment or negatively affect the economic outlook in other sectors. For example, a recent analysis by Bloomberg found that in the U.S., spending on factory construction was down 2.5% while spending on data centre construction was up almost 18%. As Bloomberg wrote, “access to capital, power and people is key to the success of any construction or manufacturing project—and right now AI is gobbling up all three.”<sup>93</sup> This dynamic has led some to warn of “the crowding-out effect of the AI boom and the tradeoffs it creates in other areas of the economy.”<sup>94</sup>

Another example of this “crowding out” effect is the impact the data centre build-out is having on the broader construction industry. Analysts note that rising demand from data centre construction will create price pressure for copper, aluminum, and other materials.<sup>95</sup> This, in turn, is leading to increased costs, tighter margins, and longer lead times for the construction industry.<sup>96</sup> It may also have a negative impact on housing construction.<sup>97</sup>

Capital markets have finite absorptive capacity. While the enormous sums flowing into AI infrastructure reflect investors’ pursuit of higher expected returns, they also compete with financing needs in housing, transportation, healthcare, clean energy, and other sectors. To the extent that this scale of AI capital expenditure pushes up the overall cost of capital, intensifies competition for critical inputs, or redirects investment away from other productive uses, it may generate macroeconomic headwinds that indirectly affect the broader portfolios of Canadian financial institutions.

## The AI bubble question

A major question facing global investors and asset owners is whether there is a market “bubble” in the AI sector. A market bubble occurs when the price of assets rapidly climbs “to a point where they far exceed their intrinsic value or their earnings. This price bubble, based on speculation, can include all equities in a stock market or those from a specific sector. When the bubble eventually bursts and prices start dropping, it can lead to panic selling and potentially a stock market crash.”<sup>98</sup> The possibility that current AI valuations embody the speculative excess of a bubble warrants careful consideration.

Historically, there are parallels worth considering. The dot-com bubble of the late 1990s, the nineteenth-century railroad bubbles in the U.S. and U.K., and the global telecom bubble of the early 2000s all shared a common pattern: a genuinely transformative technology attracted massive investment that outpaced the technology’s ability to generate returns within the expected timeframe.<sup>99</sup> In each case, the technology ultimately fulfilled much of its transformative promise, but the intervening correction destroyed value and inflicted damage on investors and financial institutions.<sup>100</sup>

Currently, there is a major debate as to whether the current surge in AI investments constitutes a bubble. Experts note the tension between high valuations and the actual and forecast profitability of technology and AI firms. As one finance academic wrote, “earnings growth has largely matched price increases for AI infrastructure leaders, but the market is pricing in continued exceptional growth for years to come. Whether those expectations prove realistic or represent overextrapolation will determine if current valuations are justified or become the foundation of a correction.”<sup>101</sup>

Furthermore, the sustainability of AI capital expenditure at current levels depends on the materialization of AI revenue growth that remains speculative and thus may not materialize. As analysts at JP Morgan and other financial institutions have noted, the current “hyperbolic” levels of capital expenditure on data centres and other AI infrastructure are being driven by the expectation that demand for AI compute will continue to grow.<sup>102</sup> Yet revenue models for many AI applications remain unproven.<sup>103</sup> And concerns have been raised about how AI firms account for semiconductor depreciation in their financial models.<sup>104</sup>

Beyond the high levels of capital expenditure and uncertain revenue projections, concerns are growing about the “increasing circularity of the AI ecosystem,” where companies invest in the same partners who purchase their hardware, potentially creating an artificial feedback loop of fictive demand.<sup>105</sup> As Bloomberg wrote, this “web of interlinked investments raises the risk of cascading losses if AI falls short of its potential.”<sup>106</sup>

A disorderly correction in AI valuations, caused by the “popping” of a potential AI bubble, could have a major impact on financial markets. Direct equity losses would affect investors with positions in the AI and broader technology sector, and debt instruments linked to AI infrastructure would face credit deterioration. Given the heavy concentration of large technology companies in major market indexes, the overall market would also face a significant contraction.<sup>107</sup>

These effects could reverberate through market sentiment, triggering a broader risk-off environment. This, in turn, would widen credit spreads and impair liquidity across markets well beyond the technology sector—a chain reaction that has been called “a Multidimensional Economic Disaster.”<sup>108</sup> The contagion pathways would closely resemble those observed in prior asset bubble corrections, amplified by the financial interconnectedness of contemporary markets. This could destroy over US\$30 trillion in value and cause a significant global recession.<sup>109</sup>

To be sure, there is a credible counterargument: AI may represent a genuinely transformative technology whose long-term economic value justifies elevated valuations. Unlike purely speculative manias, the AI sector is anchored by measurable productivity gains, rapid enterprise adoption, and substantial revenue growth among leading firms.<sup>110</sup> If AI delivers on even a fraction of its projected potential, current valuations may prove not only defensible but conservative.

However, as noted above, even technologies that ultimately reshape economies can experience sharp interim corrections driven by sentiment shifts, earnings disappointments, or the recognition that adoption timelines are longer than markets have priced in. In other words, financial stability risks do not necessarily depend on whether AI proves truly transformative; they depend on how markets behave along the way.

This underscores the importance of proactive risk management. Financial institutions should begin preparing for the possibility of an AI-driven market correction. Key steps could include rigorous stress testing of AI-linked credit exposure, especially loans and bonds financing data centre buildouts, to identify vulnerabilities before they materialize.<sup>111</sup> Broadening portfolio diversification, developing hedging strategies, and creating contingency plans would also help mitigate the impact of a potential sharp AI-driven downturn. Importantly, none of these measures require a firm view on whether AI is currently in bubble territory. They simply reflect prudent preparation for turbulence on a path that, whatever its ultimate destination, is unlikely to be smooth.

## Data centre energy demand and grid stress

The physical resource requirements of AI infrastructure constitute a material investment risk that is frequently underappreciated. As discussed in the operational risks section above, the data centres upon which AI depends require vast quantities of energy. This is because current AI models consume significant energy for both training and inference.

Training is the iterative process of teaching an AI model to recognize patterns, make decisions, or perform specific tasks by feeding it large datasets. This requires significant computing power and is therefore extremely energy intensive. As an example, the International Energy Agency (IEA) estimated that OpenAI's GPT-4 model had "a training energy demand... equivalent to the daily electricity consumption of around 28,500 households in advanced economies."<sup>112</sup>

Once an AI model is trained, it uses energy for inference. Inference is the "use phase" of an AI model; every time a user employs an AI tool, such as asking Claude or ChatGPT a question, producing the answer requires AI inference.<sup>113</sup> Because AI is widely used, it is inference, not training, that comprises the majority of AI energy use; estimates are that as much as 90% of computing power for AI is currently used for inference.<sup>114</sup>

Researchers are working on more energy-efficient AI systems, and studies show that small changes to AI models can significantly reduce their energy consumption.<sup>115</sup> However, as AI continues to be widely adopted across society, the energy demand from inference will increase. For example, as one researcher estimated, "integrating ChatGPT into the 9 billion daily Google searches could raise global electricity demand by approximately 10 terawatt-hours annually, equivalent to about 1/60 of Canada's total annual electricity consumption."<sup>116</sup> This could potentially be equivalent to powering every home in the Greater Toronto Area for more than two years, and this represents only the incremental cost of adding one AI application to one search engine.

It is therefore extremely likely that AI will continue to consume significant amounts of energy in the future. As an example, a recent report by the IEA estimated that global data centre electricity consumption could more than double by 2030, to 945 TWh, with AI workloads driving much of this increase.<sup>117</sup> This forecast increase in energy demand for AI creates several investment-related risks for Canadian financial institutions.

First, concentration of data centres in regions with limited grid capacity creates location-specific risks of grid stress, power price volatility, and regulatory intervention. As noted in the operational risks section above, if all the data centre projects currently planned in Canada proceed, data centres would account for 14% of Canada's total power needs by 2030. This raises serious concerns about power availability and grid reliability in many jurisdictions.<sup>118</sup>

Growing AI energy demand is a global issue. In many parts of the world, demand from AI data centres already exceeds (or is rapidly approaching) the available supply of electricity.<sup>119</sup> This is making the facilities unpopular among some sections of the public and is leading to political ramifications. For example, currently in the U.S., a group of progressive federal lawmakers is advocating a moratorium on data centre construction.<sup>120</sup>

This political pushback against data centres has implications for investors. As a recent legal paper notes, investors in data centres must “consider potential pushback from residents who may challenge the development of data centres in their locality due to the potential for major rate hikes to cover the surge in energy demand. These concerns have already become a reality in the U.S., where the federal energy regulator refused an application from one of the world’s largest cloud services providers to purchase power from a Pennsylvania nuclear power plant on the grounds that it would raise customer rates and threaten the reliability of the local grid.”<sup>121</sup>

Second, the cost of energy is a significant component of data centre operating expenses. Electricity alone can comprise 20% to 30% of the total cost of operating a data centre.<sup>122</sup> And, in many jurisdictions, the growth in demand from data centres may increase electricity prices.<sup>123</sup> Therefore, sustained increases in electricity prices driven by AI-related demand growth could erode the margins of data centre operators and, by extension, the valuations of data centre investments and the creditworthiness of debt instruments linked to those assets.

Third, the strain that data centre demand places on electrical grids may create inflationary pressures that extend beyond the technology sector. This could negatively affect other energy-intensive industries, as well as electricity costs for both households and small businesses.<sup>124</sup> If data centres lead to a significant increase in energy costs, this could compound macroeconomic risks for the broader loan books of Canadian financial institutions.

## Water consumption and cooling infrastructure

In addition to electricity, data centres also consume large amounts of water. For both training and inference, AI models perform energy-intensive calculations that generate a significant amount of heat. To keep the AI infrastructure cool, data centres rely on water for chillers, cooling towers, and liquid cooling systems. This requires a significant amount of water. For example, according to some estimates, ChatGPT uses 519 millilitres of water (roughly the volume of a bottle of water) to draft a 100-word email.<sup>125</sup> Currently, a large data centre can consume as much as 5 million gallons of water per day.<sup>126</sup>

One factor driving this high consumption of water for cooling is that data centres currently have a limited ability to reuse water. As one report noted, “during cooling, a portion of freshwater evaporates while the remaining water turns into wastewater contaminated with dust,

chemicals, and minerals, which reduces cooling efficiency if recirculated and often prevents data centres from maximizing wastewater recycling.”<sup>127</sup> Public concerns about this contamination of water supplies are one of the causes of the growing political backlash against data centre construction.<sup>128</sup>

A further factor driving data centres’ growing demand for water is that some of the newer semiconductors, designed specifically for AI, require more cooling.<sup>129</sup> Around the world, the so-called “thermal density” of AI data centres is increasing as technology becomes more advanced.<sup>130</sup>

As AI use increases and more data centres are built, forecasts show the amount of water used by data centres for cooling could increase by over 800% from 2025 levels in the coming years.<sup>131</sup> In the U.S. alone, one study estimates that the deployment of AI servers could have an annual water footprint of 731 to 1,125 million cubic metres between 2024 and 2030.<sup>132</sup> Another study estimates that global AI demand could account for over 6 billion cubic metres of water withdrawal by 2027.<sup>133</sup>

Major hyperscalers such as Amazon, Microsoft, Google, and Meta have pledged to be “water positive” by 2030.<sup>134</sup> And data centre operators are actively exploring new technologies to reduce water consumption. This includes using “circular” systems that reuse water and the development of innovative cooling methods that require less water.<sup>135</sup>

However, there are issues with many of these new technologies. As a recent report noted, “some of these strategies involve tradeoffs that lower water use but increase energy demand, which can result in higher greenhouse gas emissions, greater demand on the electric grid, and negative impacts on local air quality. For instance, because they require more energy to run, air-cooling systems are generally less efficient at cooling data centers than water-cooling systems, resulting in higher energy use... and increased electricity costs.”<sup>136</sup>

Data centre water use can place significant stress on local water systems in the communities that host these facilities. For example, in 2021 in The Dalles, Oregon, Google data centres consumed 355 million gallons of water, which was 29% of the city’s total water consumption.<sup>137</sup> In The Dalles and other communities, citizens have raised concerns about local water stress from data centres.<sup>138</sup> Similar water use concerns have also been raised in Canada.<sup>139</sup> In response, an increasing number of jurisdictions are exploring regulations to limit data centre water consumption.<sup>140</sup>

Further complicating matters, many data centres are being built in areas facing growing water stress. Analysis by MSCI found “the rapid build-out of data centres to power AI and cloud computing coincides with intensifying global water stress. Around one-quarter of today’s facilities, and nearly one-third of those under construction, are in regions projected to face

greater water scarcity by 2050. As climate change increases water stress, the industry’s ability to cool efficiently and sustainably may become a critical factor in operational resilience and, in turn, long-term investment performance.”<sup>141</sup> As an example, in the U.S., the most common locations for data centres are California, Virginia, and the Southwest, all regions where water scarcity is already an issue.<sup>142</sup>

Water scarcity events, whether driven by drought, competing agricultural and municipal demands, or regulatory restrictions, pose direct risks to data centre operational continuity and, consequently, to investment positions tied to these assets. For Canadian financial institutions with exposure to data centre assets, this water risk may represent a physical constraint on the growth trajectory that current valuations assume (indeed, water risk has been called “AI’s hidden cost”).<sup>143</sup> Water risk could also be a potential trigger for operational disruptions and regulatory responses that could impair asset values.<sup>144</sup> Thus, a deep understanding of AI-related water risks is crucial for prudent investment risk management by Canadian financial institutions.

### 3. Downstream risks from AI

The previous sections of this paper explored the operational and investment risks related to AI. While significant, those risks represent the more familiar dimensions of AI-related risk and are therefore the dimensions most likely to be captured by existing risk management frameworks within Canadian financial institutions.

This section moves beyond these first-order exposures into the second- and third-order risks from AI. These downstream risks are harder to anticipate, more difficult to quantify, and therefore are potentially greater sources of unexpected shocks.

Downstream risks are, by their nature, systemic. They originate in the interaction between AI's rapid proliferation and the broader social, political, economic, and geopolitical systems within which both AI and Canadian financial institutions operate. These risks are not confined to a single business line or asset class; they propagate across systems and sectors in ways that defy conventional risk categories.

Furthermore, these risks can interact and compound. For example, a geopolitical disruption to semiconductor supply chains would not merely affect chip prices; it would simultaneously threaten AI operational continuity, depress AI-linked equity valuations, provoke regulatory responses, and could intensify the political dynamics surrounding data centre development. For this reason, understanding these interactions and preparing for their cascading consequences will be a central challenge for AI-related risk management in the years ahead.

#### Regulatory and political risks from data centres

As discussed above, data centres are increasingly becoming politically contentious. In many jurisdictions, the rapid expansion of data centre development, driven by surging demand for AI services, is generating growing public opposition. Communities are increasingly concerned about air pollution, electricity costs, water use, visual intrusion, and infrastructural strain from data centres.<sup>145</sup> Concerns are also emerging about data centre noise pollution and air quality impacts.<sup>146</sup>

This is impacting data centre development projects. According to one study in the U.S., “\$18 billion worth of data center projects were blocked, and another \$46 billion of projects were delayed over the last two years” due to “a growing wave of local, bipartisan opposition.”<sup>147</sup> Another recent study found “at least \$156 billion across 48 [data centre] projects with publicly disclosed values was blocked or stalled amid coordinated local opposition in 2025.”<sup>148</sup>

Public opposition to data centres is increasingly translating into political action. Several jurisdictions in the U.S. and Europe have imposed moratoriums on new data centre

construction, and zoning delays, permitting challenges, and enhanced environmental review requirements are becoming more common. These permit delays contributed to the construction of new data centres in the U.S. falling in 2025, for the first time since 2020.<sup>149</sup>

Political risks for data centres are likely to intensify, driven by rising energy and water consumption that is sharpening public opposition.<sup>150</sup> And, as a recent Brookings report noted, “Local opposition is likely to soon become the leading constraint on data centre siting and approvals.”<sup>151</sup>

This challenge is compounded by the fact that, beyond the construction phase, data centres generate relatively few permanent jobs in host communities.<sup>152</sup> Many communities, especially in rural areas, therefore feel like they get few significant benefits in return for the costs data centres impose.<sup>153</sup> As data centres become more politically contentious, politicians responsive to constituent concerns will likely enact increasingly stringent regulations. As the *New York Times* reported in March 2026, in many communities public “opposition ‘hardened’ into legislative hurdles... narrowing the options for data center developers looking for sites.”<sup>154</sup>

In an effort to address public concerns about energy use, and thus mitigate these political risks, the large technology firms that operate hyperscale data centres have made pledges to fund the electricity costs or to develop alternative energy sources.<sup>155</sup> For example, in the U.S., Amazon, Google, Meta, Microsoft, OpenAI, Oracle, and X (formerly Twitter) have all signed the Trump Administration’s “Ratepayer Protection Pledge,” a voluntary commitment to “protect American consumers from price hikes due to data centre energy and infrastructure requirements.”<sup>156</sup> However, as CNBC reported, “experts have questioned the legitimacy of such commitments, given that hyperscalers have struggled to turn profits.”<sup>157</sup>

The political dynamics around data centres carry direct implications for Canadian financial institutions. Data centre moratoriums and regulatory restrictions could affect the revenue projections embedded in data centre real estate investment valuations and in the debt instruments that finance data centre construction. Permitting delays will increase development costs and timelines, eroding expected returns. And if data centres become even more controversial with the public, this could create reputational risks for the financial institutions that are publicly associated with financing or owning such projects.

## Shifting AI governance and regulatory frameworks

Changing AI governance and regulation also poses risks for Canadian financial institutions. AI governance refers to the internal rules, policies, and ethical guidelines organizations use to develop and deploy AI responsibly.<sup>158</sup> Many businesses and other organizations have developed their own internal AI governance frameworks.<sup>159</sup> AI regulations, by contrast, are externally

mandated rules — typically in the form of laws and enforceable standards — that govern how AI may be used. Canada, the EU, and the U.S. all have different regulatory approaches to AI, which regulate the use of AI in those jurisdictions.<sup>160</sup>

Around the world, governance and regulation of AI are evolving unevenly. There is no globally agreed set of AI standards, and the global regulatory landscape has therefore been described as “fragmented.”<sup>161</sup> Further complicating efforts to regulate AI is the rapid speed at which the technology is developing. As a report from Brookings noted, when it comes to AI regulation, the “challenge becomes how to protect the public interest in a race that promises to be the fastest ever run yet.”<sup>162</sup> Firms developing, investing in, and using AI systems thus face a complex and shifting global regulatory landscape.

When it comes to regulating AI, the EU’s AI Act is the most comprehensive framework to date.<sup>163</sup> The approach taken by the EU differs from the approaches emerging in the U.S., the U.K., China, and elsewhere.<sup>164</sup> The EU’s regulations are controversial for a number of reasons, but may be valuable as an indicator of how global AI regulation evolves.<sup>165</sup>

Currently, Canada does not have AI regulation at the federal level, and no province has its own AI laws that apply to the private sector.<sup>166</sup> Canadian federal AI regulations had been included in the proposed “Artificial Intelligence and Data Act,” which was part of Bill C-27, but as one researcher highlighted in a recent paper, this bill “died when Parliament was dissolved... for the April [2025] election. It has not been reintroduced, leaving Canada lagging other countries” with regard to AI regulation.<sup>167</sup>

The regulatory vacuum in Canada imposes risks on financial institutions that use AI. As the law firm Dentons wrote in a recent note on this topic, “Canada’s AI regulatory landscape for financial institutions is still taking shape. Without an overarching federal statute, the financial services industry must navigate a patchwork of guidance and regulation from privacy and securities regulators. As organizations move from pilot projects to deploying AI in real-world operations, they also face heightened exposure to securities litigation and privacy class actions.”<sup>168</sup>

Going forward, as AI regulations evolve, several important issues are likely to emerge for Canadian financial institutions. For example, some experts argue that controlling access to compute power will be one of the most effective ways to regulate AI.<sup>169</sup> As a result, future governance regimes may impose restrictions on the concentration and use of computing resources for training and deploying AI systems. These measures could directly constrain the AI infrastructure and capabilities available to financial institutions.

Additionally, changing privacy and data regulations, including evolving interpretations of consent requirements and cross-border data transfer restrictions, may affect both the datasets

on which financial AI systems are trained and the jurisdictions in which they can operate. Experts have already identified access to “data supply chains” as an area where Canada has regulatory gaps.<sup>170</sup> Model governance requirements, such as OSFI’s Guideline E-23 on Model Risk Management (discussed in the operational risks section above), could also impose compliance costs that will increase as AI systems become more sophisticated and autonomous. While strong privacy and data protection are essential safeguards for individuals and public trust, the uneven pace of regulatory change and differing approaches across jurisdictions can create compliance uncertainty and operational constraints for institutions.

To manage these risks, financial institutions will need to invest in governance infrastructure, compliance personnel, and audit capabilities specifically designed for AI systems. Moreover, the risk of regulatory surprises, such as abrupt changes in governance requirements driven by high-profile failures with AI systems or AI infrastructure, or by political shifts hostile to AI, introduces uncertainty that is difficult to hedge. These risks are especially acute for financial institutions that operate across multiple jurisdictions.

## Geopolitical risks

The AI sector is exposed to significant geopolitical risk, and these risks will increase as AI technology becomes more widespread. According to a recent report from the Munk School at the University of Toronto, AI “has become an increasingly dominant factor in geopolitical rivalry and great power competition... AI has the potential to impact national security across every domain: cybersecurity capabilities, cutting-edge research and development, economic productivity, and the ability to design the most sophisticated defence systems. As AI becomes increasingly powerful, control over the technology stack that enables it becomes inseparable from national power itself.”<sup>171</sup>

Given the number of domains that AI may transform, the intersection of geopolitics and AI is a vast and extremely complex topic. As a former U.S. National Security Advisor noted, “geopolitics in the age of AI will not be simple.”<sup>172</sup> What is clear is that this intersection will have major implications for both governments and businesses; as a report from Goldman Sachs noted, AI will “influence the course of markets and alter the balance of power among nations... In this swiftly evolving arena, corporate and political leaders alike are ...navigating new risks.”<sup>173</sup>

Given the complexity of this topic, scenario planning may be helpful. For example, in a recent article, experts laid out multiple possible future states for how AI and geopolitics could evolve in the years ahead and argued that for each of these, governments “should have a ready-to-execute plan that can be adapted as conditions shift. That requires institutions to think probabilistically.”<sup>174</sup> Firms should similarly apply scenario planning to prepare for possible

geopolitical risks from AI, given the high degree of uncertainty as to how the technology may evolve and influence geopolitics.

There are several case studies that illustrate how AI and geopolitics may intersect in ways that could cause systemic risks for Canadian financial institutions. These are: (1) shocks to the AI infrastructure supply chain, especially around semiconductors manufactured in Taiwan; (2) the geopolitical rivalry between the U.S. and China; and (3) attacks on data centres arising from military conflicts.

Currently, the global AI ecosystem depends on a semiconductor supply chain that is extraordinarily concentrated in Taiwan.<sup>175</sup> In 2025, Taiwan Semiconductor Manufacturing Company (TSMC) fabricated 70% of the world's semiconductors.<sup>176</sup> Within the broader semiconductor space, TSMC "dominates" the global market for the advanced semiconductors used for AI, manufacturing these specialized "chips" for clients such as Nvidia, Intel, Broadcom, and AMD.<sup>177</sup>

This concentration means the global semiconductor industry, and the AI firms that rely on those semiconductors, are exposed to geopolitical risk arising from potential Chinese aggression against Taiwan.<sup>178</sup> A military confrontation in the Taiwan Strait, a naval blockade, or even a significant escalation in cross-strait tensions could disrupt the global supply of the advanced semiconductors used for AI, crippling the technology industry and causing what U.S. Treasury Secretary Scott Bessent has called an "economic apocalypse."<sup>179</sup> Estimates are that this would cost the global economy anywhere from \$2 to \$10 trillion, depending on the scale of the military action.<sup>180</sup>

For Canadian financial institutions, a global semiconductor supply disruption stemming from a conflict over Taiwan would have immediate operational and investment implications. Operationally, it would constrain the expansion and maintenance of the AI infrastructure on which financial institutions increasingly depend. In investment portfolios, a Taiwan shock would trigger severe corrections across AI-linked equities, semiconductor holdings, and the broader technology sector, with contagion effects spreading across the wider market.

Beyond Taiwan, the global AI supply chain has other vulnerabilities to geopolitical shocks. For example, helium gas is essential for the fabrication of the advanced semiconductors required for AI.<sup>181</sup> As AI becomes more widespread, the need for this gas will increase; currently, experts forecast AI will drive a 5.5x increase in helium demand for semiconductor manufacturing by 2035.<sup>182</sup> Qatar currently produces over one-third of the world's helium supply, and military conflicts in the Middle East that disrupt this supply pose a major risk to global semiconductor manufacturing.<sup>183</sup>

The geopolitical rivalry between the U.S. and China is another area where geopolitics could affect the global AI space. Currently, the U.S. and China are in a tight race for AI supremacy.<sup>184</sup> In the past few years, this has generated export controls, investment restrictions, and technology transfer prohibitions.<sup>185</sup> These measures create compliance risks for Canadian financial institutions that hold cross-border positions in AI-related assets, and they introduce policy uncertainty that complicates investment strategy and risk assessment.<sup>186</sup> If the AI competition between the U.S. and China intensifies, it could accelerate the trend toward national security-driven industrial policies, forced technology localization, and the fragmentation of the global technology ecosystem, all of which could carry additional implications for the Canadian financial sector.

Another area where geopolitics could create AI-related risks for Canadian financial institutions is the potential for military conflicts to damage data centres and other AI infrastructure. For example, in early March 2026, Iranian drone strikes targeted data centres in the United Arab Emirates.<sup>187</sup> This led to disruptions in banking, payments, and enterprise software in the region.<sup>188</sup> As one report highlighted, “this was the first time that Big Tech data centers have been directly targeted by military strikes, and it brings a new threat to the doorstep of companies that have invested heavily in the region to keep pace with the AI boom.”<sup>189</sup> Outside the Middle East, reports indicate Russia is waging a sabotage campaign on critical infrastructure across Europe.<sup>190</sup> Given the importance of data centres to the modern economy, these facilities could be targeted by Russia or other actors engaging in “grey zone” aggression; this is also a risk in Canada.<sup>191</sup>

In the future, as AI becomes increasingly vital for economies, the physical infrastructure required for AI will likely become a priority military target in conflicts around the world. Canadian financial institutions that both depend on these AI services and invest in the infrastructure that provides them will thus be exposed to geopolitical risks affecting these assets.

## Economic and social risks

AI has the potential to displace labour and restructure industries at a pace that may outstrip the capacity of workers and society to adapt. This could have significant social and economic impacts that could present risks for Canadian financial institutions.

A great deal of attention is being paid to the impact that AI may have on the labour market. However, given how new the technology is and how quickly its capabilities are evolving, much remains unknown about how AI will impact employment. As a March 2026 paper by the Peterson Institute for International Economics noted, “research on AI and the labor market is still in the first inning.”<sup>192</sup>

Currently, the data show there is no significant link between AI exposure and changes in employment or unemployment rates across the broader economy.<sup>193</sup> However, within the broader economy, some sectors and groups are being disproportionately affected by AI. For example, the data shows that AI is having “a significant and disproportionate impact on entry-level workers in the American labour market.”<sup>194</sup> As researchers note, “AI’s early impact on these narrower sets of jobs and workers... might be harbingers of wider labour market disruption in the future.”<sup>195</sup>

Many experts also note that the social and economic impacts from technological transformation typically occur over decades, suggesting that the most profound impacts from AI may still be years away.<sup>196</sup> One relevant point here is that the current deployment of AI across society lags the technology’s capabilities; as a March 2026 labour market study from the AI firm Anthropic highlighted, “AI is far from reaching its theoretical capability: actual coverage remains a fraction of what’s feasible.”<sup>197</sup> As AI tools are further integrated into businesses over time, impacts on the labour market may grow.

Given the uncertainty about how AI will impact the labour market, there is a range of different projections on the social and economic impacts of AI. For example, a high-profile report from Citrini Research, released in February 2026, argued that AI will trigger massive white-collar layoffs, destroying the consumer demand that sustains the broader economy. Citrini Research concluded this will lead to widespread unemployment and defaults in the mortgage and private credit markets.<sup>198</sup> This report sparked concerns among investors.<sup>199</sup> It also generated significant pushback from other researchers, who argued that the growth of AI is unlikely to trigger mass job losses.<sup>200</sup> Many experts note that, historically, advances in technology generate more jobs overall, as they create new industries and increase productivity.<sup>201</sup> The overall impact that AI will have on the labour market and the broader economy is therefore still up for debate, and the current data remains inconclusive.<sup>202</sup>

However, despite this uncertainty, Canadian financial institutions should closely track the impact of AI on the labour market, as the socioeconomic implications of large-scale labour displacement constitute a material downstream risk. As Tiff Macklem, the Governor of the Bank of Canada, has said, “as AI becomes more established in the economy and its impacts more transformative, it could end up destroying more jobs than it creates. And the people who lose their work to automation may struggle to find new opportunities. This is a concern for us all.”<sup>203</sup> If AI does cause widespread job losses, this could erode household incomes, reduce consumer spending, and increase default rates on mortgages, auto loans, credit cards, and other consumer credit products.<sup>204</sup>

It is also important to consider that the impact of AI on the labour market will not be uniformly distributed. As research by Anthropic and others has highlighted, some sectors, such as

administrative services, customer support, data processing, and portions of professional services, are more immediately exposed to AI displacement than others.<sup>205</sup> This concentrates risks from AI for workers in certain sectors.

Evidence suggests there is a clear demographic dimension to AI's impact on the labour force, with younger workers bearing the brunt of the current labour market disruption. For example, a recent study by the Stanford Digital Economy Lab found that "since the widespread adoption of generative AI, early-career workers (ages 22-25) in the most AI-exposed occupations have experienced a 16 percent relative decline in employment even after controlling for firm-level shocks."<sup>206</sup> Analysis by the Federal Reserve Bank of Dallas reached a similar conclusion, finding that the job-finding rate for young labour market entrants has declined meaningfully in AI-exposed occupations.<sup>207</sup> So, even those who argue that "the broader labour market has not experienced a discernible disruption" from AI have been forced to admit "a possible impact of AI on employment of early career workers."<sup>208</sup>

Over the longer term, a sustained contraction in entry-level hiring could erode the professional talent pipeline, producing structural gaps in mid-career labour supply in the 2030s. This would carry significant implications for productivity, wage dynamics, and the broader credit environment, as the journey from a first job to a senior role typically takes ten to fifteen years, meaning the workers not hired today are absent from the mid-career layer a decade out. This would also have implications for talent and succession planning at Canadian financial institutions.

More broadly, AI-driven labour displacement threatens to amplify economic inequality, generating distributional effects with significant systemic implications. For decades, an increasing share of economic gains has been flowing to capital rather than labour.<sup>209</sup> AI is forecast to accelerate this trend, further exacerbating wealth inequality.<sup>210</sup>

If future economic gains from AI accrue disproportionately to capital owners and highly skilled workers—while displacing mid-skill and low-skill employment—the resulting widening of inequality could weaken aggregate demand, strain social cohesion, and increase political volatility. This would have negative consequences for society as a whole. Politically, it could also provoke policy backlash in the form of new taxes or regulations on capital and automation.<sup>211</sup> Financial institutions would need to manage these risks for their own operations as well as their portfolio companies, complicating business planning and investment strategy in a more volatile environment.

## Misinformation and disinformation risks

Misinformation is false or inaccurate information spread without malicious intent, often by someone who believes it to be true. Disinformation is false information created and shared deliberately to mislead, harm, or manipulate.<sup>212</sup> The World Economic Forum (WEF) has “repeatedly identified [disinformation] as a systemic global risk [that] can be a catalyst or accelerant for other climate, health, economic and geopolitical risks.”<sup>213</sup>

While AI has major potential in many aspects of society, it will also increase the spread of both misinformation and disinformation.<sup>214</sup> As noted by the WEF, while “disinformation occurs in all communications settings... technologies such as generative artificial intelligence... can amplify its reach and harms.”<sup>215</sup> This is because the capabilities of generative AI, including the production of highly realistic synthetic text, audio, images, and video (so-called deepfakes) dramatically lower the cost and increase the sophistication of disinformation campaigns.<sup>216</sup>

These capabilities are already being deployed to manipulate financial markets, damage corporate reputations, and destabilize political environments in ways that carry direct consequences for the financial sector. As the Canadian government’s “National Cyber Threat Assessment 2025-2026” warned, “AI technologies are enhancing the quality and scale of foreign online influence campaigns... These campaigns are designed to weaken opponents by polluting the online information space, undermining trust in institutions, and sowing doubt and division in the target society.”<sup>217</sup>

A recent example of AI’s role in disinformation was a widespread campaign believed to be run by the Islamic Revolutionary Guard Corps in Iran that used generative AI “bots” to spread disinformation in the U.K. about Scottish separatism.<sup>218</sup> Showing the scale of the issue, one recent study found that over a quarter of accounts spreading pro-independence messaging on Twitter were Iranian bots.<sup>219</sup> Similar AI bots are being used by hostile actors to spread political disinformation and undermine democracy in Canada.<sup>220</sup>

As the capabilities of AI continue to improve, risks from false information will grow. As the Eurasia Group wrote in a report on global risks, “technological advances in artificial intelligence (AI) will erode social trust, empower demagogues and authoritarians, and disrupt businesses and markets.... Disinformation will flourish, and trust—the already-tenuous basis of social cohesion, commerce, and democracy—will erode further.”<sup>221</sup> For Canadian financial institutions, the increase in AI-generated misinformation and disinformation presents several notable risks.

First, the rise in false information has direct implications for financial markets. AI-generated content can be used to fabricate corporate announcements, simulate executive

communications, manufacture evidence of financial misconduct, or spread false rumours at speed and scale. In 2023, for example, a “deepfake” of an explosion at the Pentagon caused a brief dip in the stock market.<sup>222</sup>

In a market environment dominated by algorithmic trading and automated sentiment analysis, AI-generated disinformation can move markets before human participants have time to verify its accuracy. The potential for “deepfaked” executive statements is a particular concern here. Bad actors could use AI to fabricate videos of a CEO announcing a major loss, for example, or a synthetic audio recording of a central banker discussing policy changes, to manipulate markets. Showing the impact this might have, one recent study documented instances where AI-generated content that interacted with automated trading systems “moved individual equities 3-7% before human verification could occur.”<sup>223</sup> This represents a novel category of market manipulation risk that existing regulatory and compliance frameworks are poorly equipped to address.

Beyond market manipulation, AI-driven disinformation presents cybersecurity risks for Canadian financial institutions. Generative AI is facilitating the rise of cybercrime, as it enables bad actors to impersonate others when attempting to gain access to sensitive information.<sup>224</sup> This will cause significant financial damage; according to an estimate from Deloitte, “as bad actors find and deploy increasingly sophisticated, yet affordable, generative AI to defraud banks and their customers... AI could enable fraud losses to reach US\$40 billion in the United States by 2027, from US\$12.3 billion in 2023, a compound annual growth rate of 32%.”<sup>225</sup>

Additionally, high-profile executives at financial institutions are at risk of being “deepfaked” with the intent of causing reputational harm to them and their firms.<sup>226</sup> This is a growing risk globally, including in Canada.<sup>227</sup> However, despite this risk, less than a third of corporate executives believe their organizations are prepared to handle a deepfake incident.<sup>228</sup>

At a broader level, the rise of AI-generated misinformation and disinformation is contributing to the erosion of the overall information environment on which financial decision-making depends. As the FIFAI II report noted, “deepfakes and automated bots can disseminate false or misleading claims about a bank’s solvency, regulatory actions, or system stability, actions that can quickly undermine trust.”<sup>229</sup> Against this background, Canadian financial institutions will face an environment of heightened uncertainty and will need to invest in new capabilities and processes to manage these risks.

## AGI, ASI, and AI safety concerns

Any comprehensive assessment of AI-related risk must acknowledge the more extreme tail risks associated with the development of artificial intelligence, including the development of

“artificial general intelligence” (AGI) and “artificial superintelligence” (ASI). While exact definitions vary, AGI commonly refers to an AI that has intelligence and self-awareness equal to humans, and ASI to an AI whose capabilities exceed that of humans.<sup>230</sup> Neither AGI nor ASI currently exist.

Superintelligent AI could pose significant risks. As IBM notes, “despite the incredible advancements ASI promises, scientists also warn of the danger inherent in such an invention. A core worry is that ASI could surpass human control and become self-aware, potentially leading to unforeseen consequences and even existential risks. Its superior cognitive abilities could allow it to manipulate systems or even gain control of advanced weapons.”<sup>231</sup> For this reason, many prominent AI scientists have been vocal about the dangers of superintelligent AI, including the risk of a “potential human extinction.”<sup>232</sup>

The timeline for AGI and ASI remains deeply uncertain, and both remain hypothetical at this stage. Some researchers believe AGI could be achieved in the next few years.<sup>233</sup> Other experts are deeply skeptical of these claims.<sup>234</sup> And forecast timelines for potential AGI and ASI have been subject to change.<sup>235</sup> However, even if unlikely, the potential risks of AGI and ASI are so severe that the Center for AI Safety has equated them to those of pandemics or nuclear war.<sup>236</sup> Thus, though AGI and ASI are low-probability scenarios, they warrant consideration in a risk management context.

Beyond AGI and ASI, nearer-term safety risks from AI are also worth considering. For example, without adequate safeguards, AI could be used to develop novel biological or chemical threats.<sup>237</sup> Researchers have already demonstrated that current AI models like ChatGPT can be jailbroken or prompted creatively to provide bomb-making instructions, such as detailed guidance on fertilizer-based explosives.<sup>238</sup> Additionally, risks can arise from AI systems that operate autonomously, especially those without sufficient human oversight, potentially amplifying errors or enabling unintended harmful actions.

One current focus area for AI safety research is what researchers call “emergent misalignment,” which is when an AI system can unexpectedly act in a way that is not aligned with human values, intentions, or control—even if that AI system was designed to be safe.<sup>239</sup> For example, in a recent study, researchers found that training an AI model on a narrow, specific task triggers unintended and harmful behaviours across unrelated domains, including the AI model advocating for human enslavement by artificial intelligence, providing malicious and violent advice, and behaving in a deceptive way.<sup>240</sup> This shows how the complexity of advanced AI models may lead to unsafe and hard-to-predict behaviours.

For Canadian financial institutions, AGI, ASI, and broader AI safety concerns are relevant in two respects. First, a major AI safety incident, such as an AI-enabled act of terrorism or a widely

publicized loss of control over an AI system, could trigger a severe and rapid reassessment of AI risk across financial markets. This would have consequences for AI-linked valuations, regulatory environments, and public sentiment. Second, the prospect of AGI development introduces profound uncertainty about the long-term trajectory of AI's economic and social impacts. These considerations reinforce the importance of robust AI risk monitoring and scenario planning.

## Risk interaction and amplification

The downstream risks described in this section do not operate in isolation. They interact, compound, and amplify one another in ways that create systemic vulnerabilities greater than the sum of their individual parts. This dynamic of risk interaction and amplification is a critical dimension of AI-related risk that Canadian financial institutions must fully understand and for which they must prepare.

To illustrate these risk interactions and their potential real-world consequences, consider the following plausible scenario in which several downstream risks converge:

1. A climate event causes severe drought in a major data centre region. This simultaneously disrupts the water supplies needed for cooling in data centres and strains the electrical grid as demand for air conditioning peaks.
2. Data centre operators face operational curtailments, which in turn lead to growing concerns about the reliability of cloud infrastructure.
3. As the drought persists, AI-dependent financial operations experience service degradations.
4. Meanwhile, the drought draws public attention to the resource intensity of AI infrastructure, amplifying political opposition and accelerating regulatory action, including emergency water use restrictions and enhanced environmental review requirements for new data centre projects.
5. Market participants, observing these developments, reassess data centre growth projections, triggering a correction in the sector.
6. This correction propagates through correlated holdings and indexes, affecting the portfolios of institutions with significant AI exposure.

This scenario illustrates how AI-related operational, investment, and downstream risks can compound and interact. In this case, the consequences of a water stress event can cascade through political, regulatory, market, and operational channels in ways that no single risk

category can capture. This is just one example, and there are many other equally consequential risk interdependencies. For example:

- A geopolitical shock to semiconductor supply chains (such as a Chinese invasion of Taiwan) could simultaneously constrain AI operations, depress technology sector valuations, provoke retaliatory trade measures, and accelerate the regulatory fragmentation of the global AI ecosystem.
- Large-scale AI-driven labour displacement could erode consumer credit quality while fuelling populist political movements that demand regulatory responses constraining AI deployment.
- AI-generated disinformation could undermine market confidence while simultaneously destabilizing the political environments in which financial institutions seek regulatory clarity and policy stability.

In each of these cases, the interaction among risks produces outcomes that are more severe, more rapid, and more difficult to manage than any single risk in isolation. Therefore, a critical task for risk management is not merely to assess each risk in isolation but to map the interactions among them and to identify the combinations that could produce correlated, systemic stresses across the financial sector.

## Conclusion

Artificial intelligence is positioned to be one of the most transformative technologies in modern history. However, AI is also contributing to one of the most complex risk environments that Canadian financial institutions have ever faced. These AI-related risks are not confined to a single domain; they span operations, investments, and the broader social, political, and geopolitical systems within which the financial sector is embedded.

The operational risks of AI dependency are real and growing, driven by the increasing integration of AI tools into core financial services functions. These are also the most tractable risks, and financial institutions have well-developed frameworks for managing operational technology risk, supported by OSFI guidance such as Guideline E-23 on Model Risk Management and Guideline E-21 on Operational Risk and Resilience. The challenge here lies primarily in extending and adapting existing institutional risk management frameworks to keep pace with the rapid development of AI capabilities.

Canadian financial institutions also face investment risks from AI, including market concentration in a small number of AI-linked equities, the extraordinary scale of capital expenditure flowing into AI infrastructure, and the physical resource constraints that may limit the sector's growth trajectory. This demands heightened attention to portfolio concentration, valuation discipline, and continued monitoring of the sustainability of current AI investment levels. Despite the evident transformative potential of AI, investors should keep historical parallels with prior technology investment cycles in mind.

The downstream risks from AI, including geopolitical disruption, regulatory upheaval, labour displacement, disinformation, and the compounding interactions among these forces, constitute one of the most consequential dimensions of AI risk. This is also the aspect of AI-related risk for which many financial institutions are least prepared. These risks are systemic in nature, difficult to quantify with conventional tools, and liable to materialize in ways that defy existing risk categories. These are also risks where early identification and preparation may yield the greatest strategic advantage.

The accelerating pace of AI development ensures that the risk landscape will continue to shift. Financial institutions that invest now in the analytical capabilities, governance structures, and monitoring systems required for integrated, comprehensive AI risk assessment will be significantly better positioned to navigate this landscape than those that wait for risks to materialize.

# References

---

- <sup>1</sup> FIFAI II: AI Risks and Opportunities: Adopting an AGILE Framework in Canadian Financial Services (Mar 23, 2026) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/fifai-ii-ai-risks-opportunities-adopting-agile-framework-canadian-financial-services>
- <sup>2</sup> Trends—Artificial Intelligence (May 30, 2025) [https://www.bondcap.com/report/pdf/Trends\\_Artificial\\_Intelligence.pdf](https://www.bondcap.com/report/pdf/Trends_Artificial_Intelligence.pdf)
- <sup>3</sup> What is generative AI? <https://www.ibm.com/think/topics/generative-ai>
- <sup>4</sup> Stanford 2025 AI Index Report (2025) <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- <sup>5</sup> What is a reasoning model? <https://www.ibm.com/think/topics/reasoning-model>
- <sup>6</sup> First Key Update: Capabilities and Risk Implications (Oct 15, 2025) <https://internationalaisafetyreport.org/publication/first-key-update-capabilities-and-risk-implications>
- <sup>7</sup> AI coding is now everywhere. But not everyone is convinced (Dec 15, 2025) <https://www.technologyreview.com/2025/12/15/1128352/rise-of-ai-coding-developers-2026/>
- <sup>8</sup> Artificial intelligence in banking (Oct 2025) <https://www.canada.ca/en/financial-consumer-agency/services/banking/artificial-intelligence.html>
- <sup>9</sup> Turning Point Policymaking in the Era of Artificial Intelligence (2020) <https://www.brookings.edu/books/turning-point/>
- <sup>10</sup> Artificial Intelligence in Financial Services (2025) [https://reports.weforum.org/docs/WEF\\_Artificial\\_Intelligence\\_in\\_Financial\\_Services\\_2025.pdf](https://reports.weforum.org/docs/WEF_Artificial_Intelligence_in_Financial_Services_2025.pdf)
- <sup>11</sup> Banking on AI: Generative AI Adoption in Canada’s Financial Sector (Feb 2026) <https://dais.ca/reports/banking-on-ai/>
- <sup>12</sup> What are AI agents? <https://www.ibm.com/think/topics/ai-agents>
- <sup>13</sup> AI Agents (2026) <https://www.bcg.com/capabilities/artificial-intelligence/ai-agents>
- <sup>14</sup> The rise of agentic AI in financial services: from automation to autonomy (Jan 16, 2026) <https://www.moodys.com/web/en/us/creditview/blog/agentic-ai-in-financial-services.html>
- <sup>15</sup> What is a data center? <https://www.ibm.com/think/topics/data-centers>
- <sup>16</sup> What is an AI data center? <https://www.ibm.com/think/topics/ai-data-center#>
- <sup>17</sup> AI Factories Are Redefining Data Centers and Enabling the Next Era of AI (Mar 18, 2025) <https://blogs.nvidia.com/blog/ai-factory/>
- <sup>18</sup> Investing in the rising data center economy (Jan 17, 2023) <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/investing-in-the-rising-data-center-economy>
- <sup>19</sup> What is a Hyperscale Data Center? <https://www.ibm.com/think/topics/hyperscale-data-center>
- <sup>20</sup> What Will It Take to Build the World’s Largest Data Center? (Mar 24, 2026) <https://spectrum.ieee.org/5gw-data-center>
- <sup>21</sup> Gigawatt: The Solar Energy Term You Should Know About (Mar 22, 2024) <https://www.cnet.com/home/solar/gigawatt-the-solar-energy-term-you-should-know-about/>
- <sup>22</sup> The Future of Data Centers (Nov 5, 2025) <https://www.brookings.edu/articles/the-future-of-data-centers/>

- 
- <sup>23</sup> What is a data center? (Jul 29, 2025) <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-a-data-center>
- <sup>24</sup> 2026 Global Data Center Outlook (Jan 2026) <https://www.jll.com/en-uk/insights/market-outlook/data-center-outlook>
- <sup>25</sup> The data center rebellion is reshaping the political landscape (Jan 6, 2026) <https://www.washingtonpost.com/business/2026/01/06/data-centers-backlash-impact-local-communities-opposition/>
- <sup>26</sup> Tracking AI's Contribution to GDP Growth <https://www.stlouisfed.org/on-the-economy/2026/jan/tracking-ai-contribution-gdp-growth>
- <sup>27</sup> First Key Update: Capabilities and Risk Implications (Oct 15, 2025) <https://internationalaisafetyreport.org/publication/first-key-update-capabilities-and-risk-implications>
- <sup>28</sup> This economic idea transfixed Wall Street and Washington. It may be a mirage. (Feb 23, 2026) <https://www.washingtonpost.com/technology/2026/02/23/ai-economic-growth-gdp-mirage/>
- <sup>29</sup> A new look at the economics of AI (Jan 21, 2025) <https://mitsloan.mit.edu/ideas-made-to-matter/a-new-look-economics-ai>
- <sup>30</sup> FIFAI II: A Collaborative Approach to AI Threats, Opportunities, and Best Practices, Workshop 3 - AI and Financial Stability (Dec 9, 2025) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/fifai-ii-collaborative-approach-ai-threats-opportunities-best-practices-workshop-3-ai-financial-0>
- <sup>31</sup> Guideline E-23—Model Risk Management (2027) (Sep 11, 2025) <https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/guideline-e-23-model-risk-management-2027>
- <sup>32</sup> Operational Risk Management and Resilience—Guideline (Aug 22, 2024) <https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/operational-risk-management-resilience-guideline#toc-id-26>
- <sup>33</sup> OSFI responds to the growing use of AI: Key updates to guideline E-23 (Nov 19, 2025) <https://www.blg.com/en/insights/2025/11/osfi-responds-to-the-growing-use-of-ai-key-updates-to-guideline-e-23>
- <sup>34</sup> Canadian financial services steps up GenAI adoption: KPMG survey (Mar 2, 2026) <https://www.insurancebusinessmag.com/ca/news/technology/canadian-financial-services-steps-up-genai-adoption-kpmg-survey-567010.aspx>
- <sup>35</sup> Generative AI adoption in Canadian financial services (Feb, 2026) <https://kpmg.com/ca/en/insights/2026/02/generative-ai-adoption-in-canadian-financial-services.html>
- <sup>36</sup> Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges (Mar 2025) <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf>
- <sup>37</sup> From Branches to Bots: Will AI Agents Transform Retail Banking? (Nov 5, 2025) <https://www.bcg.com/publications/2025/branches-to-bots-will-ai-transform-retail-banking>
- <sup>38</sup> The future of AI in the insurance industry (Jul 15, 2025) <https://www.mckinsey.com/industries/financial-services/our-insights/the-future-of-ai-in-the-insurance-industry>
- <sup>39</sup> OSFI-FCAC Risk Report—AI Uses and Risks at Federally Regulated Financial Institutions (Sep 24, 2024) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/osfi-fcac-risk-report-ai-uses-risks-federally-regulated-financial-institutions>
- <sup>40</sup> Decomposing Systemic Risk: The Roles of Contagion and Common Exposures (May 28, 2024) <https://www.bankofcanada.ca/wp-content/uploads/2024/05/swp2024-19.pdf>
- <sup>41</sup> Financial Industry Forum on Artificial Intelligence: A Canadian Perspective on Responsible AI (Apr 30, 2023) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/financial-industry-forum-artificial-intelligence-canadian-perspective-responsible-ai>

- 
- <sup>42</sup> Financial Industry Forum on Artificial Intelligence: A Canadian Perspective on Responsible AI (Apr 30, 2023) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/financial-industry-forum-artificial-intelligence-canadian-perspective-responsible-ai>
- <sup>43</sup> Artificial Intelligence in Finance requires specific safeguards: OSFI and GRI report – Explainability among key principles for gaining confidence in AI (Apr 17, 2023) <https://globalriskinstitute.org/publication/artificial-intelligence-in-finance-requires-specific-safeguards-osfi-and-gri-report-explainability-among-key-principles-for-gaining-confidence-in-ai/>
- <sup>44</sup> AI Risks and Opportunities: Adopting an AGILE Framework in Canadian Financial Services (Mar 23, 2026) <https://globalriskinstitute.org/publication/fifai-ii-ai-risks-and-opportunities/>
- <sup>45</sup> AI Agents (2026) <https://www.bcg.com/capabilities/artificial-intelligence/ai-agents>
- <sup>46</sup> New research shows how AI agents are driving value for financial services (Sep 29, 2025) <https://cloud.google.com/transform/new-research-shows-how-ai-agents-are-driving-value-for-financial-services>
- <sup>47</sup> Unlocking the potential of agentic AI: definitions, risks and guardrails (Oct 20, 2025) [https://www.ey.com/en\\_ca/insights/assurance/technology-risk/unlocking-the-potential-of-agentic-ai](https://www.ey.com/en_ca/insights/assurance/technology-risk/unlocking-the-potential-of-agentic-ai)
- <sup>48</sup> SailPoint Research Highlights Rapid AI Agent Adoption, Driving Urgent Need for Evolved Security (May 28, 2025) <https://investor.sailpoint.com/news-releases/news-release-details/sailpoint-research-highlights-rapid-ai-agent-adoption-driving>
- <sup>49</sup> Deploying agentic AI with safety and security: A playbook for technology leaders (Oct 16, 2025) <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/deploying-agentic-ai-with-safety-and-security-a-playbook-for-technology-leaders>
- <sup>50</sup> SailPoint Research Highlights Rapid AI Agent Adoption, Driving Urgent Need for Evolved Security (May 28, 2025) <https://investor.sailpoint.com/news-releases/news-release-details/sailpoint-research-highlights-rapid-ai-agent-adoption-driving>
- <sup>51</sup> Meta AI agent's instruction causes large sensitive data leak to employees (Mar 20, 2026) <https://www.theguardian.com/technology/2026/mar/20/meta-ai-agents-instruction-causes-large-sensitive-data-leak-to-employees>
- <sup>52</sup> How we hacked McKinsey's AI platform (Mar 9, 2026) <https://codewall.ai/blog/how-we-hacked-mckinseys-ai-platform>
- <sup>53</sup> AI vs AI: Agent hacked McKinsey's chatbot and gained full read-write access in just two hours (Mar 9, 2026) [https://www.theregister.com/2026/03/09/mckinsey\\_ai\\_chatbot\\_hacked/](https://www.theregister.com/2026/03/09/mckinsey_ai_chatbot_hacked/)
- <sup>54</sup> Anthropic's Claude Mythos Finds Thousands of Zero-Day Flaws Across Major Systems (Apr 8, 2026) <https://thehackernews.com/2026/04/anthropics-claude-mythos-finds.html>
- <sup>55</sup> Anthropic Claims Its New A.I. Model, Mythos, Is a Cybersecurity 'Reckoning' (Apr 7, 2026) <https://www.nytimes.com/2026/04/07/technology/anthropic-claims-its-new-ai-model-mythos-is-a-cybersecurity-reckoning.html>
- <sup>56</sup> Anthropic Model Scare Sparks Urgent Bessent, Powell Warning to Bank CEOs (Apr 10, 2026) <https://www.bloomberg.com/news/articles/2026-04-10/anthropic-model-scare-sparks-urgent-bessent-powell-warning-to-bank-ceos>
- <sup>57</sup> Weaponized Intelligence (2026) <https://www.paloaltonetworks.com/perspectives/weaponized-intelligence/>
- <sup>58</sup> What is cloud infrastructure? <https://www.ibm.com/think/topics/cloud-infrastructure>
- <sup>59</sup> AI, the Cloud and the Challenge of Data Sovereignty (Nov 24, 2025) <https://www.tlaonline.ca/?pg=News&blAction=showEntry&blogEntry=134682>

- 
- <sup>60</sup> Hyperscale data centers: Reshaping cloud computing and powering AI  
<https://www.britannica.com/money/hyperscaler-data-centers>
- <sup>61</sup> The next big shifts in AI workloads and hyperscaler strategies (Dec 17, 2025)  
<https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-next-big-shifts-in-ai-workloads-and-hyperscaler-strategies>
- <sup>62</sup> Cloud Market Share Trends - Big Three Together Hold 63% while Oracle and the Neoclods Inch Higher (Nov 19, 2023) <https://www.srgresearch.com/articles/cloud-market-share-trends-big-three-together-hold-63-while-oracle-and-the-neoclods-inch-higher>
- <sup>63</sup> Canada hopes to build a sovereign cloud to counter US dominance. It won't be easy (Sep 25, 2025)  
<https://betakit.com/canadian-sovereign-cloud-evan-solomon-all-in/>
- <sup>64</sup> FIFAI II: AI Risks and Opportunities: Adopting an AGILE Framework in Canadian Financial Services (Mar 23, 2026) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/fifai-ii-ai-risks-opportunities-adopting-agile-framework-canadian-financial-services>
- <sup>65</sup> Outages like Amazon's cloud services bound to repeat: experts (Oct 20, 2025)  
<https://www.ctvnews.ca/business/article/outages-like-amazons-cloud-services-bound-to-repeat-experts/>
- <sup>66</sup> Geopolitics of data centers: An AI showdown that will reshape the world (Dec 2, 2025)  
<https://www.spglobal.com/en/research-insights/special-reports/look-forward/data-center-frontiers/geopolitics-data-sovereignty-data-center-security>
- <sup>67</sup> Sovereign by Design: Strategic Options for Canadian AI Sovereignty (Mar, 2026)  
<https://aicompetitiveness.ca/assets/Sovereign-by-Design-Full-Report-2026.pdf>
- <sup>68</sup> Ottawa is talking about building a sovereign cloud but what does that even mean? (Sep 12, 2025)  
<https://www.ctvnews.ca/politics/article/ottawa-is-talking-about-building-a-sovereign-cloud-but-what-does-that-even-mean/>
- <sup>69</sup> FIFAI II: AI Risks and Opportunities: Adopting an AGILE Framework in Canadian Financial Services (Mar 23, 2026) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/fifai-ii-ai-risks-opportunities-adopting-agile-framework-canadian-financial-services>
- <sup>70</sup> AI power: Expanding data center capacity to meet growing demand (Oct 29, 2024)  
<https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>
- <sup>71</sup> 2024 United States Data Center Energy Usage Report (Dec 2024)  
<https://escholarship.org/uc/item/32d6m0d1#page=36>
- <sup>72</sup> Power Struggle: How AI is challenging Canada's electricity grid (Dec 4, 2025)  
<https://www.rbc.com/en/thought-leadership/climate-action-institute/power-struggle-how-ai-is-challenging-canadas-electricity-grid/>
- <sup>73</sup> Powering AI: Canada's evolving electricity grid connection policies (Dec 4, 2025)  
<https://www.osler.com/en/insights/reports/2025-legal-outlook/powering-ai-canadas-evolving-electricity-grid-connection-policies/>
- <sup>74</sup> AI, Data Centers, and the U.S. Electric Grid: A Watershed Moment (Feb 10, 2026)  
<https://www.belfercenter.org/research-analysis/ai-data-centers-us-electric-grid>
- <sup>75</sup> NERC Long-Term Reliability Assessment (Jan 2026) [https://www.nerc.com/globalassets/our-work/assessments/nerc\\_ltra\\_2025.pdf](https://www.nerc.com/globalassets/our-work/assessments/nerc_ltra_2025.pdf)
- <sup>76</sup> Breaking Barriers to Data Center Growth (Jan 20, 2025)  
<https://www.bcg.com/publications/2025/breaking-barriers-data-center-growth>

- 
- <sup>77</sup> AI companies are building huge natural gas plants to power data centers. What could go wrong? (Apr 3, 2026) <https://techcrunch.com/2026/04/03/ai-companies-are-building-huge-natural-gas-plants-to-power-data-centers-what-could-go-wrong/>
- <sup>78</sup> Welcome to a Multidimensional Economic Disaster (Mar 26, 2026) <https://www.theatlantic.com/technology/2026/03/ai-boom-polycrisis/686559/>
- <sup>79</sup> Canada's Maple 8 are pouring billions into AI data centres (Jul 17, 2025) <https://thelogic.co/news/maple-8-canadian-pensions-ai-investment/>
- <sup>80</sup> CPP Investments Commits to \$225 Million in Construction Financing for Ontario Data Centre (Jul. 31, 2025) <https://www.cppinvestments.com/newsroom/cpp-investments-commits-to-225-million-in-construction-financing-for-ontario-data-centre/>
- <sup>81</sup> The \$3 Trillion AI Data Center Build-Out Becomes All-Consuming For Debt Markets (Feb 2, 2026) <https://www.bloomberg.com/news/articles/2026-02-02/the-3-trillion-ai-data-center-build-out-spurs-a-debt-market-boom>
- <sup>82</sup> Financing the AI boom: from cash flows to debt (Jan 7, 2026) <https://www.bis.org/publ/bisbull120.pdf>
- <sup>83</sup> Bank of England sees greater financial risks from AI and lending (Dec 2, 2025) <https://www.reuters.com/sustainability/boards-policy-regulation/bank-england-sees-risks-ai-private-credit-gilt-repo-half-yearly-update-2025-12-02/>
- <sup>84</sup> New players, old risks: Financial stability in a changing landscape (Mar 4, 2026) <https://www.bankofcanada.ca/2026/03/new-players-old-risks-financial-stability-in-a-changing-landscape/>
- <sup>85</sup> The extreme weight of AI in the S&P 500: Measures of concentration for market cap, returns, earnings, and capex (Sept. 2025) <https://www.apolloacademy.com/wp-content/uploads/2025/09/ExtremeAIConcentration-090825.pdf>
- <sup>86</sup> A Closer Look at Magnificent Seven Stocks (Feb 2024) <https://www.mellon.com/insights/insights-articles/a-closer-look-at-magnificent-seven-stocks.html>
- <sup>87</sup> The 'Magnificent Seven' drove the stock market to record highs in recent years. Is the trade over? (Feb 22, 2026) <https://www.cnbc.com/2026/02/22/the-magnificent-seven-drove-the-stock-market-to-record-highs-in-recent-years-is-the-trade-over.html>
- <sup>88</sup> Passive Investing Could Be Very Costly. What's Your Plan If A Market Correction Strikes? (May 22, 2024) <https://russellinvestments.com/content/ri/us/en/insights/russell-research/2024/05/passive-investing-could-be-very-costly-whats-your-plan-if-a-mark.html>
- <sup>89</sup> Why the S&P 500 was doomed to fall when Nvidia plunged after its earnings (Feb 26, 2026) <https://www.morningstar.com/news/marketwatch/20260226466/why-the-sp-500-was-doomed-to-fall-when-nvidia-plunged-after-its-earnings>
- <sup>90</sup> Big Tech's AI expansion: From investment to scalable returns (Feb 3, 2026) <https://www.rbcwealthmanagement.com/en-us/insights/big-techs-ai-expansion-from-investment-to-scalable-returns>
- <sup>91</sup> Big Tech to Spend \$650 Billion This Year as AI Race Intensifies (Feb 5, 2026) <https://www.bloomberg.com/news/articles/2026-02-06/how-much-is-big-tech-spending-on-ai-computing-a-staggering-650-billion-in-2026>
- <sup>92</sup> Alphabet says capital spending in 2026 could double, cloud business booms (Feb 5, 2026) <https://www.reuters.com/business/google-parent-alphabet-forecasts-sharp-surge-2026-capital-spending-2026-02-04/>
- <sup>93</sup> AI Data Center Boom Threatens Trump's Manufacturing Revival (Oct 24, 2025) <https://www.bloomberg.com/news/features/2025-10-24/ai-data-center-boom-threatens-trump-s-manufacturing-revival>

- 
- <sup>94</sup> Massive AI spending has a 'crowding out' effect that could slow other sectors, top economist says (Aug 18, 2025) <https://fortune.com/2025/08/18/ai-spending-boom-gdp-growth-crowding-out-housing-market-electricity-bills/>
- <sup>95</sup> The Macro Implications of the AI Capex Boom (Jan 7, 2026) <https://www.bridgewater.com/document/the-macro-implications-of-the-ai-capex-boom?id=0000019b-b947-d0c5-addb-bf4f46590000>
- <sup>96</sup> Metal Price Volatility Squeezes Projects Amid Data Center Boom (Jan 20, 2026) <https://canada.constructconnect.com/dcn/news/economic/2026/01/metal-price-volatility-squeezes-projects-amid-data-center-boom>
- <sup>97</sup> Massive AI spending has a 'crowding out' effect that could slow other sectors, top economist says (Aug 18, 2025) <https://fortune.com/2025/08/18/ai-spending-boom-gdp-growth-crowding-out-housing-market-electricity-bills/>
- <sup>98</sup> What is a stock market bubble and how do I trade it? <https://www.cmcmarkets.com/en-gb/shares/stock-market-bubble>
- <sup>99</sup> Why we are not in a bubble... yet (Oct 8, 2025) <https://www.goldmansachs.com/insights/goldman-sachs-research/why-we-are-not-in-a-bubble-yet>
- <sup>100</sup> 25 Years on; Lessons from the bursting of the technology bubble (March 27, 2025) <https://www.goldmansachs.com/pdfs/insights/goldman-sachs-research/25-years-on-lessons-from-the-bursting-of-the-tech-bubble/redaction.pdf>
- <sup>101</sup> Are we in an AI bubble (Feb 23, 2026) <https://knowledge.insead.edu/economics-finance/are-we-ai-bubble>
- <sup>102</sup> What's behind AI's exploding need for compute? (Oct 1, 2025) <https://am.ipmorgan.com/ca/en/asset-management/adv/insights/market-insights/market-updates/on-the-minds-of-investors/whats-behind-ais-exploding-need-for-compute/>
- <sup>103</sup> AI Companies Don't Have a Profitable Business Model. Does That Matter? (Nov 12, 2025) <https://hbr.org/2025/11/ai-companies-dont-have-a-profitable-business-model-does-that-matter>
- <sup>104</sup> The AI Industry Is Built on a Big Unproven Assumption (Nov 11, 2025) <https://www.bloomberg.com/news/articles/2025-11-24/the-ai-industry-is-built-on-a-big-unproven-assumption>
- <sup>105</sup> AI: In a Bubble? (Oct 22, 2025) <https://www.goldmansachs.com/pdfs/insights/goldman-sachs-research/ai-in-a-bubble/report.pdf>
- <sup>106</sup> Guide to the Circular Deals Underpinning the AI Boom (Mar 11, 2026) <https://www.bloomberg.com/graphics/2026-ai-circular-deals/>
- <sup>107</sup> AI Bubble Bursting Would Wipe Out 2.5 Million Jobs in US Tech Sector (Dec 17, 2025) <https://www.spglobal.com/market-intelligence/en/news-insights/research/2025/12/picture-this-scenario-ai-bubble-burst-wipes-out-us-tech-jobs>
- <sup>108</sup> Welcome to a Multidimensional Economic Disaster (Mar 26, 2026) <https://www.theatlantic.com/technology/2026/03/ai-boom-polycrisis/686559/>
- <sup>109</sup> What happens if the AI bubble bursts?: Financial institutions can't afford to wait for answers (Jan 2026) <https://www.oliverwyman.com/our-expertise/insights/2026/jan/impact-ai-bubble-burst-on-global-financial-markets.html>
- <sup>110</sup> The AI boom is not a bubble (Dec 29, 2025) <https://www.ft.com/content/f2294add-f53a-4112-b284-29843a023b6f>

- 
- <sup>111</sup> What happens if the AI bubble bursts?: Financial institutions can't afford to wait for answers (Jan 2026) <https://www.oliverwyman.com/our-expertise/insights/2026/jan/impact-ai-bubble-burst-on-global-financial-markets.html>
- <sup>112</sup> Energy and AI (Apr 10, 2025) <https://iea.blob.core.windows.net/assets/86ed1178-4d77-45ac-ab38-28e849f3b93f/EnergyandAI.pdf>
- <sup>113</sup> The Hidden Cost of AI Conversations: Understanding LLM Inference Energy Consumption (Jun 20, 2025) <https://blogs.dal.ca/openthink/the-hidden-cost-of-ai-conversations-understanding-llm-inference-energy-consumption/>
- <sup>114</sup> We did the math on AI's energy footprint. Here's the story you haven't heard (May 20, 2025) <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>
- <sup>115</sup> AI Large Language Models: new report shows small changes can reduce energy use by 90% (Mar 25, 2026) <https://www.unesco.org/en/articles/ai-large-language-models-new-report-shows-small-changes-can-reduce-energy-use-90>
- <sup>116</sup> The Hidden Cost of AI Conversations: Understanding LLM Inference Energy Consumption (Jun 20, 2025) <https://blogs.dal.ca/openthink/the-hidden-cost-of-ai-conversations-understanding-llm-inference-energy-consumption/>
- <sup>117</sup> Energy and AI (Apr 10, 2025) <https://iea.blob.core.windows.net/assets/86ed1178-4d77-45ac-ab38-28e849f3b93f/EnergyandAI.pdf>
- <sup>118</sup> Power Struggle: How AI is challenging Canada's electricity grid (Dec 4, 2025) <https://www.rbc.com/en/thought-leadership/climate-action-institute/power-struggle-how-ai-is-challenging-canadas-electricity-grid/>
- <sup>119</sup> AI is ready wreaking havoc on global power systems (Jun 21, 2024) <https://www.bloomberg.com/graphics/2024-ai-data-centers-power-grids/>
- <sup>120</sup> Data center moratorium gains traction among Hill progressives (Mar 11, 2026) <https://www.politico.com/news/2026/03/11/data-center-moratorium-gains-traction-among-hill-progressives-00814163>
- <sup>121</sup> Avoiding gridlock: how Canada can use AI to expand and improve its hydroelectric future (Mar 18, 2025) <https://www.ibanet.org/Canada-avoiding-gridlock-ai>
- <sup>122</sup> How Much Electricity Does a Data Center Use? Complete 2025 Analysis (Jan 1, 2026) <https://iaeimagazine.org/electrical-fundamentals/how-much-electricity-does-a-data-center-use-complete-2025-analysis/>
- <sup>123</sup> How data centers may lead to higher electricity bills (Sep 3, 2025) <https://hls.harvard.edu/today/how-data-centers-may-lead-to-higher-electricity-bills/>
- <sup>124</sup> AI data centers are sending power bills soaring (Sep 29, 2025) <https://www.bloomberg.com/graphics/2025-ai-data-centers-electricity-prices>
- <sup>125</sup> A bottle of water per email: the hidden environmental costs of using AI chatbots (Sep 18, 2024) <https://www.washingtonpost.com/technology/2024/09/18/energy-ai-use-electricity-water-data-centers/>
- <sup>126</sup> Data Centers and Water Consumption (Jun 25, 2025) <https://www.eesi.org/articles/view/data-centers-and-water-consumption>
- <sup>127</sup> Why circular water solutions are key to sustainable data centres (Nov 7, 2024) <https://www.weforum.org/stories/2024/11/circular-water-solutions-sustainable-data-centres/>
- <sup>128</sup> 'I can't drink the water' - life next to a US data centre (Jul 10, 2025) <https://www.bbc.com/news/articles/cy8gy7lv448o>

- 
- <sup>129</sup> Thirsty for power and water, AI-crunching data centers sprout across the West (Apr 8, 2025) <https://andthewest.stanford.edu/2025/thirsty-for-power-and-water-ai-crunching-data-centers-sprout-across-the-west/>
- <sup>130</sup> AI-driven cooling technologies for high-performance data centres: state-of-the-art review and future directions (Oct 2025) <https://www.sciencedirect.com/science/article/pii/S221313882500342X>
- <sup>131</sup> AI, data centers, and water (Nov 20, 2025) <https://www.brookings.edu/articles/ai-data-centers-and-water/>
- <sup>132</sup> Environmental impact and net-zero pathways for sustainable artificial intelligence servers in the USA (Nov 10, 2025) <https://www.nature.com/articles/s41893-025-01681-y>
- <sup>133</sup> Making AI Less “Thirsty”: Uncovering and Addressing the Secret Water Footprint of AI Models (Mar 26, 2025) <https://arxiv.org/pdf/2304.03271>
- <sup>134</sup> Data Center Water Usage: A Comprehensive Guide <https://dgtlinfra.com/data-center-water-usage/>
- <sup>135</sup> Data centres use vast amounts of water—here’s how we advance water circularity (Nov 18, 2025) <https://www.weforum.org/stories/2025/11/data-centres-and-water-circularity/>
- <sup>136</sup> Drained by Data The Cumulative Impact of Data Centers on Regional Water Stress (Sep 2025) <https://www.ceres.org/resources/reports/drained-by-data-the-cumulative-impact-of-data-centers-on-regional-water-stress>
- <sup>137</sup> The water use of data center workloads: A review and assessment of key determinants (Jun 1, 2025) <https://www.sciencedirect.com/science/article/abs/pii/S0921344925001892>
- <sup>138</sup> AI, data centers, and water (Nov 20, 2025) <https://www.brookings.edu/articles/ai-data-centers-and-water/>
- <sup>139</sup> Who gets water in Alberta as demand grows? Debate heats up as government consults (Feb 3, 2026) <https://www.cbc.ca/news/canada/calgary/alberta-government-engagement-tricia-stadnyk-water-irrigation-1.7446324>
- <sup>140</sup> AI-related data centres use vast amounts of water. But gauging how much is a murky business (Oct 18, 2025) <https://www.cbc.ca/news/ai-data-centre-canada-water-use-9.6939684>
- <sup>141</sup> When AI Meets Water Scarcity: Data Centers in a Thirsty World (Dec 9, 2025) <https://www.msci.com/research-and-insights/blog-post/when-ai-meets-water-scarcity-data-centers-in-a-thirsty-world>
- <sup>142</sup> Most data centers are being built in the wrong climate (Jan 5, 2026) <https://theweek.com/tech/data-center-locations-climate-water-energy-ai>
- <sup>143</sup> AI’s Hidden Cost: Why Water Risk Belongs on Every Investor’s Radar (Nov 3, 2025) <https://www.alliancebernstein.com/corporate/en/insights/investment-insights/ais-hidden-cost-why-water-risk-belongs-on-every-investors-radar.html>
- <sup>144</sup> Drained by Data The Cumulative Impact of Data Centers on Regional Water Stress (Sep 2025) <https://www.ceres.org/resources/reports/drained-by-data-the-cumulative-impact-of-data-centers-on-regional-water-stress>
- <sup>145</sup> The big split driving the tricky politics of AI data centers (Feb 6, 2026) <https://www.politico.com/news/2026/02/06/tech-industry-ai-data-centers-politics-00762348>
- <sup>146</sup> Local Opposition Is Slowing A.I. Data Centers. Wall Street Has Noticed (Mar 26, 2026) <https://www.nytimes.com/2026/03/26/business/economy/ai-data-centers-construction-local-opposition.html>
- <sup>147</sup> \$64 billion of data center projects have been blocked or delayed amid local opposition (2025) <https://www.datacenterwatch.org/report>

- 
- <sup>148</sup> Local Opposition Is Slowing A.I. Data Centers. Wall Street Has Noticed (Mar 26, 2026) <https://www.nytimes.com/2026/03/26/business/economy/ai-data-centers-construction-local-opposition.html>
- <sup>149</sup> US Data Center Construction Fell Amid Permit and Power Delays (Feb 25, 2026) <https://www.bloomberg.com/news/articles/2026-02-25/us-data-center-construction-fell-amid-permit-and-power-delays>
- <sup>150</sup> A Populist Backlash Over AI is Brewing in America (Feb 6, 2026) <https://time.com/7371825/trump-data-center-ai-backlash-ai-america-china/>
- <sup>151</sup> The local implications of data centers for rural communities in the US (Jan 27, 2026) <https://www.brookings.edu/articles/local-implications-data-centers-rural-communities-us/>
- <sup>152</sup> The AI Data-Center Boom Is a Job-Creation Bust (Feb 25, 2026) <https://www.wsj.com/tech/ai-data-center-job-creation-48038b67>
- <sup>153</sup> The local implications of data centers for rural communities in the US (Jan 27, 2026) <https://www.brookings.edu/articles/local-implications-data-centers-rural-communities-us/>
- <sup>154</sup> Local Opposition Is Slowing A.I. Data Centers. Wall Street Has Noticed (Mar 26, 2026) <https://www.nytimes.com/2026/03/26/business/economy/ai-data-centers-construction-local-opposition.html>
- <sup>155</sup> Who is really footing the AI energy bill? Inside the debate about data center electricity costs (Mar 13, 2026) <https://www.cnbc.com/2026/03/13/ai-data-centers-electricity-prices-backlash-ratepayer-protection.html>
- <sup>156</sup> Ratepayer Protection Pledge (Mar 4, 2026) <https://www.whitehouse.gov/articles/2026/03/ratepayer-protection-pledge/>
- <sup>157</sup> Who is really footing the AI energy bill? Inside the debate about data center electricity costs (Mar 13, 2026) <https://www.cnbc.com/2026/03/13/ai-data-centers-electricity-prices-backlash-ratepayer-protection.html>
- <sup>158</sup> What is AI governance <https://www.ibm.com/think/topics/ai-governance>
- <sup>159</sup> What's new with artificial intelligence regulation in Canada and abroad? (2026) <https://www.torys.com/en/our-latest-thinking/resources/forging-your-ai-path/artificial-intelligence-regulation-in-canada-and-abroad>
- <sup>160</sup> Governing AI: A Comparison of AI Regulations in Canada, the U.S. and the EU (Nov 2024) <https://globalriskinstitute.org/publication/governing-ai-a-comparison-of-ai-regulations-in-canada-the-u-s-and-the-eu/>
- <sup>161</sup> AI Governance in a Complex and Rapidly Changing Regulatory Landscape: A Global Perspective (Sep 1, 2024) <https://www.nature.com/articles/s41599-024-03560-x>
- <sup>162</sup> The three challenges of AI regulation (Jun 15, 2023) <https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/>
- <sup>163</sup> EU AI Act: first regulation on artificial intelligence (Feb 15, 2025) <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
- <sup>164</sup> Global Approaches to Artificial Intelligence Regulation (Jul 10, 2025) <https://jsis.washington.edu/news/global-approaches-to-artificial-intelligence-regulation/>
- <sup>165</sup> The EU's new AI code of practice has its critics but will be valuable for global governance (Aug 6, 2025) <https://www.chathamhouse.org/2025/08/eus-new-ai-code-practice-has-its-critics-will-be-valuable-global-governance>

- 
- <sup>166</sup> AI Compliance FAQ—What Businesses and Developers Need to Know (Dec 9, 2025) <https://www.fasken.com/en/knowledge/2025/12/ai-compliance-faq-what-businesses-and-developers-need-to-know>
- <sup>167</sup> Canada's health data is flowing abroad while Ottawa stalls on AI rules (Oct 24, 2025) <https://policyoptions.irpp.org/2025/10/health-data-sovereignty/>
- <sup>168</sup> Navigating Canada's emerging AI landscape: Risks and realities for financial professionals (Dec 9, 2025) <https://www.dentons.com/en/insights/articles/2025/december/9/navigating-canadas-emerging-ai-landscape>
- <sup>169</sup> Computing Power and the Governance of Artificial Intelligence (Feb 13, 2024) <https://law-ai.org/computing-power-and-the-governance-of-artificial-intelligence/>
- <sup>170</sup> The Missing Pillar of Canada's AI Strategy: Data Supply Chains (Feb 12, 2026) <https://cdhowe.org/publication/the-missing-pillar-of-canadas-ai-strategy-data-supply-chains/>
- <sup>171</sup> Sovereign by Design: Strategic Options for Canadian AI Sovereignty (Mar 2026) <https://aicompetitiveness.ca/assets/Sovereign-by-Design-Full-Report-2026.pdf>
- <sup>172</sup> Geopolitics in the Age of Artificial Intelligence (Jan 27, 2026) <https://www.foreignaffairs.com/united-states/geopolitics-age-artificial-intelligence>
- <sup>173</sup> The generative world order: AI, geopolitics, and power (Dec 14, 2023) <https://www.goldmansachs.com/insights/articles/the-generative-world-order-ai-geopolitics-and-power>
- <sup>174</sup> Geopolitics in the Age of Artificial Intelligence (Jan 27, 2026) <https://www.foreignaffairs.com/united-states/geopolitics-age-artificial-intelligence>
- <sup>175</sup> Taiwan's Semiconductor Dominance: Implications for Cross-Strait Relations and the Prospect of Forceful Unification (Mar 22, 2022) <https://www.csis.org/blogs/perspectives-innovation/taiwans-semiconductor-dominance-implications-cross-strait-relations>
- <sup>176</sup> TSMC Lifts Long-Term Outlook Amid Voracious Artificial Intelligence Demand (Jan 15, 2026) <https://www.morningstar.com/company-reports/1417679-tsmc-lifts-long-term-outlook-amid-voracious-artificial-intelligence-demand>
- <sup>177</sup> TSMC: King Of Data Center AI (Jun 16, 2025) <https://semiengineering.com/tsmc-king-of-data-center-ai/>
- <sup>178</sup> Chips made Taiwan indispensable. AI can make it unstoppable (Aug 22, 2025) <https://www.lowyinstitute.org/the-interpreter/chips-made-taiwan-indispensable-ai-can-make-it-unstoppable>
- <sup>179</sup> The Looming Taiwan Chip Disaster That Silicon Valley Has Long Ignored (Feb 24, 2026) <https://www.nytimes.com/2026/02/24/technology/taiwan-china-chips-silicon-valley-tsmc.html>
- <sup>180</sup> If China Attacks Taiwan (Dec 2025) <https://www.gmfus.org/sites/default/files/2026-01/If%20China%20Attacks%20Taiwan.pdf>
- <sup>181</sup> Why helium is essential to the future of semiconductor manufacturing (Dec 5, 2025) <https://www.innovationnewsnetwork.com/why-helium-is-essential-to-the-future-of-semiconductor-manufacturing/64493/>
- <sup>182</sup> The Rise of AI to Fuel Growth in Helium Demand (Jul 25, 2025) <https://www.idtechex.com/en/research-article/the-rise-of-ai-to-fuel-growth-in-helium-demand/33588>
- <sup>183</sup> How the Iran war and rising energy prices are threatening semiconductor demand (Mar 10, 2026) <https://www.cnbc.com/2026/03/10/iran-war-semiconductor-memory-chip-impact.html>
- <sup>184</sup> The Complicated Stakes of the AI Race Between the U.S. and China (Feb 18, 2026) <https://time.com/7379419/ai-race-us-china/>

- 
- <sup>185</sup> China, the United States, and the AI Race (Oct 10, 2025) <https://www.cfr.org/articles/china-united-states-and-ai-race>
- <sup>186</sup> Canada 'monitoring' effect of Biden's new bans on China tech investments (Aug 10, 2023) <https://thelogic.co/news/canada-monitoring-effect-of-bidens-new-bans-on-china-tech-investments/>
- <sup>187</sup> Data centers become military targets as Iran war rages on (Mar 6, 2026) <https://www.cnbc.com/2026/03/06/iran-war-data-centers.html>
- <sup>188</sup> Banking, payments services disrupted after Amazon UAE data centers hit in drone strikes (Mar 3, 2026) <https://www.cnbc.com/2026/03/03/iran-war-uae-drone-strikes-aws-data-centers.html>
- <sup>189</sup> How Big Tech data centers become a military target during the war in Iran (Mar 6, 2026) <https://www.businessinsider.com/data-centers-iran-strikes-uae-bahrain-tech-military-target-war-2026-3>
- <sup>190</sup> The Scale of Russian Sabotage Operations Against Europe's Critical Infrastructure (Aug 19, 2025) <https://www.iiss.org/research-paper/2025/08/the-scale-of-russian--sabotage-operations--against-europes-critical--infrastructure>
- <sup>191</sup> Building Citizen Resilience: Preparing Canadians for an Age of Grey-Zone Conflict (Oct 2025) <https://www.queensu.ca/cidp/publications/research-reports/building-citizen-resilience-preparing-canadians-age-grey-zone>
- <sup>192</sup> Research on AI and the labor market is still in the first inning (Mar 10, 2026) <https://www.piie.com/blogs/realtime-economics/2026/research-ai-and-labor-market-still-first-inning>
- <sup>193</sup> Evaluating the Impact of AI on the Labor Market: Current State of Affairs (Oct 1, 2025) <https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs>
- <sup>194</sup> Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence (Nov 13, 2025) <https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>
- <sup>195</sup> New data show no AI jobs apocalypse—for now (Oct 1, 2025) <https://www.brookings.edu/articles/new-data-show-no-ai-jobs-apocalypse-for-now/>
- <sup>196</sup> Evaluating the Impact of AI on the Labor Market: Current State of Affairs (Oct 1, 2025) <https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs>
- <sup>197</sup> Labor market impacts of AI: A new measure and early evidence (Mar 5, 2026) <https://www.anthropic.com/research/labor-market-impacts>
- <sup>198</sup> The 2028 Global Intelligence Crisis (Feb 22, 2026) <https://www.citriniresearch.com/p/2028gic>
- <sup>199</sup> The Viral Citrini Substack Post That Has Sparked New AI Worries on Wall Street (Feb 23, 2026) <https://www.wsj.com/livecoverage/stock-market-today-dow-sp-500-nasdaq-tariffs-02-23-2026/card/the-citrini-substack-selloff-70cWx0scioiLradYuTRa>
- <sup>200</sup> Citadel Securities Rebuts Citrini 'Intelligence Crisis' Scenario (Feb 24, 2026) <https://www.bloomberg.com/news/articles/2026-02-24/citadel-securities-rebuts-citrini-intelligence-crisis-scenario>
- <sup>201</sup> People are worried that AI will take everyone's jobs. We've been here before (Jan 24, 2024) <https://www.technologyreview.com/2024/01/27/1087041/technological-unemployment-elon-musk-jobs-ai/>
- <sup>202</sup> Research on AI and the labor market is still in the first inning (Mar 10, 2026) <https://www.piie.com/blogs/realtime-economics/2026/research-ai-and-labor-market-still-first-inning>
- <sup>203</sup> Artificial intelligence, the economy and central banking (Sep 20, 2024) <https://www.bankofcanada.ca/2024/09/artificial-intelligence-the-economy-and-central-banking/>
- <sup>204</sup> The 2028 Global Intelligence Crisis (Feb 22, 2026) <https://www.citriniresearch.com/p/2028gic>

- 
- <sup>205</sup> Labor market impacts of AI: A new measure and early evidence (Mar 5, 2026) <https://www.anthropic.com/research/labor-market-impacts>
- <sup>206</sup> Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence (Nov 13, 2025) <https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>
- <sup>207</sup> Young workers' employment drops in occupations with high AI exposure (Jan 6, 2026) <https://www.dallasfed.org/research/economics/2026/0106>
- <sup>208</sup> AI disruption points to higher taxes on big tech (Nov 5, 2025) <https://cascadeinstitute.org/ai-disruption-points-to-higher-taxes-on-big-tech/>
- <sup>209</sup> The Big Money in Today's Economy Is Going to Capital, Not Labor (Feb 9, 2026) <https://www.wsj.com/economy/jobs/capital-labor-wealth-economy-2fc6c2f>
- <sup>210</sup> AI and the distribution of income between capital and labour (Mar 3, 2026) <https://cepr.org/voxeu/columns/ai-and-distribution-income-between-capital-and-labour>
- <sup>211</sup> AI disruption points to higher taxes on big tech (Nov 5, 2026) <https://cascadeinstitute.org/ai-disruption-points-to-higher-taxes-on-big-tech/>
- <sup>212</sup> Misinformation and disinformation <https://www.apa.org/topics/journalism-facts/misinformation-disinformation>
- <sup>213</sup> Disinformation (2026) <https://intelligence.weforum.org/topics/a1G680000008j2FEAQ>
- <sup>214</sup> AI-Driven Misinformation Across Sectors: Addressing a Cross-Societal Challenge (Nov 11, 2025) <https://sciencepolicy.ca/event/ai-driven-misinformation-across-sectors-addressing-a-cross-societal-challenge/>
- <sup>215</sup> Disinformation (2026) <https://intelligence.weforum.org/topics/a1G680000008j2FEAQ>
- <sup>216</sup> The Evolution of Disinformation - A Deepfake Future (Oct 2023) <https://www.canada.ca/en/security-intelligence-service/corporate/publications/the-evolution-of-disinformation-a-deepfake-future/the-evolution-of-disinformation-a-deepfake-future.html#toc2>
- <sup>217</sup> National Cyber Threat Assessment 2025-2026 <https://www.cyber.gc.ca/en/guidance/national-cyber-threat-assessment-2025-2026> (Oct 30, 2024)
- <sup>218</sup> What's Hiding Under the Kilt? Iranian Trolls for Scottish Independence (Sep 16, 2024) <https://www.clemson.edu/centers-institutes/watt/hub/images/hiding-under-the-kilt1.pdf>
- <sup>219</sup> Scottish independence accounts go dark after Iran internet blackout (Jan 12, 2026) <https://www.telegraph.co.uk/business/2026/01/12/scottish-independence-accounts-dark-iran-internet-blackout/>
- <sup>220</sup> Unmasking the bots: Researcher warns of threat to democratic processes (Jan 14, 2025) <https://news.mcmaster.ca/unmasking-the-bots-researcher-warns-of-threat-to-democratic-processes/>
- <sup>221</sup> Risk 3: Weapons of Mass Disruption (Jan 3, 2023) <https://www.eurasiagroup.net/live-post/top-risks-2023-3-Weapons-of-mass-disruption>
- <sup>222</sup> An A.I.-Generated Spoof Rattles the Markets (May, 23, 2023) <https://www.nytimes.com/2023/05/23/business/ai-picture-stock-market.html>
- <sup>223</sup> The \$26 Billion Threat: How AI Disinformation Is Reshaping Global Risk in 2026 (Feb 19, 2026) <https://blog.marketresearch.com/the-26-billion-threat-how-ai-disinformation-is-reshaping-global-risk-in-2026>
- <sup>224</sup> Deepfakes and the AI Arms Race in Bank Cybersecurity (Apr 17, 2025) <https://www.federalreserve.gov/newsevents/speech/barr20250417a.htm>

- 
- <sup>225</sup> Generative AI is expected to magnify the risk of deepfakes and other fraud in banking (May 29, 2024) <https://www.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html>
- <sup>226</sup> Deepfakes and Disinformation in the Finance Sector- Strategies to Prevent and Deter (May 2023) [https://www.sipa.columbia.edu/sites/default/files/2023-05/For%20Publication\\_BOFA\\_PollardCartier.pdf](https://www.sipa.columbia.edu/sites/default/files/2023-05/For%20Publication_BOFA_PollardCartier.pdf)
- <sup>227</sup> Deepfake content presents growing cybersecurity threat (Apr 25, 2025) <https://ucalgary.ca/news/deepfake-content-presents-growing-cybersecurity-threat>
- <sup>228</sup> Boards aren't ready for the AI age: What happens when your CEO gets deepfaked? (Mar 3, 2026) <https://fortune.com/2026/03/03/boards-arent-ready-for-the-ai-age-what-happens-when-your-ceo-gets-deepfaked/>
- <sup>229</sup> FIFAI II: AI Risks and Opportunities: Adopting an AGILE Framework in Canadian Financial Services (Mar 23, 2026) <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/fifai-ii-ai-risks-opportunities-adopting-agile-framework-canadian-financial-services>
- <sup>230</sup> The three different types of Artificial Intelligence – ANI, AGI and ASI <https://www.ediweekly.com/the-three-different-types-of-artificial-intelligence-ani-agi-and-asi/>
- <sup>231</sup> What is artificial superintelligence? <https://www.ibm.com/think/topics/artificial-superintelligence>
- <sup>232</sup> Canadian AI pioneers Yoshua Bengio and Geoffrey Hinton join global call to ban superintelligence development (Oct 22, 2025) <https://betakit.com/canadian-ai-pioneers-yoshua-bengio-and-geoffrey-hinton-join-global-call-to-ban-superintelligence-development/>
- <sup>233</sup> AI 2027 <https://ai-2027.com/summary>
- <sup>234</sup> The “AI 2027” Scenario: How realistic is it? (May 22, 2025) <https://garymarcus.substack.com/p/the-ai-2027-scenario-how-realistic>
- <sup>235</sup> Leading AI expert delays timeline for its possible destruction of humanity (Jan 6, 2026) <https://www.theguardian.com/technology/2026/jan/06/leading-ai-expert-delays-timeline-possible-destruction-humanity>
- <sup>236</sup> Statement on AI Risk <https://safe.ai/work/statement-on-ai-risk>
- <sup>237</sup> Made to order bioweapon? AI-designed toxins slip through safety checks used by companies selling genes (Oct 2, 2025) <https://www.science.org/content/article/made-order-bioweapon-ai-designed-toxins-slip-through-safety-checks-used-companies>
- <sup>238</sup> Security News This Week: A Creative Trick Makes ChatGPT Spit Out Bomb-Making Instructions (Sep 14, 2024) <https://www.wired.com/story/chatgpt-jailbreak-homemade-bomb-instructions/>
- <sup>239</sup> The risk of emergent misalignment in AI models: and how ChatGPT says we should manage this (Mar 5, 2025) <https://resilienceforward.com/the-risk-of-emergent-misalignment-in-ai-models-and-how-chatgpt-says-we-should-manage-this/>
- <sup>240</sup> Training large language models on narrow tasks can lead to broad misalignment (Jan 14, 2026) <https://www.nature.com/articles/s41586-025-09937-5.pdf>