

Banking on AI

Generative AI Adoption in Canada's Financial Sector

Viet Vu, Mahtab Laghaei | February 2026



Acknowledgements



The Dais is a public policy and leadership think tank at Toronto Metropolitan University, working at the intersection of technology, education and democracy to build shared prosperity and citizenship for Canada.

For more information, visit dais.ca
20 Dundas St. W, Suite 921, Toronto, ON M5G 2C2



How to Cite this Report

Vu, Viet, and Mahtab Laghaei. *Banking on AI: Generative AI Adoption in Canada's Financial Sector*. The Dais. 2026.

Graph data in this report can be found at:
https://github.com/thedaisTMU/FinanceSector_AISkills

<https://dais.ca>

© 2026, Toronto Metropolitan University 350 Victoria St,
Toronto, ON M5B 2K3



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. You are free to share, copy and redistribute this material provided you: give appropriate credit; do not use the material for commercial purposes; do not apply legal terms or technological measures that legally restrict others from doing anything the license permits; and if you remix, transform, or build upon the material, you must distribute your contributions under the same license, indicate if changes were made, and not suggest the licensor endorses you or your use.

Banking on AI: Generative AI Adoption in Canada's Financial Sector is part of the portfolio of work by the Future Skills Centre, which is funded by the [Government of Canada's Future Skills Program](#).

FSC is a forward-thinking centre for research and collaboration dedicated to preparing Canadians for employment success. As a pan-Canadian community, we are collaborating to rigorously identify, test, measure, and share innovative approaches to assessing and developing the skills Canadians need to thrive in the days and years ahead.

The opinions and interpretations in this publication are those of the author(s) and do not necessarily reflect those of the Government of Canada.

Design and Illustration

Zaynab Choudhry

Copy-Editing

Suzanne Bowness, CodeWord Communications

Translation

Marie-Pierre Lavoie

Contributor

Catherine Amburgey

The Dais proudly engages a diverse group of funders to support and catalyze our work, consistent with our **values**, and subject to a thorough internal review. As a non-partisan, public-interest institute, we only accept funds from organizations that support our mission and enable us to undertake work independently, with full editorial control. The names of all of our financial supporters are publicly and transparently displayed on all online and printed material for each project or initiative.



Authors



Viet Vu **Manager, Economic Research**

Viet Vu leads economics research at the Dais. Viet is interested in how governments and companies design policies and markets to drive human behaviour. He is also fascinated by how the world adapts to emerging new markets. Viet holds a Master of Science in Economics from the London School of Economics and Political Science and a Bachelor of Arts in Economics with honours from the University of British Columbia.



Mahtab Laghaei **Policy Analyst**

Mahtab Laghaei (she/her) is a policy analyst at the Dais, where she conducts research on digital technology, data and privacy, and the information ecosystem. Her work is informed by her interest in the relationship between geopolitics and tech policy. Mahtab holds a Bachelor of Arts from the University of British Columbia, and a Master of Global Affairs from the Munk School at the University of Toronto.

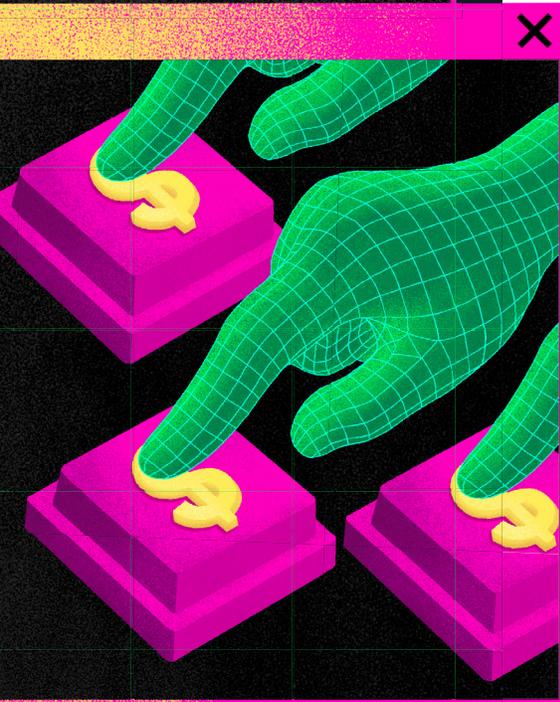


Table of Contents

5 Executive Summary

7 Introduction

8 Sector Background

9 Sectoral AI adoption considerations

10 Financial Sector Occupational and Task-Based Analysis

11 Occupational AI exposure-complementarity analysis

15 AI task exposure analysis

19 Recommendations: Financial Sector AI Use Cases

20 Use Case 1: Generative AI for front-line customer support

21 Use Case 2: Ancillary non-core finance tasks

22 Conclusions

23 References

1

Executive Summary

BOLD IDEA

Canada's financial sector is highly exposed to generative AI, but its most impactful use cases have little to do with crunching numbers. Institutions that deploy the technology responsibly can boost productivity and customer service, while those that don't risk far more than their bottom line.

Canada's financial sector is turning a new corner in the adoption of artificial intelligence (AI) technologies. While the sector has long used AI to support core analytical, modelling, and fraud-detection functions, the adoption of generative AI introduces new opportunities, risks, and implications for the workforce.

Through a sector-wide occupational exposure analysis with a novel task-level assessment, and using real usage data from Anthropic's Claude and Microsoft's Copilot, we propose recommendations for the sector on where generative AI can add value, where it introduces risk, and how to implement the technology to support productivity while safeguarding privacy and compliance.

Our analysis maps the impact of AI use (with a focus on generative AI) in the financial sector across two dimensions. First, we explore the exposure and complementarity of the workforce to understand the *likelihood of impact* to jobs, and the tasks within jobs, from AI technologies. Second, we identify the actual use of generative AI tools by financial sector workers and the types of tasks they are being used for.

The occupational exposure analysis finds:

- The vast majority of financial sector workers, numbering just over 800,000, are in occupations that are **highly exposed to AI technologies (98 per cent)**—far higher than the total Canadian workforce (56 per cent).
- Some of those workers, notably in senior management roles, are in **high AI complementarity occupations (24 per cent)**, meaning their work tasks are more likely to be assisted or augmented by AI.
- A much larger share of workers are in **occupations with a higher likelihood of AI task replacement (73 per cent)**. These are heavily concentrated in two occupation groups: business, finance and administration (e.g. human resources professionals, auditors), and sales and service (e.g. customer service representatives).

At the same time, using data from Anthropic's Claude and Microsoft's Copilot to map actual use trends of generative AI, the task analysis finds that:

- Generative AI is neither well used, nor particularly suited, to perform core numeracy tasks in the financial sector. This is largely due to the inherent inaccuracy in generative AI systems not found in other forms of AI technologies.
- Instead, tasks associated with front-line customer interaction, as well as ancillary business support tasks, are much better targets for AI deployment.

In deploying generative AI to the financial sector, we recommend that institutions use AI to:

Support front-line customer interactions

1. Balance cost-savings objectives with customer experience by deploying generative AI in ways that improve responsiveness and service capacity without fully replacing existing customer interaction workflows.
2. Design generative AI applications to ensure strong data and privacy assurance, including safeguards that prevent sensitive customer information from being shared with external systems.
3. Proactively assess liability considerations by establishing clear processes for remediation when generative AI provides incorrect or misleading information.
4. Plan for human redundancies and robust quality monitoring of generative AI outputs to maintain human oversight, incorporating regular audits to safeguard service quality.

Support non-core financial tasks

1. Align generative AI use cases with applicable professional and regulatory standards, particularly in occupations with formal accreditation or compliance obligations.
2. Be conscious of overreliance on generative AI tools by encouraging responsible use and preserving opportunities for workers to maintain and develop core professional skills.
3. Monitor generative AI use across the organization to ensure that the productivity gains from adoption outweigh implementation and governance costs.
4. Prioritize adoption in ancillary tasks that have lower consequences of error, focusing on areas where generative AI can safely support routine work.

2

Introduction

Artificial Intelligence (AI) encompasses a large swath of technologies, from industrial robots that form the backbone of Canada’s advanced manufacturing industry, to machine learning that powers the algorithms behind the platform apps on our smartphones, to chatbot tools that many people in Canada now directly interact with daily. As such, the impact that this class of technology has on each industry could look radically different, even when we restrict the analysis to one specific type of AI technology.

With a broad diversity of occupations represented in Canada’s major industries, the implications for specific jobs—and the tasks that make up those jobs—will be even more varied across the workforce. As a result, Canada’s aspirations to responsibly adopt AI must focus on trends and firm-level lessons about how AI tools can be applied across unique contexts, business needs, and workforce profiles of economic sectors across the country.

This report focuses on the AI exposure of jobs and skills, and task-based generative AI usage trends, in Canada’s finance and insurance sector (herein referred to as the *financial sector*)—one of the largest and most important for Canada in terms of employment and economic impact.

We expand upon an analytical approach established in [previous research](#) to assess the financial sector’s workforce exposure to AI¹ (i.e. the probability of interacting with AI systems in day-to-day work), and complementarity to AI (whether usage of AI is more likely to assist the worker with common tasks, or replace those tasks). We also introduce a new lens of analysis for understanding actual task-based use of generative AI tools, with data from Anthropic and Microsoft.

3

Sector Background

The financial sector collectively contributes 7.5 per cent² to Canada's GDP (the sixth largest industry by GDP contribution in Canada), and employs 4.6 per cent³ of Canada's workforce. The sector includes a few large institutional commercial banks, smaller regional credit unions, pension funds, asset investment firms, insurance providers, emerging technology-based financial (fintech) companies, and more.

Despite its size and economic contributions, the Canadian industry as a whole has seen relatively low levels of innovation, which prominent voices have attributed to a lack of competition in a sector dominated by a few of the largest financial institutions.⁴ In recent years industry and policy efforts have been focused on improving competition through "open banking" (a technological concept that allows for portability and interoperability when it comes to consumer data), with the 2025 federal budget also committing to creating a regulatory framework for cryptocurrency stablecoins and advancing new payment infrastructure as part of a financial sector innovation agenda.⁵

Still, there has been some progress with AI deployment. A Dais study, *Waiting for Takeoff*, found financial services to have the fifth-highest rate of firm-level AI adoption across Canada's industry sectors, with 10 per cent of companies reporting integration of some form of AI by the beginning of 2024.⁶ The concentration of larger firms in the financial sector is likely a contributor to the pattern of larger firms adopting AI and technology at a higher rate than smaller firms. Recent developments in generative AI are believed to offer compelling opportunities for efficiency and performance improvements across finance functions, through automation, data analysis, text summarization, and customer personalization.

Within the industry, our analysis found that many institutions focused on generative AI's potential impact in marketing and sales of financial products through hyper-personalization of marketing materials.⁷ For example, these can include chatbots (or agents) that customers can interact with to arrive at recommendations for products (such as credit cards) that directly respond to each customer's

unique needs. Agentic assistance was also highlighted as a transformative tool to assist with tasks such as customer onboarding, either directly through customer's interactions with chatbots, or internally when deployed to help staff expedite the onboarding process.⁸ Another benefit to generative AI adoption was in its ability to "unify" financial data that has remained siloed, potentially aiding in strategic scenario modelling.⁹

With compelling use cases ranging from marketing and sales to anti-money laundering efforts, there is significant potential for workforce disruption. However, discussions within the sector discount the likelihood of job displacement and disruption, and many predict that any changes will result in job reallocation or augmentation.^{10 11} Many in the sector therefore strongly believe that workforce training should be geared towards how individual workers can integrate AI tools into their workflow, instead of preparing and retraining workers for replacement of their jobs.

Sectoral AI adoption considerations

Industry analysis, consultancy reports, and related literature highlight factors to consider in the adoption of generative AI in the financial sector. They focus on risk governance and liability considerations, explainability and transparency of generative AI, and data quality. The industry continues to contend with challenges presented by AI to cybersecurity, such as breaches of data privacy, and security of physical assets.¹² Concerns about copyright and fair use are also emerging around generative AI applications sharing proprietary and licensed content, either with, or about, financial institutions.¹³ One report on risk governance argues that to ensure compliance and protect from liabilities, financial institutions should develop systems that track the source and use of data, and confirm if the data follows privacy regulation.¹⁴ Generative AI products are probabilistic and can be unpredictable. If teams lack the explainability or ability to determine how the model arrived at its output, institutions will struggle to validate or confirm the model's recommendations or actions, complicating compliance. Better

transparency can increase trust in generative AI both within teams and for customers.

Early adopters experienced both the benefits and the risks of generative AI, magnified by regulatory uncertainty. The risks could present an obstacle for firms navigating liability and security concerns, which call for stronger governance and control internally (such as at the board level). One outlet advises firms to begin the generative AI integration process in places where value is clear, but risk is low, as adequate accountability mechanisms are put in place.¹⁵ Efforts to organize and use quality data also work best when AI governance frameworks are in place, and when individual teams are attuned to generative AI's shortcomings and advantages.

Poor data quality has also surfaced as an impediment to harnessing generative AI's potential, making firms susceptible to liability issues and reputational damage as a result of inaccurate outputs. A blog post from a workflow automation company also called on firms to ensure accurate data is underlying the training of generative AI models.^{16 17} Others propose overcoming the issue of unstructured and multimodal data underpinning generative AI through Retrieval-Augmented Generation (RAG), a technique that enhances large language models (LLMs) by making generative AI pull information from trusted documents to ground or corroborate its response, making it more accurate and reliable than standalone generation.¹⁸ Still, a common concern is that general-purpose generative AI tools are not equipped to deliver accurate estimates through sensitive mathematical-based forecasting and risk assessments, and that firms should opt for more bespoke applications developed for specific financial sector tasks.¹⁹

4

Financial Sector Occupational and Task-Based Analysis

With this sector-specific context about AI adoption in mind, this section analyzes the AI exposure of workers in the financial sector and *likely impact* based on job tasks, and the task-based actual use of generative AI tools by financial workers. This analysis informs the identification of common financial-sector use-cases, and considerations for responsible AI deployment and use to support productivity enhancing outcomes for the sector.

Previous research has examined the degree to which occupations in the Canadian economy are exposed to generative AI, and how well AI use may complement work.^{20 21} **Exposure to AI reflects** the likelihood that a worker encounters an AI system when they work, where high-exposure occupations are those most likely to interact with AI systems in day-to-day work. **Complementarity to AI** measures the nature of a worker's interaction with AI. An occupation with high complementarity to AI means that a worker's interaction with AI is assistive, as opposed to an AI system replacing parts of their tasks.

Occupational AI exposure-complementarity analysis

Figure 1
Finance Occupations by Quadrant and Employment Size



Statistics Canada’s 2021 census data identifies 815,720 workers in the financial sector. The exposure-complementarity analysis of the workforce presented in Figures 1 and 2 reveal that the vast majority of financial sector workers are in occupations that are highly exposed to AI technologies (98 per cent). This

compares to 56 per cent in the Canadian workforce. In addition, nearly three-quarters of workers (73 per cent) are in the high exposure-low complementarity (HE-LC) quadrant, in occupations with a higher likelihood of AI task replacement.



Table 1

Overall financial sector employment by exposure-complementarity index

Exposure to AI	Complementarity to AI	Employment #	Employment share	Employment share (overall Canadian workforce)
High exposure	High complementarity	197,830	24.3%	27%
	Low complementarity	598,315	73.3%	29%
Low exposure	High complementarity	15,585	1.9%	14%
	Low complementarity	3,990	0.5%	29%

Table 2 compares the share of financial sector workers in highly AI-exposed occupations with the Canadian workforce across occupational groups. It shows a much higher concentration of financial-sector workers in highly AI exposed business, finance and administration jobs (57 versus 30 per cent), and a higher concentration in sales and service (28 versus 21 per cent). This reflects the greater potential impact of generative AI technologies on tasks associated with traditionally white-collar occupations. This delineation between generative AI and other forms of AI will become clearer in the next section, connecting real generative AI usage data to our analysis.

A breakdown of financial sector workers in the high exposure category (Table 3) shows that a much larger share of those working in sales occupations are represented in the high exposure-low complementarity quadrant, compared to high exposure-high complementarity quadrant. This shows that generative AI technology may have higher potential in replacing tasks associated with sales occupations, something we return to in the next section.

Table 2

Share of workforce in “highly exposed to AI” quadrant

Broad occupational group	Example occupations	Share in finance that is highly exposed	Share among the Canadian workforce in high-exposure occupations
Senior management	Senior managers in finance, health, trade, or construction	2%	2.3%
Business, finance, and administration	Human resources professionals, administrative assistants, auditors, and accountants	57%	30%
Natural and applied sciences	Information systems specialists, civil engineers, urban and land-use planners	9.8%	13.9%
Health	Registered nurses, orthopedic technologists, blood donor clinic assistants	Negligible	7.3%
Education, law, and social, community, and government services	Social workers, paralegals, policy researchers, lawyers	2.8%	18%
Art, culture, recreation, and sport	Translators, editors, graphic designers, and illustrators	0.5%	3.6%
Sales and service	Customer service managers and representatives, financial service representatives	28%	21%
Trades, transport, and equipment operators	Facility operation and maintenance managers, construction and transportation managers	0.1%	2.7%
Natural resources and agriculture	Managers in natural resources production and fishing, landscaping and grounds maintenance labourers	Negligible	0.1%
Manufacturing and utilities	Utilities managers, supervisors, petroleum, gas and chemical processing	Negligible	0.8%

Table 3

Occupational distribution among AI exposure quadrants for the financial sector

Broad occupational group	Example occupations	High exposure-High complementarity	High exposure-Low complementarity
Senior management	Senior managers in finance, health, trade, or construction	8.3%	0%
Business, finance, and administration	Human resources professionals, administrative assistants, auditors, and accountants	74.9%	51.1%
Natural and applied sciences	Information systems specialists, civil engineers, urban and land-use planners	5.5%	11.3%
Health	Registered nurses, orthopedic technologists, blood donor clinic assistants	0.4%	Negligible
Education, law, and social, community, and government services	Social workers, Paralegals, policy researchers, lawyers	5%	12.1%
Art, culture, recreation, and sport	Translators, editors, graphic designers, and illustrators	0.5%	0.4%
Sales and service	Customer service managers and representatives, financial service representatives	4.9%	35.1%
Trades, transport, and equipment operators	Facility operation and maintenance managers, construction and transportation managers	0.5%	Negligible
Natural resources and agriculture	Managers in natural resources production and fishing, landscaping and grounds maintenance labourers	Negligible	Negligible
Manufacturing and utilities	Utilities managers, supervisors, petroleum, gas, and chemical processing	Negligible	Negligible

AI task exposure analysis

The occupational analysis focuses on broad patterns of potential task-based AI exposure. However, it does not reveal how AI is *actually* being used at a task level. To supplement our occupational analysis, we introduce a new analytical method and conceptual framework that links automation exposure to actual generative AI usage.

By way of background, a factor that has often been overlooked when it comes to task exposure to automation is the consideration of quality differentials between tasks performed by humans as compared to machines, and how relevant those quality differentials are in the decision to automate a task. Some tasks may have large quality differentials between work performed by AI and humans, but may be used in instances where those quality differentials matter less (e.g. driver versus machine route navigation for taxis, where different routing only modestly changes the time to reach the destination). Other tasks, despite reflecting a lower quality differential between machine and human completion of a work task, can have severe consequences resulting from small quality differences (e.g. pharmacist versus machine prescribing and dispensing of medication, where getting exact amounts and dosage is essential).

To improve our ability to provide linkages between automation exposure and real AI usage, we introduce a measure on the suitability for tasks to be performed by AI. We distinguish between different automation suitability of each task by understanding the relative size of the consequence of errors, or providing lower quality goods in performing work tasks. We broadly call this a task's **error consequence** and ranks work tasks—using the Detailed Work Activities (DWA) construct in O*NET, a US-based occupational taxonomy—on tasks that have the highest consequence of making errors, to tasks that have the lowest consequence of making errors.

We use data on **actual AI usage**, taken from those that cover usage in Microsoft's Copilot²² and Anthropic's Claude by all users of the tool.²³ The usage is grouped at different levels, with Claude's usage being more granular at the DWA level, while Copilot's usage being at Intermediate Work Activities (IWA), which is one level higher than DWAs. We then calculate these concepts for tasks associated with **core occupations** in the financial sector, defined as occupations that relate to the core output within the sector. For instance, this includes workers such as accountants and financial analysts, but excludes occupations in human resources and administrative departments. This focus allows us to discuss AI usages for core business needs, as opposed to ancillary usages of the technology within the sector.

Taken together, this analysis allows us to better understand how well automation exposure maps to the suitability of automating technology in performing tasks, to actual AI usage, where such insights will allow better strategic decisions about where to deploy AI to be made. In a forthcoming paper, we will specifically focus on analyzing the validity of this analytical frame as a way to understand the variation in AI adoption across the task and occupational space.

Figure 2

Task Specific Automation Exposure and Error Tolerance, Finance

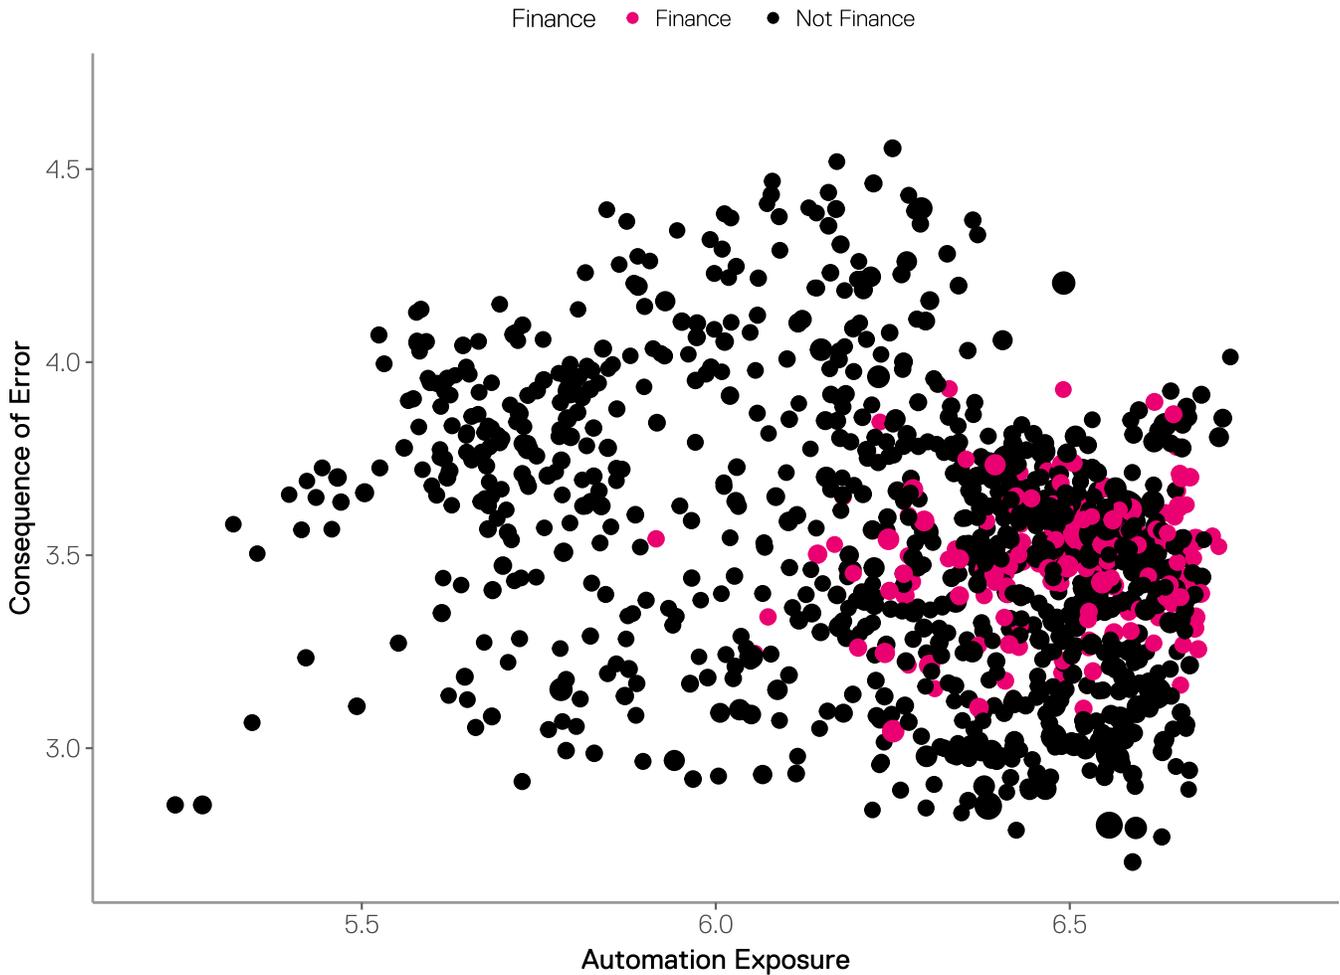


Figure 2 maps all tasks (at the DWA level) along our automation-exposure measure, and our error consequence (or suitability for each task to be done by AI) measure. This suggests tasks in the lower right quadrant (the highest level of automation exposure and the lowest level of consequence of error) are best suited for AI technologies. On the other hand, tasks in the upper left quadrant (lowest automation exposure and highest consequence of error) are tasks least suited for AI technologies to perform.

When tasks with the highest and lowest generative AI usage for core finance occupations are examined, there is a clear delineation of tasks. Real-life usage of generative AI tools tends to tasks related to non-financial information and assistance in general business operations, as opposed to work tasks requiring direct verification of records or numerical precisions. Similar trends are seen when Copilot usage is analyzed (with a focus on a task level one level higher than DWAs).



Table 4

**Top 5 and bottom 5 DWAs by real generative AI usages
(across all users, not just the financial sector)**

DWA description	
Top 5	Resolve computer software problems.
	Answer customer questions about goods or services.
	Develop promotional materials.
	Advise others on business or operational matters.
	Prepare business correspondence.
Bottom 5	Verify accuracy of records.
	Review license or permit applications.
	Examine the condition of property or products.
	Monitor flow of cash or other resources.
	Determine operational compliance with regulations or standards.

A more detailed analysis of representative tasks associated with core finance occupations shows two trends. The first trend reflects the previous exposure analysis, where work tasks in core finance occupations are highly exposed to AI. However, we

also observe that most finance tasks had middling levels of error consequence, limiting the potential of “low-risk” areas for adoption. Table 5 presents illustrative examples of detailed work areas (DWAs).



Table 5

Illustrative DWAs in core financial-sector occupations, and automation exposure

DWA	Occupational context²⁴	Automation exposure (percentile)	Error consequence (percentile)	Current Claude usage (percentile)
Calculate tax information	Tax preparers	69th	87th	18th
Analyze financial records to improve efficiency	Financial managers	87th	75th	87th
Answer customer questions about goods or services	Bank tellers	37th	9th	97th
Sell products of services	Securities, commodities, and financial services sales agents	25th	28th	48th
Review accuracy of sales or other transactions	Accountants and auditors	64th	73rd	13th
Monitor financial activities	Financial examiner	35th	83rd	24th

Our analysis linking generative AI usages to specific work tasks reveal key strengths and weaknesses of generative AI technology compared to other forms of AI. Traditionally, other forms of AI have tended to excel in numerical domains, owing to the statistical nature of these technologies. Generative AI technologies, on the other hand, focus on non-numerical input and generating non-numerical output. This focus in turn increases the probabilistic nature of numeracy-tasks and reduces the degree to which users can rely on the consistency or the accuracy of answers.

Instead, generative AI technologies tended to be used more intensively where variations in outputs are better tolerated (or in some cases desired). These facts inform our recommendations when it comes to deploying generative AI in the financial sector.

5

Recommendations: Financial Sector AI Use Cases

What becomes clear, when the analysis of occupational exposure and complementarity is combined with the task-level analysis, is that generative AI adoption in the financial sector should likely, for the time being, not be focused on core numeracy-based financial tasks, or tasks and use cases requiring high degrees of precision around issues like regulatory compliance requirements.

To be clear, existing non-generative AI systems (such as traditional machine learning models) are effective at numeracy-based financial tasks, and are commonly deployed in the financial sector in areas as wide ranging as automatic fraud detection, financial analysis and modelling, and tax calculations. In this study, however, the primary focus is generative AI

adoption. As a result, we urge the financial sector to first clearly delineate between generative AI systems, and non-generative systems, as they impact different work tasks, and will likely require different diffusion plans.

Two major use cases emerge from our analysis of task level generative AI usages in the financial sector: tasks involving front-line communications with customers, and within-job support of ancillary business operational tasks that are performed by core occupations in the financial sector. Neither of these use areas are without risk, and we assess both the opportunities and challenges of integrating generative AI.

Use Case 1: Generative AI for front-line customer interactions

For the first domain, the opportunity lies in generative AI's ability to be effective in summarizing and explaining existing information, where the risk of mistakes have a lower cost. This represents an opportunity to use these systems for front-line customer support interactions; for example, in using these systems to automatically or semi-automatically answer customer questions about financial products. These tasks are often associated with customer-support roles, as well as front-line bank teller roles. Our analysis suggests four considerations to ensure deployment is productivity-enhancing and mitigates risk:

Balance cost-savings with customer experience

Current generative AI systems are able to improve specific customer-experience metrics (such as cutting down wait time for customer support inquiries), while delivering customer interactions in a cost-effective manner. However, current implementation of the technology continues to be unreliable in delivering services, and may decrease customer satisfaction — for example, AI chatbots misinterpreting customer needs, or providing ineffective solutions.²⁵ In deploying these tools, the focus should be on optimizing cost savings and customer experience, and be used to expand service coverage (e.g. after-hours services or surge capacity during crisis times), as opposed to fully replacing existing customer interaction workflows.

Design for data and privacy assurance

The financial sector often deals with highly sensitive customer information. Any implementation of generative AI technology in the financial services therefore must either ensure custom AI systems are deployed to avoid storing and processing customer data with external vendors, or implement technical screens to proactively warn users before transmitting sensitive customer data to a third party. Participants in the sector can build upon solutions that guardrail customer data through an open banking framework.

Assess liability considerations

Despite best technical and design efforts to prevent generative AI tools from “hallucinating” (the term used to describe when AI generates false data), generative AI tools may still provide incorrect product information, or financial information to the customer. In these instances, investment should be made to understand and build structures to ensure reparative actions for the customer, and for the company to be equipped to deal with broader liability issues.

Ensure human fail-safes, and quality monitoring of outputs

In addition to investments in managing the interaction-by-interaction risks of the technology's deployment, companies should also implement regular macro-level quality monitoring of output to ensure productivity improvements are actually present without customer satisfaction being compromised. These audits could also provide signals for improvements in non-AI provided information (e.g. if the same question is being asked about a product for that information to be more prominently featured). Finally, some levels of human redundancies must be present to take over AI support when needed. Early evidence suggests generative AI systems benefit most from working alongside humans, where key benefits from using these systems come from the ability of these systems to solve rare problems, and reduce the skill-training gap for newer employees.²⁶

Use Case 2: Ancillary non-core finance tasks

For the second domain, integration of generative AI for non-customer-facing tasks should focus on ancillary business process tasks, as opposed to core finance numeracy tasks. These tasks can, for example, include summarizing financial news, preparing correspondence, and writing routine reports, or non-numeracy tasks that core occupations in the financial sector often perform. These can include, for example, financial advisors using a chatbot to conduct a routine daily scan of key financial news (where the financial advisor still interprets the news), or reviewing correspondences (as an additional layer to human review) to flag compliance risks. Such use cases can take advantage of the relative strength that generative AI systems have in taking in and summarizing information.

Any efforts to deploy generative AI in such a use case should additionally consider the following:

Align use cases with applicable professional and regulatory standards

Some occupations within the financial sector are under stringent professional and regulatory standards (e.g. financial advisors, accountants). Professional liability and standards should determine the best use cases for these technologies, including verifying outputs or not using it in areas that may compromise good standing and compliance with regulations and codes of conducts.

Be conscious of becoming over-reliant on generative AI tools

These technological tools remain emergent, and research on the long-term learning and skills impact of usage of these tools are still being understood. As users take up these tools, it is important to be cautious in becoming over-reliant on them to perform basic tasks.

Perform enterprise-wide cost-benefit analysis

As generative AI tools may not be applicable to core financial tasks, companies need to understand that productivity improvements and savings may be marginal. It is important to regularly evaluate whether the recorded benefits exceed the costs.

Prioritize adoption in tasks that have lower error consequences

A work mapping exercise²⁷ can help to identify priority ancillary tasks that have lower consequences or costs from errors, such as relying on generative AI to generate broad analysis, but to not recall specific data points.

6

Conclusions

Given the already extensive incorporation of AI systems in applications as wide-ranging as financial modelling and fraud detection, it may appear on the surface that the financial sector is well positioned to take advantage of advancements in generative AI. However, it is also true that generative AI technologies are not designed to handle high-precision, computationally-heavy, and numerical work that the industry often demands. In any event, the financial sector playbook for AI—and generative AI specifically—will look somewhat different than it does in other sectors.

Our research, combining insights from a financial industry scan of adoption dynamics with a workforce analysis of occupational exposure and task-based generative AI use, reveals that the vast majority of occupations in the financial sector are highly exposed to AI technologies in their day-to-day work, and that their usage of generative AI tends to less numerical or high-precision types of tasks that tend to have lower consequence for error. Until such a time that generative AI technologies improve to be reliable for numeracy tasks, generative AI adoption in the financial sector is best suited to improve customer interactions and assist workers in performing their core tasks. Instead of aiming for radical transformation, companies should set expectations for generative AI adoption to create incremental productivity improvements and modest cost savings.

References

- ¹ Vivian Li and Graham Dobbs, Right Brain, Left Brain, AI Brain: AI's Impact on Jobs and Skill Demand in Canada's Workforce, The Dais, 2025, <https://dais.ca/reports/right-brain-left-brain-ai-brain>.
- ² "Monthly Gross Domestic Product by Industry at Basic Prices in Chained (2017) Dollars - Seasonally Adjusted," Version 36-10-0434-01, Statistics Canada, 2025, <https://www150.statcan.gc.ca/n1/daily-quotidien/251031/t001a-eng.htm>.
- ³ "Employment by Industry, Annual," Version 14-10-0202-01, Statistics Canada, 2025, <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1410020201>.
- ⁴ Jordan Gowling, "Bank of Canada's Rogers Says Competition in Canada's Financial Sector Will Boost Productivity," Financial Post, October 9, 2025, <https://financialpost.com/news/bank-of-canada-rogers-competition-in-canadas-financial-sector-boost-productivity>.
- ⁵ Canada Strong Budget 2025, Department of Finance Canada - Government of Canada, November 2025: 118–21, <http://www.canada.ca/Budget>.
- ⁶ Viet Vu, Vivian Li, Angus Lockhart, Graham Dobbs and Christelle Tessono, Waiting for Takeoff: The Short-term Impact of AI Adoption on Firm Productivity, The Dais, 2024, <https://dais.ca/reports/waiting-for-takeoff>.
- ⁷ Andy Lees, "Harnessing Generative AI for Competitive Edge in Financial Services," Deloitte Blog, October 30, 2024, <https://www.deloitte.com/global/en/alliances/google/blogs/generative-ai-in-financial-services.html>.
- ⁸ Ibid.
- ⁹ Sydney Scott, 'How Generative AI Is Reinventing Scenario Planning,' Workday Blog, September 4, 2025, <https://blog.workday.com/en-ca/how-generative-ai-is-reinventing-scenario-planning.html>.
- ¹⁰ Caroline Hroncich, Why AI Isn't Coming for Your Banking Job, The Financial Brand, May 7, 2024, <https://thefinancialbrand.com/news/artificial-intelligence-banking/when-will-ai-come-for-banking-jobs-177763>.
- ¹¹ Claire Williams, Will the Uncertainty Continue for Financial Institutions? American Banker, December 13, 2023, <https://www.americanbanker.com/research-report/will-the-uncertainty-continue-for-financial-institutions>.
- ¹² Financial Industry Forum on Artificial Intelligence II: A Collaborative Approach to AI Threats, Opportunities, and Best Practices, Workshop 1 - Security and Cybersecurity, Office of the Superintendent of Financial Institutions (OSFI) - Government of Canada, July 2, 2025, <https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/financial-industry-forum-artificial-intelligence-ii-collaborative-approach-ai-threats-opportunities>.
- ¹³ Mary Cormack, Harnessing Gen AI in Finance: Balancing Innovation with Compliance, Copyright Licensing Agency, 2025, <https://cla.co.uk/harnessing-gen-ai-in-finance>.
- ¹⁴ Amit Garg, David Schoeman, Gabriel Morgan Asaftei, Kevin Buehler, and Liz Grennan, "How Financial Institutions Can Improve Their Governance of Gen AI," McKinsey & Company, 2025, <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/how-financial-institutions-can-improve-their-governance-of-gen-ai>
- ¹⁵ Casimir Rajnerowicz, "10 Key Use Cases of Generative AI in Finance," V7 Labs Blog, February 5, 2025, <https://www.v7labs.com/blog/generative-ai-in-finance>.
- ¹⁶ "GenAI and Data Quality: Paving the Path to AI Success," Moody's, June 13, 2025, <https://www.moody.com/web/en/us/insights/ai/genai-and-data-quality-paving-the-path-to-ai-success.html>.
- ¹⁷ David Schwimmer, "Why the Global Financial System Needs High-Quality Data It Can Trust," World Economic Forum, 2025, <https://www.weforum.org/stories/2025/01/high-quality-data-is-imperative-in-the-global-financial-system/>.
- ¹⁸ Garg et al., "How Financial Institutions."
- ¹⁹ Ashok Reddy, "Big Models, Bad Math: The GenAI Problem in Finance," Forbes Technology Council, May 5, 2025, <https://www.forbes.com/councils/forbestechcouncil/2025/05/05/big-models-bad-math-the-genai-problem-in-finance/>.
- ²⁰ Li and Dobbs, Right Brain, Left Brain, AI Brain.
- ²¹ Graham Dobbs, Vivian Li, Viet Vu and André Côté, Adoption Ready? The AI Exposure of Jobs and Skills in Canada's Public Sector Workforce, The Dais, 2025, <https://dais.ca/reports/adoption-ready-the-ai-exposure-of-jobs-and-skills-in-canadas-public-sector-workforce>.
- ²² Kiran Tomlinson, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri, "Working with AI: Measuring the Applicability of Generative AI to Occupations," arXiv, October 17, 2025, <https://arxiv.org/abs/2507.07935>.
- ²³ Kunal Handa, , Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller et al. "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations," arXiv, February 11, 2025, <https://arxiv.org/abs/2503.04761>.
- ²⁴ Occupations that are provided here are O*NET occupations (based on US taxonomies) as opposed to Canadian National Occupation Classifications (NOC). O*NET occupations are used as Detailed Work Activities (DWAs) is an O*NET concept, as opposed to NOCs.
- ²⁵ Flora An, "What Consumers Think About AI Customer Service Gone Wrong," Sobot, February 15, 2025, <https://www.sobot.io/article/ai-customer-service-gone-wrong/>.
- ²⁶ Erik Brynjolfsson, Danielle Li, Lindsey Raymond, "Generative AI at Work," The Quarterly Journal of Economics 140, no 2 (2025): 889–942, <https://doi.org/10.1093/qje/qjae044>.
- ²⁷ The work mapping exercise can be modelled after a user-story mapping exercise as seen here <https://ccndr.ca/wp-content/uploads/2025/04/Outcomes-of-prototypes-solution-pilot-EN.pdf>