


Waiting for Takeoff

The short-term impact of AI adoption on firm productivity

Viet Vu, Vivian Li, Angus Lockhart, Graham Dobbs and Christelle Tessono | December 2024



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

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>

Graph data in this report can be found at <https://github.com/thedaisTMU/waiting-for-takeoff>.

© 2024, 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.

This work is funded by:



Waiting for Takeoff: The Short-term Impact of AI Adoption on Firm Productivity is part of the portfolio of work by the Future Skills Centre, which is funded by the Government of Canada's Future Skills Program.

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: Mariana Rodrigues
Copy-Editing : Suzanne Bowness, CodeWord Communications
Contributor: Nina Rafeek Dow

We thank the participants of two roundtables we held in relation to this research who gave valuable feedback and perspectives. We in particular thank Magnet and its staff for co-organizing one of the roundtables.

We also thank the members of the research advisory committee who provided their expertise:

Ryan Kelly
Steven Tobin
Louise Levonian
Sarah Gagnon-Turcotte
Tony Bonen
Zachary Rose
Michelle Alexopoulos
Beryl Li

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.



Vivian Li - Senior Economist

Vivian Li is a Senior Economist at the Dais. Her areas of focus include examining labour markets within Canada's innovation economy and technology sector. She is interested in analyzing the relationship between economic factors and people, and the socioeconomic impacts of policies and events.

Vivian has completed a Bachelor of Arts in Economics from the University of Waterloo, and a Master of Arts in Economics from the University of Toronto.



Angus Lockhart - Senior Policy Analyst

Angus Lockhart is a senior policy analyst at the Dais. Angus researches the adoption of innovative technologies in both the public and private sectors and the role that policy can play in accelerating uptake.

Angus holds a Bachelor of Arts in Political Science from the University of British Columbia and a Master of Arts in Political Science from Simon Fraser University.



Graham Dobbs - Senior Economist

As a Senior Economist at the Dais, Graham Dobbs explores educational and technological innovations in the Canadian labour force and its impact on occupational distributions and transitions.

Graham has a degree in Economics and holds a Master of Arts in Economic Policy from McMaster University.

He was also a sessional lecturer at the undergraduate health sciences program at McMaster, teaching the economics of healthcare.



Christelle Tessono - Policy and Research Assistant

Christelle Tessono conducts research at the intersections of digital technology, human rights, and governance. This has led her to work on a variety of projects related to political advertising on social media platforms, gig work, facial recognition technology, privacy, and AI policy.

Christelle holds a Bachelor of Arts in political science from McGill University and is currently pursuing graduate studies at the University of Toronto's Faculty of Information.



Table of Contents

5	EXECUTIVE SUMMARY
7	INTRODUCTION
9	LITERATURE REVIEW
13	METHODOLOGY
13	Firm-level productivity estimation
14	Causal estimation (triple difference-in-difference estimation)
17	DATA
17	Survey of Digital Technology and Internet Use
18	Business Research Microdata
19	FINDINGS
19	Descriptive results
21	Causal modelling
23	CONCLUSION
24	APPENDIX A
26	APPENDIX B



1

Executive Summary

BOLD IDEA: Many private-sector and policy leaders are betting big on AI as the silver bullet for Canada’s productivity crisis, but AI is a long game, not a quick fix. Companies must integrate it thoughtfully into their operations to see real gains.

Artificial intelligence (AI) is the most discussed technology of recent years. Advocates promise that it will help overcome productivity challenges and radically transform the economy through increased wage gains and higher economic output, among other benefits. This conversation about the intersection of AI and productivity growth is particularly important in Canada today, amidst concern about a “**productivity crisis**.”

Productivity is a key ingredient in future economic growth and standard of living, as it offers the potential to increase output without increasing inputs—like worker hours, natural resources, and investment costs. Yet, in past waves of innovation, we have seen patterns where a technology achieves widespread adoption, without any evidence of it increasing productivity.

The late American economist Robert Solow famously remarked that “You can see the computer age everywhere but in the productivity statistics.” Will this time be different?

Understanding whether AI will follow the same trends as previous waves of innovation is essential. Gaining this understanding will inform economic policy, business investment decisions, workforce planning, and broader discussions about AI’s benefits and drawbacks.

In this study, we tackle the critical question of whether AI adoption leads to productivity improvement at the firm level.

Key Findings

Using high-quality data from Statistics Canada's Survey of Digital Technology and Internet Use and the Business Research Microdata, accessed through the Canadian Research Data Centre Network, we looked at firms that adopted AI between early 2020 and late 2021, to understand whether the adoption of AI technologies translated to short-term firm-level productivity gains. When we define "short-term" as one to two years post adoption, we find the following:

- **Evidence of productivity gains from AI use is mixed.** There is no conclusive evidence of a strong positive or negative relationship between AI adoption and short-term productivity improvement.
- **There was no significant relationship between the adoption of AI in this period and either Total Factor Productivity (TFP) levels or TFP growth** (efficiencies in output production which do not stem from added labour or capital inputs) in the short-term.
- **The set of firms that adopted AI were already more productive than their peers**, but the decision to adopt AI did not increase the rate at which their productivity grew.

While this is the first report in Canada to provide a look into the relationship between AI adoption and firm productivity, the overall rates of AI adoption in the Canadian economy remain low, and at an early stage.

As applications in AI become more widespread and are increasingly embedded across various operations, there could be an increased chance for potential efficiencies to translate into increased productivity.

Most notably, this research focuses on the impacts of AI adoption before the public launch of OpenAI's ChatGPT in late 2022, which sparked widespread interest in the latest generation of generative AI technologies. As generative AI tools like Large Language Models (LLMs) offer different capabilities from other types of AI, further research is needed to assess its impact on productivity growth in Canadian firms. Data from the next iteration of the Statistics Canada survey conducted in late 2023 through early 2024 will make this possible.

Many private-sector and policy leaders presume that business adoption of AI can be a silver bullet in addressing Canada's productivity growth challenge. Our findings call for caution in asserting that AI adoption at the firm level results in short-term productivity gains. We look forward to continuing to analyze and research how the deployment of these fast-changing technologies affects the course of the Canadian economy.





2

Introduction

Since the Industrial Revolution and the first generation of automating technologies powered by the steam engine, concerns about the implications of those technologies for the economy, and in particular, for workers, have been constant. Now, almost 300 years later, we are faced with the same questions, this time due to automation technologies powered by electricity and silicon-based semiconductor chips. This conversation about the intersection of artificial intelligence (AI) and productivity growth is particularly important in Canada today. Many voices are raising concern about a “productivity crisis” that threatens Canada’s future economic growth and standard of living, which is more dependent on the resource sector than our G7 peers.¹² Many have argued that AI can be an essential piece of the growth puzzle.^{3,4} Yet, to date, more research has focused on whether these technologies will replace workers, with few studies assessing whether the adoption of AI leads to improvements in productivity.⁵

There are many reasons to take the inquiry into estimating AI’s productivity impact seriously. Robert Solow, a Nobel laureate and leading American theorist on economic growth, famously remarked

about the widespread adoption of personal computers: “You can see the computer age everywhere but in the productivity statistics.” Termed the “Solow paradox,” many economists have sought to explain why productivity growth showed little relation to technological advancement.⁶ Will this be the case with the current wave of AI as well? With relatively low AI adoption rates in Canada and in other countries,⁷ and macroeconomic research projecting a modest contribution from AI to overall productivity growth,⁸ an objective examination of how AI adoption is affecting productivity in Canada is needed.

Robert Solow, a Nobel laureate and leading American theorist on economic growth, famously remarked about the widespread adoption of personal computers: “You can see the computer age everywhere but in the productivity statistics.”

The urgency is amplified by a new feature of the latest generation of AI technologies—a class of models known as Generative Pre-trained Transformers (GPTs). GPTs are a subset of Large Language Models (LLMs), or AI models that are trained on vast quantities of human-generated natural data such as text or images.

GPTs are showing signs that they are different from the previous waves of AI. Instead of replacing routine tasks (or tasks that can be described in a consistent, repeatable, step-by-step manner to achieve similar results), GPT-based AI seems to complement cognitive workers in non-routine tasks. Early experiments in introducing generative AI in workplaces also show that instead of hurting the lowest skilled workers in a given occupation, it reduces the performance gap between the bottom performer and the top performer in the same job task.^{9 10 11}

In this study, we tackle the critical question of whether AI adoption leads to productivity improvement at the firm level. We rely on unique, high-quality survey data on AI usage in Canadian businesses linked to their tax filings to causally estimate the change in productivity as a result of adopting AI. The survey data, which tracks firm adoption of AI between early 2020 and late 2021, is limited because it does not capture the more recent emergence of generative AI, with the launch of OpenAI's ChatGPT in late 2022. Still, it offers a unique opportunity to study the short-term productivity impacts of business AI adoption, and provides a potential baseline for longitudinal tracking of AI-related productivity trends as the pace of AI adoption accelerates in Canada.

Early experiments in introducing generative AI in workplaces also show that instead of hurting the lowest skilled workers in a given occupation, it reduces the performance gap between the bottom performer and the top performer in the same job task.



3

Literature Review

Technological innovation often surfaces a tension between short-term disruption and long-term benefits. So, a review of literature that examines both periods provides a holistic sense of AI's potential impact.

In the short run, anxieties centre upon the potential negative impacts on employment as new technologies become more common across workplaces and industries. At the same time, those who focus on the long run often note the potential benefits of technology for productivity improvement and overall prosperity, as technological innovation in the long run has often proved beneficial. For example, the “spinning jenny” yarn spinning machine that allowed a single operator to spin multiple threads simultaneously helped usher in the Industrial Revolution.¹² Chemical and biological technologies spurred revolutions in agricultural production.¹³

The acceleration in both the creation and diffusion of technological innovation makes short- and long-run impacts harder to predict. Therefore, scholars focus on whether the existing models used to understand technology's impact on the economy need to be updated.

This highlights the need to synthesize our understanding of the varying degrees in which technological impact is recognized. Our paper aims to bridge these two literatures—the long-term productivity impact and the short-term employment impact—by focusing on the short-run productivity impact of adopting AI.

What is AI?

Before we proceed, it is important to define the bounds of what we consider AI and clearly delineate where researchers have treated it as separate from or similar to previous automation technologies. The Organisation for Economic Co-operation and Development (OECD) defines AI as “a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.”¹⁴ The delineation between AI and other technologies is often ignored in the long-term economic growth literature, which concerns the general level of technological change and how that impacts long-term productivity.

For the purpose of this study, however, AI is defined using our primary data source, Statistics Canada's Survey of Digital Technology and Internet Use (SDTIU): "systems that display intelligent behaviour by analyzing their environment and taking actions - with some degree of autonomy - to achieve specific goals. AI-based systems can be purely software-based or embedded in a device."¹⁵ Notably, the survey, directed at technology decision-makers in Canadian businesses, includes a question about AI adoption.

The survey question explicitly asks about the use of specific AI-powered applications like machine learning, virtual agents, automatic speech recognition, face recognition systems and other image-analysis software, hardware with integrated AI, and technologies that automate workflows or assist in decision-making. Respondents can also report other forms of AI used in their firms. The definition used in the SDTIU ensures that businesses identified as AI adopters are using it to a meaningful degree.

The productivity paradox

Literature focused on the short-term distributional effects of technology adoption has characterized the current wave of AI technologies through its immediate impact on workers (through displacement/job losses or shifts in skills demand). However, short-term distributional effects are often not representative of a technology's impact in the long run.

Robert Solow's famous quote, referring to the wider adoption of personal computers in the economy, was first stated in response to the observance of a slowdown in productivity despite rapid technological progress in the 1980s.

Many theories to explain this seemingly counterintuitive idea were developed, including an uneven distribution of labour productivity gains, mismeasurement of firm output or employment growth due to adoption, and implementation lags.¹⁶ In addition, it was theorized that a reallocation effect was taking hold, whereby workers who are

AI is defined using our primary data source, Statistics Canada's Survey of Digital Technology and Internet Use (SDTIU): "systems that display intelligent behaviour by analyzing their environment and taking actions - with some degree of autonomy - to achieve specific goals. AI-based systems can be purely software-based or embedded in a device."

replaced by new technologies move from productive technology-adopting sectors to less productive non-adopting sectors. This movement then cancels out any expected productivity gains.¹⁷

Studies have shown that automation has the potential to reduce wage dispersion among groups of workers who perform similar tasks and reduce the gap between the most highly skilled and the least highly skilled workers within the same occupation. Automation tends to replace those tasks that command higher wages, which drives between-group wage inequality.¹⁸ The result of the changing premium of skills through the automation of tasks (and replacing of workers as a result) has a dampening effect on the wage growth that would otherwise come from increased productivity.

A profile of firm-level AI adopters in Canada

Our understanding of the potential for productivity gains from AI throughout the broader economy is currently limited. This is due to (a) low and uneven AI uptake among firms and (b) inconsistent access to data for researchers. Evidence from the OECD across 11 member countries (not including Canada) finds that AI adoption is heavily concentrated among larger and, to some extent, younger firms – and less in smaller or more established firms.¹⁹ The wider use of AI across large firms may be related to the more considerable endowments or capabilities to use intangibles and other complementary assets needed to fully leverage AI's potential. A large firm may already bring assets such as high levels of, or investment in, Information and Communications

Technology (ICT) skills and training, firm-level digital capabilities, and digital infrastructure—this provides a stronger foundation for AI use.

This pattern of uneven adoption holds true in the Canadian context as well. Previous research found that as of 2021, only 3.7 percent of Canadian firms had adopted AI in their business practice.²⁰ Similarly, adoption was uneven across businesses, with large firms adopting significantly more than smaller firms.²¹ A recent study focusing on French firms produced similar conclusions, noting that larger firms were able to yield greater productivity gains because they had access to significantly more resources to leverage the potential of the AI systems adopted.²² Given the overlap between adoption and firm size, the productivity gains by large firms do not seem to reflect the use of AI alone, and reflect other business characteristics.

Globally, AI adoption is more common among younger firms, and in ICT and professional services industries.²³ Start-ups often bring more radical innovations to the market, especially when new technological paradigms emerge. The concentration of AI adopters within certain industries and firm types suggests AI's full potential as a general-purpose technology has yet to fully materialize across the wider economy.

Furthermore, as raised by Brynjolfsson et al. (2023),²⁴ finding evidence of the productivity effects of AI-based technologies is challenging to measure at the macroeconomic scale, due to the different ways firms choose to adopt AI. As a result, evidence of significant productivity gains is yet to be demonstrated by scholarship and the little evidence available cannot be reasonably applied across the economy.

Evidence from the OECD across 11 member countries (not including Canada) finds that AI adoption is heavily concentrated among larger and, to some extent, younger firms.

Given the overlap between adoption and firm size, the productivity gains by large firms do not seem to reflect the use of AI alone, and reflect other business characteristics.

Measuring productivity and AI adoption

While there are a number of different ways to measure productivity, we focus on the concept of Total Factor Productivity (TFP) for this study, which measures the efficiencies in producing output that does not stem from the addition of an extra unit of labour or capital.²⁵ TFP doesn't just measure labour productivity but includes productivity metrics for non-labour factors as well (such as machinery). In other words, productivity is the ratio of the value of an output, and the value of all the inputs that were involved in making the output.²⁶

Scholars have had a difficult time measuring AI productivity due to the lack of accessible data about firms' AI adoption.²⁷ Information about the specific types of AI-based technologies, and in addition to details about the contexts in which they are used by firms, are underreported. Two different methods are used to define AI use and firm output.

In one approach, researchers have used patent applications and data on specific AI-based technology adoption such as robots and data-driven managerial decision-making.²⁸ Alderucci et al. (2020) analyzed patent grants in the United States to assess the prevalence of AI-related inventions by companies. The authors linked that data with U.S. Census Bureau microdata to identify patterns in firm labour demand and growth. They found that firms with AI-related innovations experienced a 25 percent faster employment growth and a 40 percent faster revenue growth than similar firms that had not developed and adopted AI tools.²⁹

As a second approach, researchers have used firm revenue and operating income to assess productivity gains from AI adoption. For instance, Kim, Park and Kim (2022) identified 105 firms in the United States

that adopted AI for automating and augmenting tasks between 2000 and 2015.³⁰ The study found that over time, there was no evidence of cost reduction and increased profits from AI adoption, which signifies no real impact on productivity outcomes as a result. However, another study using operating income to measure output looked at the impact of AI penetration on TFP among A-share-listed³¹ Chinese manufacturing firms between 2010 and 2021. The study found that a one percent increase in the density of AI use led to a 14.2 percent increase in TFP.³²

A recent study used panel data from the European Commission's 2018 Community Innovation Survey (CIS) to assess AI-related productivity gains in Germany.³³ The study found that firms using AI had increases in sales of their products, indicating a short-term increase in productivity linked to AI adoption.

Sector-specific impacts of AI adoption

Scholars have also focused on productivity gains across different sectors of the economy in different jurisdictions. In health care, there have been studies that have found that AI enhances employee productivity in the delivery of services to patients.³⁴ In agriculture, AI has been adopted to support farmers in assessing yield predictions and allocate resources in a cost-effective manner. There are several studies that identify its potential for enhanced productivity as a result.³⁵ Another study by Kanazawa et al. (2022) found that an AI-based navigator increased taxi drivers' productivity by making it faster to find customers. However, the increase in productivity was concentrated in low-skilled drivers, narrowing the productivity gap between them and high-skilled drivers by an estimated 14 percent.³⁶

Firm-level productivity is often driven by efficiencies in worker-level tasks. In particular, generative AI and LLMs are highly applicable in supporting writing-based and coding-based tasks. For instance, there are studies that analyze the productivity gains of using assistive chatbots like OpenAI's ChatGPT in the workplace. One study, which assigned 453 college-

educated professionals mid-level writing tasks, showed that ChatGPT raised average productivity, with time taken to complete a task decreasing and writing output quality increasing.³⁷ In addition, Peng et al. (2023) conducted a study which found that software engineers who were provided GitHub CoPilot were able to complete a task twice as fast compared to those who did not have access to it.³⁸ However, other studies show that using generative AI tools in some tasks (such as coding), significantly increased error and bug rates.³⁹

In summary, research on productivity gains resulting from AI use is limited by data availability challenges and differing approaches to measuring AI adoption and firm productivity. The existing literature produces mixed results. There is no conclusive evidence, at the economy-wide or sectoral-level, of a strong positive relationship between the technology's adoption and productivity improvement. The research findings and limitations also highlight the need to approach productivity assessments with nuance and particular attention to the skills and working conditions of workers using these tools.

In summary, research on productivity gains resulting from AI use is limited by data availability challenges and differing approaches to measuring AI adoption and firm productivity. The existing literature produces mixed results. There is no conclusive evidence, at the economy-wide or sectoral-level, of a strong positive relationship between the technology's adoption and productivity improvement.



4

Methodology

Firm-level productivity estimation

A key outcome measure in our study is a company's level of productivity growth. In this work, we conceptually refer to Total Factor Productivity (TFP) whenever we use the word "productivity". TFP is defined as the efficiencies in output production which do not stem from an added unit of labour or capital input. The preferred approach, therefore, first popularized by Robert Solow, is the idea of a "Solow residual"—that TFP is whatever is "left over" after we account for the value of all the input from the value of the output.⁴⁰

This is a two-step procedure. First, industry-specific estimates of the ratios of capital and labour used are produced. The resulting parameters for the industry-specific production are then used in the second stage, alongside firm-specific attributes (such as payroll, revenue, and capital valuation). This allows us to arrive at firm-level productivity estimates. Appendix A explains the specific process used in more detail, and why this approach, called the Wooldridge-Levinsohn-Petrin process or WLP, is the preferred approach in economics.

Causal estimation (triple difference-in-difference estimation)

Using the firm-level productivity growth measure from the previous section, we now describe our empirical strategy to obtain the causal impact of AI adoption on productivity. In particular, we use a variation of a popular causal estimation in economics—the difference-in-difference (diff-in-diff) estimator.

Conceptually, diff-in-diff allows researchers to focus on the difference in outcome between two groups, where only one has been exposed to a “causal event”. A quasi-experimental method allows a researcher to argue for a causal relationship between an exposure (a policy intervention or a particular event), and an outcome (desired policy goal).

It does so by focusing on situations where the exposure to a policy intervention is targeted to a specific sub-sample. Its causal claim relies on the “parallel-trend assumption” where companies (in this case) that adopted AI and companies that did not adopt AI were on a similar trajectory. If AI adopters had not adopted AI, it is assumed they would have remained on the same trajectory. Its key advantage is that it allows for a limited causal interpretation where random exposure to an intervention (in this case, AI adoption) is absent. Technical details on the difference-in-difference model used in this report can be found in Appendix A.

We identified two distinct subgroups of Canadian firms: adopters of AI between 2020 and 2021, and non-adopters who remained in the sample.

Variable selection

Variable of interest

Our variable of interest identifies firms adopting AI technologies over the three-year panel sample. Simply put, we identified two distinct subgroups of Canadian firms: adopters of AI between 2020 and 2021, and non-adopters who remained in the sample. Those who had adopted AI before 2020 are excluded, as we cannot identify when they adopted the technology. The matrix below illustrates the inclusion criteria and groups of interest in our panel sample:

Table 1: Summary of AI adoption treatment conditions

AI adoption panel sample restrictions	Did not adopt AI by 2021	Adopted AI by 2021
Did not adopt AI by 2019	Control group	Treatment group
Adopted AI by 2019	Excluded from sample	Excluded from sample

The panel restrictions allow us to understand the unique effect of AI adoption among Canadian firms' TFP relative to firms that did not adopt over the three-year analysis period.

In our alternative modelling specifications, we create an interaction variable as our treatment variable using the firm tax reporting year and binary indicator identifying AI adoption in 2021. This provides an understanding of the independent effect of AI adoption on both the level and change of TFP.

Controls

We derive the variable for industry by creating a binary flag for firms in service industries by the first digit of the North American Industry Classification System codes (NAICS). Our previous research has found that service industries hire and employ workers with skills more likely to be impacted by AI and have a wider variety of value propositions for its application.

We cluster our standard error estimates on firm size to account for unobservable characteristics among firms adopting AI. This allows us to account for time-varying unobservable characteristics not captured in the control and treatment variables based on the firm's size. Our previous works find AI adoption among Canadian firms is skewed heavily toward firms with high employee headcounts.

We use the business size variable from the SDTIU,⁴¹ a categorical variable indicating if a business employs less than 20, 20 to 99, or 100 or more employees.⁴²

Robustness

To ensure the robustness of the estimation of the causal effect, we also leverage a fixed-effect ordinary least squares (OLS) regression model, instrumental variable, and quantile regressions at the 25th, 50th, and 75th percentiles.

This estimates the coefficient on TFP for the in-treatment group of AI adopters in the year 2021 as an interaction variable between the post-treatment period (2021) and firms in the treatment group (AI adopters in 2021 but not 2019). Similar to the difference-in-difference model, a binary flag for the firm's goods or services industry was added as a control, with clustered standard errors on firm size. In the case of quantile regressions, we add firm size as one of the control variables. More information about the form of this regression can be found in Appendix A.

The period of change we consider we study is those between 2019 and 2021, notably between a period of significant economic disruption brought about by the COVID-19 pandemic. We discuss here how our findings should be characterized in the context of this event.

We note that difference-in-difference methodology aims to control for common shocks and trends that affect all businesses. As a result, we are confident in our causal claim itself not being contaminated by the COVID-19 pandemic. We discuss potential concerns about interpreting our results outside of the COVID-19 context.

However, the COVID-19 pandemic may have created a unique environment for firms adopting AI. Given the economic disruptions caused by the pandemic, the motivation for firms to adopt AI may have been different than it would be under more normal circumstances. For instance, firms may have turned to automating technologies to help maintain productivity despite requirements for social distancing and increased health and safety protocols. Similarly, it is possible that given the slack in the labour market produced by the pandemic, some firms that would have otherwise implemented AI to address hiring difficulties decided against it.

If firms adopted AI during this period to address new challenges, or failed to adopt it for reasons that would normally have motivated them to do so, it is possible that in a different context adoption of AI may have yielded different results. Replicating this work for the period after the pandemic will be important to verify whether these results continue to be true or if they change outside of the COVID-19 context.



5

Data

Survey of Digital Technology and Internet Use

The main source of data we will be using for our research will be the Survey of Digital Technology and Internet Use (SDTIU). The SDTIU is a mandatory survey conducted by Statistics Canada, targeted at businesses with the goal to understand the uptake of different technologies in Canadian businesses. The modern iteration of the survey was first conducted in 2019 (into 2020), with a follow-up wave in 2021 (into 2022), and a third wave that concluded data collection in March of 2024 (not yet published). Businesses are legally required to answer the survey, and around 15,000 firms answer each survey wave.

The survey is answered by the technology decision-maker at the businesses and covers technology use cases that are “core to business processes” (e.g. an employee using Google to search for information, and thus interacting with AI that ranks search results, would not be covered). The answer to this question, particularly as it relates to AI, forms our primary measure of AI adoption in firms.

Another feature we exploit to perform our causal analysis is the non-trivial number of firms who answered both the 2019 and 2021 SDTIU—totalling more than 2,000 firms. This allows us to examine productivity growth in companies that adopted AI between 2019 and 2021 (that is, they had not adopted AI in 2019 but have adopted it by 2021).

The SDTIU by itself, however, does not provide any financial and employment data needed to estimate firm-level productivity. To obtain this information, we rely on the fact that the SDTIU exists on Statistics Canada’s Business-Linkable File Environment (B-LFE). This data ecosystem links business survey data to administrative data sources.

Using the B-LFE, we connect unique firm-level identifiers to key business attributes in the Statistics Canada Business Register (such as their industry and location) and financial information from the General Index of Financial Information (data on revenue, employment, and capital from business tax filings). This allows us to estimate firm-level productivity growth for businesses in the SDTIU dataset.

Business Research Microdata

Business Research Microdata (BRM), previously known as the National Accounts Longitudinal Microdata File (NALMF), is a core research administrative data source that Statistics Canada maintains. It contains longitudinal data on all registered businesses in Canada, with data from administrative sources, such as the Business Register, tax data, and the Export Register.

We primarily use the BRM to perform the first of the two-stage estimation process for firm-level productivity. As we discussed in the methods section, the method we use to estimate firm-level productivity requires estimating a unique model for each of the 324 detailed industry codes (called the North American Industry Classification System or NAICS), to then apply each of these 324 models to the individual company attributes (such as revenue and employment). The sample size we have access to through the SDTIU simply does not allow us to rigorously perform this step. As a result, we will use data on all registered companies in Canada to generate these industry-specific estimation models (from 2000 to 2020).





6

Findings

Descriptive results

Before we present the modelling of the effect of adopting AI on productivity, we start by establishing baseline characteristics of the two sample groups by comparing key descriptive statistics.

First, to understand historically how input decisions have varied across industries, we calculate industry-specific allocations to labour and capital using a wider sample from the complete data set from the Business Research Microdata from 2001 to 2022.

When we apply these estimates to the treatment group (firms that adopted AI during the 2020 to 2021 period) to calculate firm-level productivity (a measure of revenue produced by a set of inputs), we see that the AI-adopting firms are overall more productive than firms not using AI both before and after the treatment period.

Among AI adopters, productivity levels have been 10 to 35 percent higher than non-adopters in the years observed in our study, as seen in Figure 1. This itself is a notable finding and suggests that AI adoption is correlated with greater firm productivity.

But does AI adoption cause greater firm productivity? Is AI adoption associated with increased firm Total Factor Productivity after adoption? Looking at annual growth in TFP, which is to say the difference in a firm's TFP between one period to the next, we see in Figure 2 that TFP at AI adopter firms is not consistently growing any faster than other firms to a significant degree (with the exception of 2022).⁴³

Figure 1: Total factor productivity by year and AI adoption

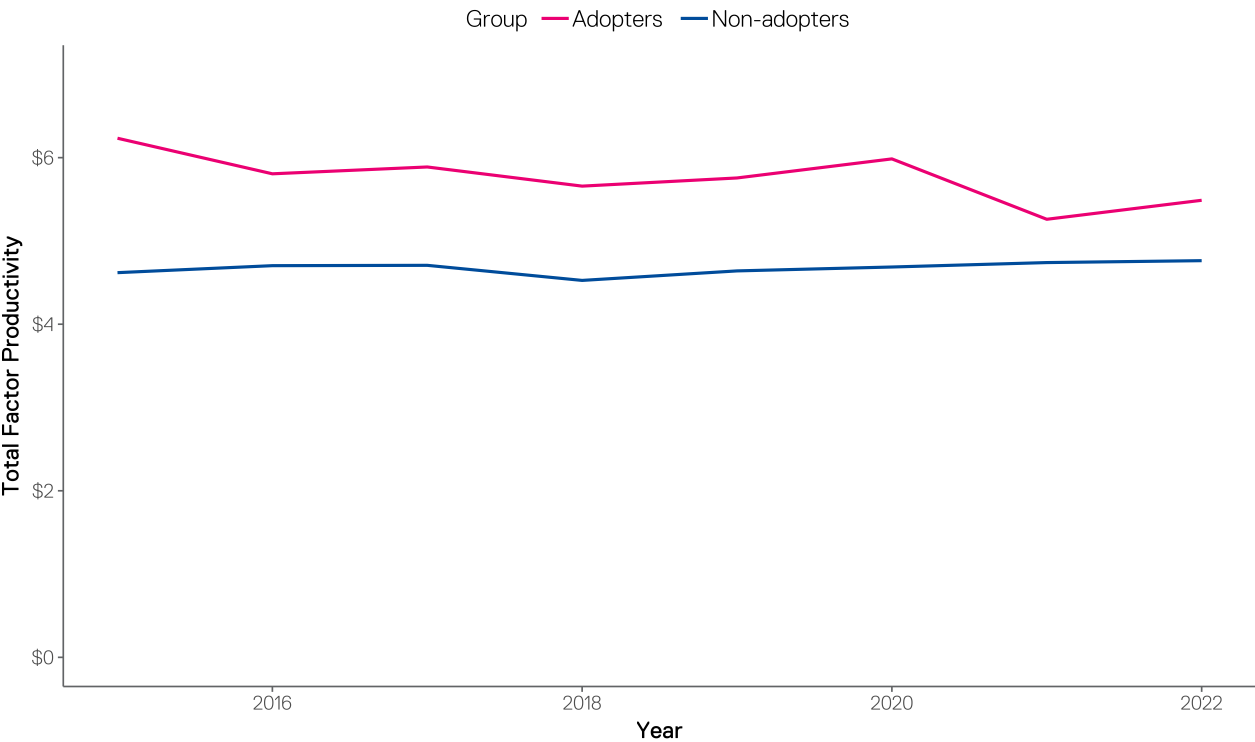
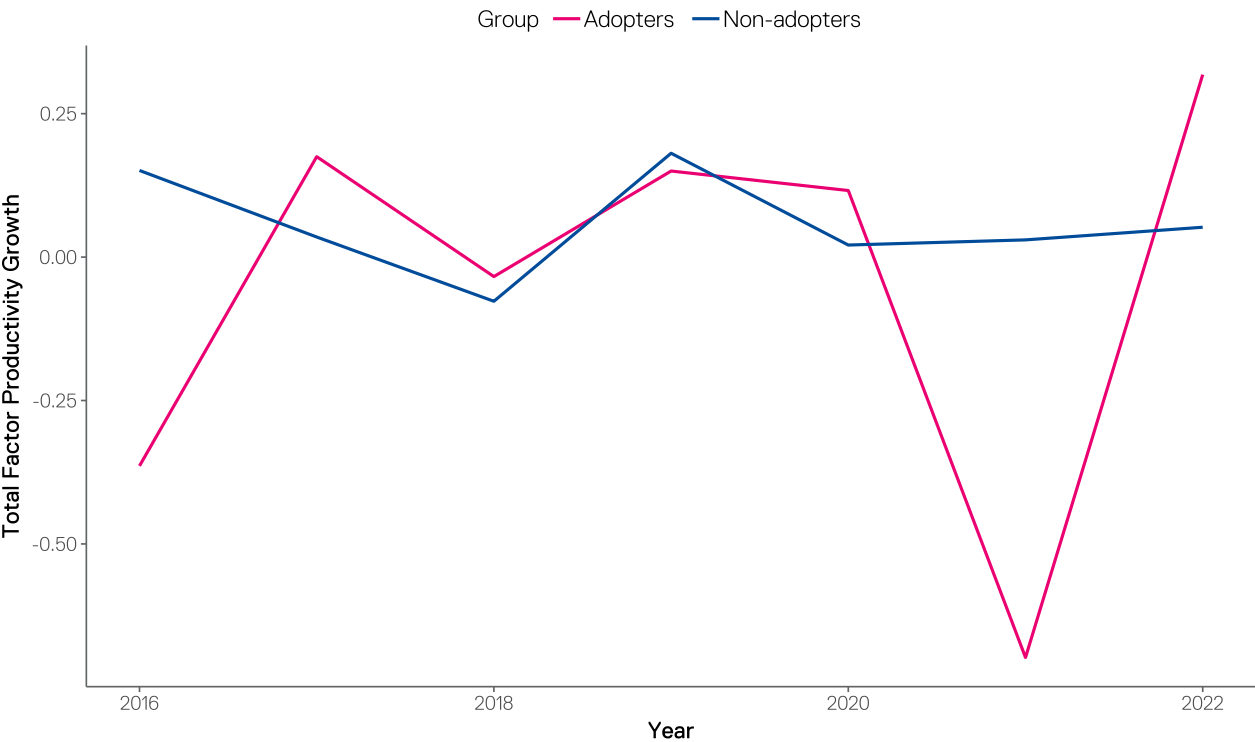


Figure 2: Total factor productivity growth by year and AI adoption



Causal modelling

Using a difference-in-difference approach, we find that the set of firms that adopted AI saw less productivity growth after adopting AI than other firms in Canada. Figure 3 shows that in 2022, when all firms in the treatment group had adopted AI, TFP growth was significantly lower within the treatment group than for the control group of non-adopting firms (represented in Figure 3 by the line at 0.0). However, in the period when firms adopted AI, the results are mixed—in 2020 growth was significantly higher and in 2021 it was significantly lower.


Our model, linking data from the Survey of Digital Technology and Internet Use with firm-level financial data, estimates the impact of a firm adopting AI on its subsequent productivity growth. This model compares the difference in productivity growth between firms that adopted AI between early 2020 and late 2021 with firms that did not adopt AI.⁴⁴

We highlight two important considerations. First, this data covers the period before the launch of

generative AI. Second, the exact time each firm in the treatment group adopted is unknown. While ideally the treatment (adoption of AI) would have been applied to all firms at the same time, in this case, it occurred at some unknown time between the end of 2019 (when the 2019 iteration of the SDTIU concluded) and the end of 2021 (when the 2021 survey concluded). That means the first year any impacts could show up would be in 2020, and all firms in the treatment group will have adopted AI by the beginning of 2022.

Using a difference-in-difference approach, we find that the set of firms that adopted AI saw less productivity growth after adopting AI than other firms in Canada. Figure 4 shows that in 2022, when all firms in the treatment group had adopted AI, TFP growth was significantly lower within the treatment group than for the control group of non-adopting firms (represented in Figure 4 by the line at 0.0). However, in the period when firms adopted AI, the results are mixed—in 2020 growth was significantly higher and in 2021 it was significantly lower.

Figure 3: Impact on adopting AI on productivity growth (relative to non-adopting firms)



Other ways of modelling this impact produce different results. Appendix B shows detailed results for five alternative models examining the relationship between AI adoption and TFP growth. While the specifications differ across the models, they consistently show no significant increase in TFP growth for firms adopting AI. Notably, however, they also do not show the same significant negative relationship that the difference-in-difference model produced.

Taken as a whole, these models suggest that the benefits of AI shown in more experimental work are not realized broadly among firms that have started using AI in their daily work over the study's short-term horizon (up to two years from the time of adoption). We also look at the relationship between AI adoption in firms and overall TFP rather than the rate of annual TFP growth. Results from these models are similar to those from the models for TFP growth. In most cases, there is no significant relationship between the adoption of AI by a firm in 2020 or 2021 and any difference in TFP.

When estimating the impact of AI adoption and the overall firm TFP, we find a significant and positive relationship between firms starting to use AI and other firms in our sample. The positive relationship among models estimating TFP level suggests that AI adopters are generally more productive than AI non-adopters, even before adoption. They also continue to be more productive on average after adoption.

In sum, the treatment group of AI adopters were already highly productive firms, but the decision to adopt AI did not increase the rate at which their productivity was growing.

Still, artificial intelligence adoption by Canadian businesses is in the early stages. As AI innovations continue, the technology diffuses to different business functions, and as more businesses integrate it, the existing pattern may change. None of this precludes TFP growth among the treatment firms over a longer time horizon, or TFP growth from future adoption by firms that are less productive to begin with.

The findings do call for caution in presuming that business adoption of AI can be a silver bullet in addressing Canada's productivity growth challenge in the near term.

In sum, the treatment group of AI adopters were already highly productive firms, but the decision to adopt AI did not increase the rate at which their productivity was growing.

7

Conclusion

The exuberance over AI and its potential benefits for the economy and society stems from an optimism that it can drive productivity growth, with the potential to improve work and boost the standard of living. Yet, the Solow paradox observes that, over various different waves of innovation, productivity growth has lagged or even declined in periods of technological adoption. Will this time be different?

The findings of this study, while preliminary, suggest that those who share this optimism should exercise caution about the productivity benefits of AI adoption in the near term. The research literature on productivity gains from AI use shows mixed results, with no conclusive evidence of a strong positive or negative relationship between the technology's adoption and productivity improvement.

Our economic model, examining a treatment group of Canadian firms that adopted AI in 2020 or 2021, finds no significant relationship between AI adoption and any difference in Total Factor Productivity (TFP) levels or TFP growth. Moreover, the subset of AI adopters were already highly productive firms, but the decision to adopt AI did not increase the rate at which their productivity was growing.

Yet, this research has limitations. Different waves of innovation are not necessarily predictive of future waves, so it is hard to extrapolate findings from one period to the next, and gains often occur over time. Given that adoption is relatively low in the broader Canadian economy, it is more difficult to derive insight into the impacts of AI on productivity—particularly compared to the impact of other more widely adopted general-purpose technologies.

Our economic model, examining a treatment group of Canadian firms that adopted AI in 2020 or 2021, finds no significant relationship between AI adoption and any difference in Total Factor Productivity (TFP) levels or TFP growth. Moreover, the subset of AI adopters were already highly productive firms, but the decision to adopt AI did not increase the rate at which their productivity was growing.

Measuring the impact of artificial intelligence on productivity through firm adoption in the SDTIU does not capture the extent to which the technologies are being used in core business functions. As applications in AI become more widespread and are increasingly embedded across various operations, there could be an increased chance for potential efficiencies to translate into increased productivity.

Importantly, this research also focuses on the impacts of AI before the latest boom in generative AI. Generative AI has a broader application range in the economy, but also larger implications for the economic welfare of Canadians, with early evidence suggesting decreasing demand for skills and tasks that are replaceable with generative applications.⁴⁵ Further research is needed into the impact of generative AI on productivity growth in Canadian firms. With data from the next iteration of the Survey of Digital Technology and Internet Use conducted in late 2023 through early 2024, this will be possible.

Appendix A

The methodology consists of a descriptive analysis and a causal analysis. The descriptive analysis provides a surface-level overview of the type of firms deploying AI, and their productivity on average for the most recent year of the SDTIU. This gives us a sense of the existing state of AI adoption in Canada. This includes a sub-population of 15,683 enterprises in 2021 (and with 14,127 enterprises in 2019), with a constructed TFP variable deconstructed by firm characteristics such as enterprise size and NAICS industry. For this analysis, we would require using the Business Research Microdata (BRM) in 2021, and firm-level responses in the 2021 SDTIU.

However, to understand the full extent of the impact of AI on business efficiency, we need to isolate the treatment effect of AI adoption. This would involve a causal analysis using a firm-level TFP variable as the dependent variable, with the treatment effect being the uptake in AI technologies by firms.^{46 47 48} As will be detailed in the data requirements section, for this step, we will leverage the BRM and responses on AI adoption from the SDTIU.

When examining the economy and the production process, economists often conceptualize what is called the “production function”, which transforms inputs (labour provided by people and entrepreneurs and machines, as well as other forms of capital) into outputs (final physical goods or services). TFP is intuitively understood as what enables the value of the output to exceed the total sum of values of the inputs. It is sometimes interpreted as technology or innovation, and growth in TFP allows the same set of inputs to create even more output than before, underlying the key logic around why economists and policymakers focus on this measure.

Empirically, estimating TFP is difficult, as it is conceptually fuzzy. The preferred approach, first popularized by American economist Robert Solow, is the idea of a “Solow residual” - that TFP is whatever is “left over” after we account for the value of all the input from the value of the output.⁴⁹ We follow that

similar approach in our work. In particular, we rely on previous work that explored estimating firm-level productivity^{50 51} (the Wooldridge-Levinsohn-Petrin or WLP process) that has also been shown to work in the Canadian data.⁵² This calculation is an adjusted approach on the Levinsohn & Petrin estimation method (2003),^{53 54} as laid out by Wooldridge (2009).⁵⁵

Specifically, the WLP process proposes a two-stage estimation where industry-specific production function is estimated. The resulting parameters for the industry-specific production are then used in the second stage, alongside firm-specific attributes (such as payroll, revenue, and capital valuation) to arrive at the firm-level productivity estimates. This involves estimating the following set of equations using an instrumental variable approach:

$$\ln(Y_{it}) = \beta_0 + \beta_L \ln(L_{it}) + \beta_K \ln(K_{it}) + \beta_m \ln(m_{it}) + \varepsilon_{it} \quad (1)$$

$$\ln(Y_{it}) = \beta_0 + \beta_L \ln(L_{it}) + \beta_K \ln(K_{it}) + f(g[K_{it-1}, m_{it-1}]) + \varepsilon_{it} \quad (2)$$

with Y_{it} as total firm revenue for the year t for firm i , L_{it} as total payroll for the year t and firm i , K_{it} as the total assets of the firm in year t for firm i , and m_{it} as the total cost of sale and intermediate goods in year t for firm i . In Equation 1, the full form of the production function is outlined, with output dependent on a firm’s decisions about labour, capital and intermediate inputs. In Equation 2, the function $f(g[K_{it-1}, m_{it-1}])$ is our productivity measure, which is a function of capital and intermediate input decisions in time $t-1$. We will follow WLP in assuming a third-degree polynomial functional form for these functions.

We observe two sets of firms: a control group of firms which did not take up AI technologies in 2019 or 2021, and a treatment group which did not deploy AI technologies in 2019 but did in 2021. In difference-in-difference estimation, we observe the impact of this intervention on TFP in the years post-implementation of AI. This allows us to observe the change to TFP in time t compared to time $t-1$, controlling for other time-specific and firm-specific factors such as the type of industry the firm is in (e.g. goods or services), firm size, etc.

Using the firm-level TFP estimates, we will run the following triple differences-in-differences (given that TFP is already a difference) estimation on the impact of AI adoption, taking the following form in Equation 3:


$$\Delta TFP_{it} = \beta_0 + \beta_1 Post + \beta_2 Treatment_{it} + \beta_3 (Treatment_{it} \times Post_t) + \beta_4 Service_{it} + \varepsilon_{it} \quad (3)$$

with $Post$ as a binary indicator representing the in-treatment group (firms who were an AI adopter in 2021 but not in 2019), $Treatment_{it}$ representing whether the period of TFP estimation was post-implementation of AI (in the year 2021), as the interaction term measuring the treatment effect on the treatment group, and $Service_{it}$ representing a binary indicator of a firm being in a goods or services industry.

In addition, fixed-effect Original Least Squares (OLS), an instrumental variable, and quantile regression models were constructed, which take the following form in Equation 4:

$$\Delta TFP_{it} = \beta_0 + \beta_1 Treatment_{it} + \beta_2 Year_{2021} + \beta_3 (Treatment_{it} \times Year_{2021}) + v_t + \beta_4 Service_{it} + \varepsilon_{it} \quad (4)$$

In the fixed-effect OLS, time-invariant firm-specific characteristics are cancelled out with TFP being a growth variable between periods t and $t-1$. When using TFP levels instead of growth, these characteristics are captured through a residual. Time-fixed effects in year-specific characteristics are captured with the residual for TFP levels and growth.



Appendix B

Table B-1: Difference-in-difference estimation, TFP growth controlling for goods/services industry indication and standard errors clustered on employment size groups

Time-period	Average Treatment Effect on the Treated (ATT)	Standard error	95% confidence interval (lower bound)	95% confidence interval (upper bound)
2017	0.52	0.22	0.22	0.82
2018	-0.09	0.91	-1.32	1.13
2019	-0.34	0.66	-1.22	0.54
2020	0.64	0.35	0.16	1.11
2021	-0.51	0.1	-0.64	-0.38
2022	-0.36	0.17	-0.59	-0.14

*Values in black font are statistically significant at the 5% level

Table B-2: Difference-in-difference estimation, TFP levels controlling for goods/services industry indication and standard error clustered on employment size groups

Time-period	Average Treatment Effect on the Treated (ATT)	Standard error	95% confidence interval (lower bound)	95% confidence interval (upper bound)
2016	-0.34	0.25	-0.77	0.09
2017	0.13	0.02	0.09	0.16
2018	0.26	0.71	-0.95	1.47
2019	-0.22	0.41	-0.93	0.48
2020	0.38	0.10	0.21	0.56
2021	0.14	0.28	-0.34	0.61
2022	0.29	0.18	-0.03	0.60

*Values in black font are statistically significant at the 5% level

Table B-3: Fixed-effects OLS regression, TFP growth controlling for goods/services industry indication and standard error clustered on employment size groups

Variable	Estimate	Standard error	T- value	P-value
Treatment group	-126.86	352.02	-0.36	0.75
Post-treatment period	0.00	0.03	0.10	0.93
Interaction: Treatment group & Post-treatment period	0.06	0.17	0.36	0.75

**Values in black font are statistically significant at the 5% level*

Table B-4: Fixed-effects OLS regression, TFP levels controlling for goods/services industry indication and standard error clustered on employment size groups

Variable	Estimate	Standard error	T- value	P-value
Treatment group	-521.43	262.76	-1.98	0.19
Post-treatment period	0.00	0.06	-0.03	0.98
Interaction: Treatment group & Post-treatment period	0.26	0.13	1.99	0.18

**Values in black font are statistically significant at the 5% level*

Table B-5: Instrumental variable (2nd stage) TFP growth controlling for goods/services industry indication and standard error clustered on employment size groups

Variable	Estimate	Standard error	T- value	P-value
Treatment group	-126.86	352.03	0.36	0.72
Post-treatment period	0.003	0.031	0.10	0.92
Interaction: Treatment group & Post-treatment period	0.06	0.17	0.04	0.97

**Values in black font are statistically significant at the 5% level*

Table B-6: Instrumental variable (2nd stage) TFP level controlling for goods/services industry indication and standard error clustered on employment size groups

Variable	Estimate	Standard error	T- value	P-value
Treatment group	521.43	262.76	0.03	0.98
Post-treatment period	-0.002	0.064	1.98	0.05
Interaction: Treatment group & Post-treatment period	0.26	0.13	1.99	0.05

**Values in black font are statistically significant at the 5% level*

Table B-7: Quantile regression, TFP growth controlling for goods/services industry indication and employment size groups

Quantile	Interaction: Treatment group & Post-treatment period	Standard error	T-Value	P-Value
25th	0.01	0.01	1.08	0.28
50th	0.01	0.01	0.86	0.39
75th	0.002	0.013	0.15	0.88

Table B-8: Quantile regression, TFP level controlling for goods/services industry indication and employment size groups

Quantile	Interaction: Treatment group & Post-treatment period	Standard error	T-Value	P-Value
25th	-0.04	0.22	0.18	0.86
50th	0.08	0.35	0.24	0.81
75th	0.37	1.27	0.29	0.77

Values in black font are statistically significant at the 5% level

Endnotes

- ¹ Carolyn Rogers, “Time to Break the Glass: Fixing Canada’s Productivity Problem,” (speech, Halifax Partnership, Bank of Canada, March 26, 2024), <https://www.bankofcanada.ca/wp-content/uploads/2024/03/remarks-2024-03-26.pdf>.
- ² Oliver Loertscher and Pau S. Pujolas, “Canadian Productivity Growth: Stuck in the Oil Sands,” *Canadian Journal of Economics* 57, no. 2 (May 2024): 478–501, <https://onlinelibrary.wiley.com/doi/10.1111/caje.12707>.
- ³ The Conference Board of Canada, “Real Talk: How Generative AI Could Close Canada’s Productivity Gap and Reshape the Workplace—Lessons From the Innovation Economy,” (February 20, 2024), https://www.conferenceboard.ca/wp-content/uploads/2022/10/real-talk_2024.pdf.
- ⁴ Mercer, “Workforce 2.0: Unlocking Human Potential in a Machine-Augmented World: Global Talent Trends,” (2024), <https://www.mercer.com/assets/global/en/shared-assets/local/attachments/pdf-2024-global-talent-trends-report-en.pdf>.
- ⁵ Viet Vu, *Race Alongside the Machines: Occupational Digitalization Trends in Canada: 2006–2021*, The Dais, (November 2022), <https://dais.ca/reports/race-alongside-the-machines>.
- ⁶ Erik Brynjolfsson, “The Productivity Paradox of Information Technology,” *Communications of the ACM* 36, no. 12 (1993): 66–77, <https://dl.acm.org/doi/10.1145/163298.163309>.
- ⁷ Angus Lockhart, *Automation Nation? AI Adoption in Canadian Businesses*, The Dais (2023), <https://dais.ca/reports/automation-nation-ai-adoption-in-canadian-businesses>.
- ⁸ Daron Acemoglu, “The Simple Macroeconomics of AI,” NBER Working Paper (2024), <http://doi.org/10.3386/w32487>.
- ⁹ Erik Brynjolfsson, Danielle Li and Lindsey Raymond, “Generative AI at Work,” arXiv, (April 23, 2023), <http://arxiv.org/abs/2304.11771>.
- ¹⁰ Fabrizio Dell’Acqua, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon and Karim R. Lakhani, “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” *Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013*, (2023), <http://dx.doi.org/10.2139/ssrn.4573321>.
- ¹¹ Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381, no. 6654 (2023): 187–92, <https://www.science.org/doi/10.1126/science.adh2586>.
- ¹² Peter Maw, Peter Solar, Aidan Kane, and John S. Lyons, “After the Great Inventions: Technological Change in UK Cotton Spinning, 1780–1835,” *The Economic History Review* 75, no. 1 (2022): 22–55. <https://doi.org/10.1111/ehr.13082>.
- ¹³ Getachew Amare and Bizuayehu Desta, “Coloured Plastic Mulches: Impact on Soil Properties and Crop Productivity,” *Chemical and Biological Technologies in Agriculture* 8, no. 1 (January 19, 2021): 4, <https://doi.org/10.1186/s40538-020-00201-8>.
- ¹⁴ Stuart Russell, Karine Perset and Marko Grobelnik, “Updates to the OECD’s Definition of an AI System Explained,” OECD AI Policy Observatory Portal, (November 23, 2023), <https://oecd.ai/en/wonk/ai-system-definition-update>.
- ¹⁵ Statistics Canada, *Survey of Digital Technology and Internet Use - 2021*, https://www23.statcan.gc.ca/imdb/p3Instr.pl?Function=assembleInstr&lang=en&Item_Id=1317562.
- ¹⁶ Roberta Capello, Camilla Lenzi and Giovanni Perucca, “The Modern Solow Paradox. In Search for Explanations,” *Structural Change and Economic Dynamics* 63 (December 2022): 166–80. <https://doi.org/10.1016/j.strueco.2022.09.013>.
- ¹⁷ Daron Acemoglu and Pascual Restrepo, “Automation and Rent Dissipation: Implications for Wages, Inequality, and Productivity,” NBER Working Paper (2024), <https://doi.org/10.3386/w32536>.
- ¹⁸ Ibid.
- ¹⁹ F. Calvino and L. Fontanelli, *A Portrait of AI Adopters Across Countries: Firm Characteristics, Assets complementarities and Productivity*, OECD Science, Technology and Industry Working Papers, No. 2023/02, (Paris: OECD Publishing, April 11, 2023), https://www.oecd.org/en/publications/2023/04/a-portrait-of-ai-adopters-across-countries_76004dec.html.
- ²⁰ Angus Lockhart, *Automation Nation? AI Adoption in Canadian Businesses*, The Dais, (August 14, 2023), <https://dais.ca/reports/automation-nation-ai-adoption-in-canadian-businesses/>.
- ²¹ Ibid.
- ²² Flavio Calvino and Lucia Fontanelli, “Artificial Intelligence, Complementary Assets and Productivity: Evidence from French Firms,” *LEM Working Paper Series*, No. 2023/35, <https://www.econstor.eu/bitstream/10419/297143/1/2023-35.pdf>.
- ²³ F. Calvino and L. Fontanelli, *A Portrait of AI Adopters Across Countries*, OECD, (April 11, 2023), https://www.oecd.org/en/publications/2023/04/a-portrait-of-ai-adopters-across-countries_76004dec.html.
- ²⁴ Erik Brynjolfsson, Danielle Li and Lindsey Raymond, “Generative AI at Work,” arXiv, (April 23, 2023), <http://arxiv.org/abs/2304.11771>.
- ²⁵ While TFP captures the efficiencies of more innovative technologies, it also includes other factors such as increased skill levels of workers or better resource allocation.
- ²⁶ Tim Sargent, “Productivity Growth in Canada: What Is Going On?” University of Calgary School of Public Policy Briefing Paper, Volume 17: 17 (October 2024), <https://doi.org/10.55016/ojs/sppp.v17i1.80080>.
- ²⁷ Dirk Czarnitzki, Gastón P. Fernández and Christian Rammer, “Artificial Intelligence and Firm-Level Productivity,” *Journal of Economic Behavior & Organization* 211 (July 2023): 188–205. <https://doi.org/10.1016/j.jebo.2023.05.008>.
- ²⁸ See Dirk Czarnitzki, Gastón P. Fernández, and Christian Rammer, “Artificial Intelligence and Firm-Level Productivity,” *Journal of Economic Behavior & Organization* 211 (July 2023): 188–205, <https://doi.org/10.1016/j.jebo.2023.05.008>; Giacomo Damioli, Vincent Van Roy and Daniel Vertesy, “The Impact of Artificial Intelligence on Labor Productivity,” *Eurasian Business Review* 11, no. 1 (2021): 1–25, <https://doi.org/10.1007/s40821-020-00172-8>; E. Brynjolfsson, Lorin M. Hitt, and Heekyoung Hellen Kim, “Strength in Numbers: How Does Data-driven Decision Making Affect Firm Performance?” (April 22, 2011), Available at SSRN: <http://dx.doi.org/10.2139/ssrn.1819486>; Maryam Ghasemaghahi and Goran Calic, “Does Big Data Enhance Firm

Innovation Competency? The Mediating Role of Data-driven Insights,” *Journal of Business Research* 104, (2019), 69–84, <https://www.sciencedirect.com/science/article/abs/pii/S0148296319304138?via%3Dihub>.

²⁹ Dean Alderucci, Lee Branstetter, Eduard Hovy, Andrew Runge and Nikolas Zolas, “Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Microdata,” (2020), <https://www.aeaweb.org/conference/2020/preliminary/paper/Tz2HdRna>.

³⁰ Taekyun Kim, Yejin Park and Wonjoon Kim, “The Impact of Artificial Intelligence on Firm Performance,” (Portland: 2022 Portland International Conference on Management of Engineering and Technology, August 7-11, 2022), pp. 1-10, <https://doi.org/10.23919/PICMET53225.2022.9882634>.

³¹ Listed on the Shenzhen or Shanghai stock exchanges.

³² Xueyuan Gao, and Hua Feng, “AI-Driven Productivity Gains: Artificial Intelligence and Firm Productivity,” *Sustainability* 15, no. 11 (2023), 8934, <https://doi.org/10.3390/su15118934>.

³³ Dirk Czarnitzki, Gastón P. Fernández, and Christian Rammer, “Artificial Intelligence and Firm-Level Productivity,” *Journal of Economic Behavior & Organization* 211 (July 2023): 188–205, <https://doi.org/10.1016/j.jebo.2023.05.008>.

³⁴ Fatima Shaikh, Gul Afshan, Rana Salman Anwar, Zuhair Abbas, and Khalil Ahmed Chana, “Analyzing the Impact of Artificial Intelligence on Employee Productivity: The Mediating Effect of Knowledge Sharing and Well-Being,” *Asia Pacific Journal of Human Resources* 61, no. 4 (2023): 794–820, <https://doi.org/10.1111/1744-7941.12385>.

³⁵ Vijaya Lakshmi and Jacqueline Corbett, “How Artificial Intelligence Improves Agricultural Productivity and Sustainability: A Global Thematic Analysis,” (Hawaii International Conference on System Sciences 2020, HICSS-53, January 7, 2020), https://aisel.aisnet.org/hicss-53/os/ai_and_sustainability/3.

³⁶ Kyogo Kanazawa, Daiji Kawaguchi, Hitoshi Shigeoka and Yasutora Watanabe, “AI, Skill, and Productivity: The Case of Taxi Drivers,” NBER Working Paper Series, (October 2022), <https://doi.org/10.3386/w30612>.

³⁷ Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381, no. 6654 (July 2023): 187–92, <https://doi.org/10.1126/science.adh2586>.

³⁸ Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,” (2023), arXiv Preprint, <https://arxiv.org/pdf/2302.06590>.

³⁹ Uplevel Data Labs, “Can Generative AI Improve Developer Productivity?” (2024), <https://uplevelteam.com/hubfs/Content%20Assets/Can%20Generative%20AI%20Improve%20Developer%20Productivity.pdf>.

⁴⁰ Robert M. Solow, “Technical Change and the Aggregate Production Function,” *The Review of Economics and Statistics* 39, no. 3 (1957): 312–20, <https://doi.org/10.2307/1926047>.

⁴¹ Statistics Canada, Classification of Size of Enterprise - Technology and Internet Use, <https://www23.statcan.gc.ca/imdb/p3VD.pl?Function=getVD&TVD=1262337&CVD=1262337&-CLV=0&MLV=1&D=1>.

⁴² One exception is applied in the case of quantile regression estimates, for which we include the variable as a control parameter in the model specification.

⁴³ A dip in productivity growth in 2021 among AI adopters is likely due to the impacts of COVID-19.

⁴⁴ For full model details, see Appendix A.

⁴⁵ Vivian Li and Graham Dobbs, Right Brain, Left Brain, AI Brain: AI Impact on Skill Demand in Canada’s Economy, The Dais, (2024).

⁴⁶ Zhonghua Cheng and Xiangwei Meng, “Can Carbon Emissions Trading Improve Corporate Total Factor Productivity?” *Technological Forecasting and Social Change* 195 (October 2023): 122791, <https://doi.org/10.1016/j.techfore.2023.122791>.

⁴⁷ Sam Hainsworth, “Does Artificial Intelligence Benefit UK Businesses? An Empirical Study of the Impact of AI on Productivity,” arXiv, (October 6, 2023), <https://doi.org/10.48550/arXiv.2310.05985>.

⁴⁸ Huwei Wen, Keyu Hu, Xuan-Hoa Nghiem and Alex O. Acheampong, “Urban Climate Adaptability and Green Total-factor Productivity: Evidence from Double Dual Machine Learning and Differences-in-Differences Techniques,” *Journal of Environmental Management* 350 (January 15, 2024): 119588, <https://doi.org/10.1016/j.jenvman.2023.119588>.

⁴⁹ Robert M. Solow, “Technical Change and the Aggregate Production Function,” *The Review of Economics and Statistics* 39, no. 3 (1957): 312–20, <https://doi.org/10.2307/1926047>.
⁵⁰ James Levinsohn and Amil Petrin, “Estimating Production Functions Using Inputs to Control for Unobservables,” *The Review of Economic Studies* 70, no. 2 (2003): 317–41, <https://doi.org/10.1111/1467-937X.00246>.

⁵¹ Jeffrey M Wooldridge, “On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables,” *Economics Letters* 104, no. 3 (September 2009): 112–14, <https://doi.org/10.1016/j.econlet.2009.04.026>.

⁵² Steven Denney, Viet Vu and Ryan Kelly, “Into the Scale-up-Verse: Exploring the Landscape of Canada’s High Performing Firms,” Brookfield Institute for Innovation + Entrepreneurship (2021), <https://dais.ca/reports/scale-up-verse>.

⁵³ James Levinsohn and Amil Petrin, “Estimating Production Functions Using Inputs to Control for Unobservables,” *The Review of Economic Studies* 70, no. 2 (2003): 317–41, <http://www.jstor.org/stable/3648636>.

⁵⁴ Inui Tomohiko and Kodama Naomi, “The Effects of Japanese Customer Firms’ Overseas Outsourcing on Supplier Firms’ Performance,” REITI Discussion Paper Series 16-E-106, (December 2016), <https://www.rieti.go.jp/jp/publications/dp/16e106.pdf>.

⁵⁵ Jeffrey M. Wooldridge, “On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables,” *Economics Letters* 104, no. 3 (September 2009): 112–14, <https://doi.org/10.1016/j.econlet.2009.04.026>.