Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives

Salomé Baslandze* Zachary Edwards John R. Graham
Ty McClure Michael Sparks Brent Meyer
Sonya Ravindranath Waddell Daniel Weitz

May 27, 2026

Abstract

We survey 734 executives to study how artificial intelligence (AI) affects productivity and labor. AI adoption is widespread but shallow at most firms, mostly focused on subscriptions. For workforce, companies on net anticipate small near-term AI-driven aggregate employment declines: larger (smaller) companies expect to reduce (increase) routine clerical (technical) positions more. Labor productivity gains are driven mainly by innovation- and demand-oriented channels. Like with past general-purpose technologies, we document a productivity paradox, likely reflecting a delay in revenue realizations. Unlike earlier technology waves, however, the typical firm exhibits little capital deepening, an indication that AI adoption mostly involves renting intangible capital.

JEL Classifications: O33, D22, J24.

Keywords: Artificial intelligence, productivity, technological change, labor markets, workforce composition, occupations.

*Baslandze: Federal Reserve Bank of Atlanta and CEPR; McClure, Meyer, Sparks, Weitz: Federal Reserve Bank of Atlanta; Edwards, Waddell: Federal Reserve Bank of Richmond; Graham: Duke University, The Fuqua School of Business and NBER (corresponding author, john.graham@duke.edu). We appreciate the excellent research assistance by Kathleen Barrow, Akshara Bassi, Mac Gilliam, Megan Trillo, and Ayush Vats. We thank the editor David Thesmar, the referee, Tania Babina, Jill Grennan, and participants in University of Maryland AI conference, the Federal Reserve Bank of Atlanta brown bag seminar and the Duke University brown bag seminar for helpful comments. We thank FEI, Nasdaq, and Duke University for surveying their financial executives to supplement our panel of CFOs. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of Richmond, and the Federal Reserve System. Internet Appendix can be found www.richmondfed.org/-/media/richmondfedorg/research/national_economy/cfo_survey/academic_publications/ai_survey_appendix.pdf. Explore deidentified survey data at the `AI Impact on Jobs & Tasks Dashboard`.

Conflict-of-interest disclosure statement

Salomé Baslandze: I have nothing to disclose.
Zachary Edwards: I have nothing to disclose.
John R. Graham: I have nothing to disclose.
Ty McClure: I have nothing to disclose.
Michael Sparks: I have nothing to disclose.
Brent Meyer: I have nothing to disclose.
Sonya Ravindranath Waddell: I have nothing to disclose.
Daniel Weitz: I have nothing to disclose.

1 Introduction

The rapid acceleration of artificial intelligence (AI) technologies and their deployment in business operations has drawn increasing attention from researchers and policymakers. Given the potentially transformative role of AI in the workplace, understanding its implications for corporate performance and labor market outcomes—including labor productivity, workforce composition, and employment size—has become a central economic question. Recent research emphasizes that the economic consequences of AI will depend critically on how the technology is deployed, with some views highlighting the risk of large-scale automation and worker displacement, but also the potential for AI to augment human labor, raise productivity, and increase economic growth.[1]

In this paper, we survey 734 financial executives ("CFOs") to document the speed and depth of AI adoption in its initial stages, as well as how AI usage will affect productivity and the size and composition of the workforce. An advantage of using survey data is that we can directly ask corporate decision-makers for timely current data as well as their expectations of the future. We specifically ask for their *AI-attributed changes* in workforce and productivity outcomes (rather than deducing AI effects using standard outcome data that likely reflect the combined effect of many micro- and macroeconomic forces). Our data focus primarily on the effects of using AI for the typical company in the economy, shedding light on initial adoption of AI across firms and its implications for productivity and workforce outcomes, rather than focusing on the relatively small set of firms responsible for developing AI technologies or building large-scale AI infrastructure.

Our data show that AI adoption is already widespread, though initial adoption is generally shallow (small dollar investments). Most spending–especially among smaller firms–takes the form of subscriptions, services, and training rather than hardware investment. We focus on two key dimensions of AI's impact: workforce size and composition, and labor productivity.

We examine in detail the effects of AI on the labor force, highlighting both workforce size and its composition across tasks and occupations. In the near term, AI has not led—and is not expected to lead—to meaningful reductions in aggregate employment. Overall effects are modest: firm-size- and sector-weighted employment is expected to decline by less than 0.4% due to AI in 2026. However, responses are heterogeneous: large firms expect workforce reductions, while smaller firms anticipate modest AI-driven employment growth.

Despite limited effects on total employment, CFOs expect a shift in the composition of

[1]See, for example, What Will U.S. Capitalism Look Like in 50 Years?", 2025; WSJ Monday, September 29. page R18; and Aghion et al. (2017).

their workers in terms of the tasks employees perform: away from routine clerical work and towards skilled-technical roles. By 2028, the share of employees that perform routine clerical work is expected to decline by more than 2 percentage points (mostly among large firms), with partially offsetting increases (mostly among small firms) in skilled technical roles (e.g., engineers, data analysts, scientists). This evidence suggests that labor force reallocation may occur less within-firm and more across the economy.

To further characterize these compositional changes, we analyze CFOs' open-ended responses describing the roles and responsibilities expected to be replaced or enhanced by AI tools. We map these responses to BLS Standard Occupational Classification (SOC) groups and construct a Negative Exposure Index (NEI) that compares the frequency with which firms describe AI as replacing versus enhancing work in each occupation. We find substantial heterogeneity across occupations and sectors: Office and administrative support roles exhibit the most negative exposure, consistent with the automation of routine clerical activities such as data entry. In contrast, many professional, technical, and sales-related occupations are frequently described as being enhanced by AI tools, suggesting complementarity with analytical and decision-oriented tasks.

Overall, our findings indicate that AI adoption is beginning to reshape the allocation of tasks and occupations, primarily by substituting for routine clerical activities while complementing higher-skill analytical and managerial work.

In the second half of the paper, we explore the effects of AI on productivity. We begin with revenue-based measures of labor productivity (revenue per employee) and find that AI-attributed productivity growth was already positive in 2025 (about 0.6 percent) and is expected to increase further in 2026. Productivity gains are largest in high-skill services and finance—where they exceed 2 percent— and are also positive across sectors, including manufacturing, construction, and lower-skill services.

To understand the sources of these gains, we decompose changes in revenue per worker into capital deepening (i.e., increased capital per labor) versus other channels. In contrast to previous IT revolutions, at the typical firm, only a small share of near-term productivity growth relates to capital deepening, reflecting that much AI spending—especially among smaller companies—takes the form of operating expenses (subscriptions) rather than capital investment. As a result, residual labor productivity—what we refer to as revenue-based TFP—accounts for the bulk of observed productivity gains, capturing potential improvements in efficiency and product quality, as well as changes in intermediate input use and markups.

In terms of mechanisms, productivity gains are most strongly associated with innovation- and demand-oriented channels–developing or improving products and services, and reaching or serving customers more effectively. These demand-driven gains raise revenues without leading to broad near-term employment declines. Operational efficiency, factor upgrading, and labor cost-cutting are not robustly associated with productivity gains, although the latter is somewhat more important for larger firms.

Finally, comparing these results with CFO-reported (that is, perceived) improvements in labor productivity due to AI, we find that perceived gains exceed those implied by revenue and employment data. This wedge likely reflects delayed output realization and quality improvements not yet captured in measured revenues, reminiscent of the classic "productivity paradox," whereby transformative technologies are widely viewed as important well before their effects are fully reflected in measured productivity.

The conclusion compares early AI adoption with past technology waves, emphasizing familiar patterns—uneven diffusion, productivity lags, and task reallocation—alongside distinct features: broad but shallow subscription-based adoption, upstream concentration of capital investment, and early productivity gains driven by innovation and demand-side channels.

2 Survey Data and AI Adoption

Near year-end 2025, we collected data on AI adoption and its implications for labor, capital, and productivity through two surveys of senior financial executives, most of whom are CFOs. The primary survey was fielded through The CFO Survey, conducted jointly by Duke University and the Federal Reserve Banks of Richmond and Atlanta, and was supplemented by a survey of senior financial decision-makers affiliated with FEI, NASDAQ, and Duke University alumni, yielding 734 total responses. The CFO data span 48 states, a wide range of firm sizes, and every major nonfarm nongovernmental industry. Appendix Section A1 provides details on survey design, coverage, summary statistics, and validation exercises.

We start by documenting stylized facts about firms' AI investment and usage–see Appendix Section A2 and a concise summary in Table 1. AI adoption is already widespread and expected to expand rapidly. Investment magnitudes vary widely across companies, with large firms (defined as those with 500 or more workers) spending substantially more in absolute terms, though investment per employee is comparable or higher among smaller companies. Most AI spending—especially among small firms—is directed toward operational expenses such as software subscriptions, services, and training, rather than hardware purchases or internal software development. Obstacles to investing in AI include lack of workforce training, a belief that AI technology is not yet sufficiently mature to be beneficial or has uncertain capabilities,[2] and concerns about privacy.

Table 1: Initial Stages of AI Adoption, Workforce, and Productivity

	Small Firms	**Large Firms**	**Evidence**
Panel A: Initial Stages of AI Adoption			
Adoption	Moderate in 2025, growing sharply in 2026.	Already widespread in 2025; expected even higher in 2026.	Figures A4, A6
Why not investing in AI?	Top reasons (similar for small and large firms): immature technology, untrained workforce, privacy concerns, and uncertain AI capabilities.		Figures A4, A5
Dollar investment	Median expected 2026 spending: $12,500; most companies plan to spend $20,000 or less.	Median expected 2026 spending: $300,000; about 30% plan to spend over $1 million.	Table A1; Figure A7
Investment intensity	Median per employee AI spending in 2026: $357.	Median per employee AI spending in 2026: $258.	Figures A8, A9
Type of spending	Operational expenses (subscription, services) dominate (65%).	Operational expenses dominate (54%); more internal development/hardware.	Figure A10
Use cases	Highest for marketing/product development, planning, forecasting, and reporting.	Similar top uses, with some edge in supply chain, quality control, and payment methods.	Figure A14
Takeaways: AI adoption	AI investment is broad but shallow: many firms are adopting or planning to adopt, but typical spending remains modest and largely subscription-based, especially among smaller firms.		
Panel B: Initial Stage AI Effects on Workforce and Productivity			
Number of Employees	Small increase	Modest decrease	Figures 2, A13
Workforce Composition	Reduction in routine-clerical (bigger for large firms); Increase in skilled-technical (bigger for small firms)		Figures 3, A12
Implied Productivity	1.62% increase in 2026	2.41% increase in 2026	Figure 4
Productivity Mechanism	Product innovation; reach/serve customers	Product innovation; cost reduction	Figure 4; Table A8

Notes: Large firms are those with at least 500 employees. Adoption in 2025 refers to firms reporting AI expenditures or financial investments over the previous 12 months; adoption in 2026 refers to firms reporting positive expected AI investment over the next 12 months. Reasons for not investing are reported among firms that did not invest in AI in 2025. Operational spending includes AI subscriptions, services, and training. Appendix figures and tables are available in the Internet Appendix, found at www.richmondfed.org/-/media/richmondfedorg/research/national_economy/cfo_survey/academic_publications/ai_survey_appendix.pdf.

In terms of outcomes from AI investment, Figure 1 reports responses to survey questions on realized (2025) and expected (2026) company outcomes *attributable to AI use* (Survey questions 8 and 9, described in Appendix Section A5). Executives expect AI to improve

[2] Menkhoff (2025) shows that information on AI's productivity potential updates firm beliefs and increases adoption.

decision speed and accuracy, output per worker, time spent on higher value-added tasks, customer satisfaction, and development of new products, while having relatively small if any effects in the near term on total employment and costs. In the next sections, we explore these patterns in detail to better understand the implications of AI investment on the workforce and labor productivity.

3 AI and the Labor Force

A large literature on automation and information technologies documents how past technological waves reallocated labor across tasks and occupations, with important consequences for employment, wages, and skill demand (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2022; Liu et al., 2025).

Artificial intelligence represents a potentially distinct technological shift, and its labor market implications remain uncertain. We provide initial-stage evidence on firms' realized and expected employment effects of AI, then turn to workforce composition, tasks, and roles.

3.1 Employment Effects of AI

We examine AI-driven employment changes using two complementary sources of qualitative and quantitative information. First, companies provide open-ended descriptions of roles and responsibilities that have been, or are expected to be, replaced or complemented by AI tools (survey questions 10 and 11; see Appendix Section A5). Second, we use quantifiable expected and realized changes in total employment attributable to AI, based on the survey questions 8 and 9.

At the extensive margin, most firms do not anticipate job replacement from AI. Based on the open-ended responses to a question asking which roles and responsibilities will be replaced by AI, 50% of companies explicitly report that AI will not replace any roles or responsibilities, and an additional 6% report being unsure. Only 44% write answers that indicate some degree of job replacement due to AI.[3]

These qualitative responses are consistent with firms' reported numerical employment changes attributed to AI: Firms indicating role replacement in the open-ended responses also report lower employment growth due to AI in the quantitative question. On average, companies reporting role replacement experience modest AI-related employment declines,

[3] A data dashboard summarizing responses to the open-ended questions is available here: `https://people.duke.edu/~jgraham/CFO_Survey_AI_Dashboard.html`.

Figure 1: Changes Attributable to Using AI

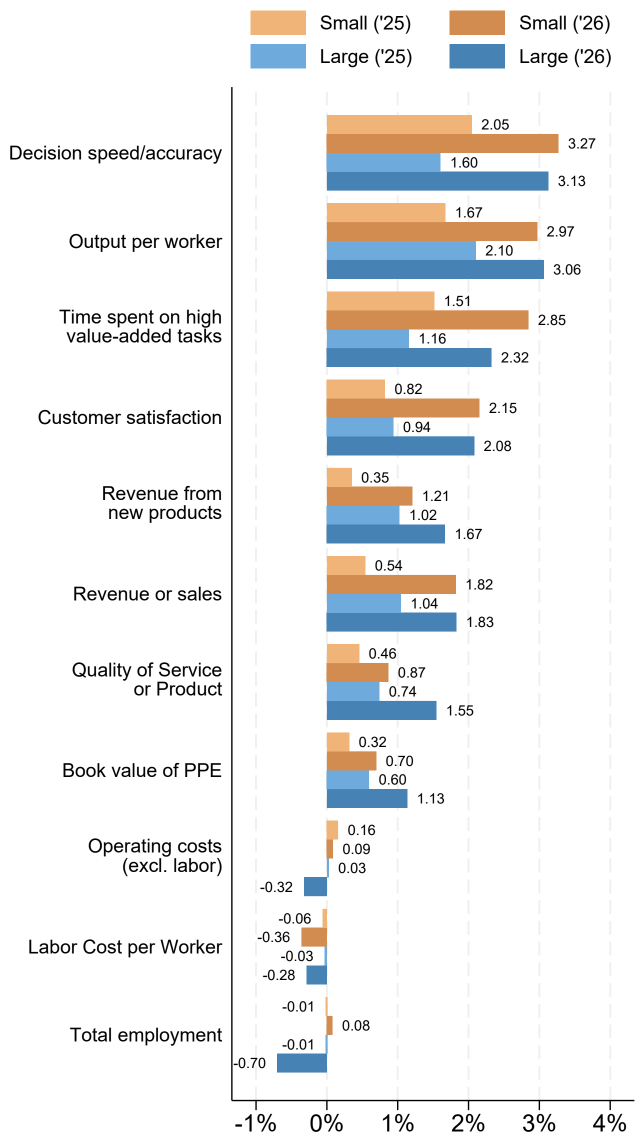

Notes: This figure depicts the realized (expected) percentage change in 2025 (2026) attributable to using AI usage for various company performance and cost metrics. The responses are categorical, with choices being "increased significantly (more than 10%)", "increased moderately (5.1 to 10%)", "increased somewhat (1 to 5%)", "little to no change", "decreased somewhat (-1 to -5%)", "decreased moderately (-5.1 to -10%)", "decreased significantly (more than -10%)". The numbers in this chart are based on the midpoint of each range, and ±5% above/below the highest/lowest option to calculate numeric averages across respondents. Responses indicating "no change" are coded as zero, while "not sure" responses are treated as missing. As examples of reading this chart, large companies (500^+ employees) expect a 3.13% increase in decision speed/accuracy in 2026 and a reduction in total employment of 0.7% *due to the usage of AI.*

whereas firms reporting no replacement show small positive employment effects (see Appendix Table A2).

To unpack this heterogeneity in employment effects, we examine AI-driven employment changes across sectors and by firm size, distinguishing between large and small companies. Figure 2 reveals that expected job losses increase with firm size and are concentrated in the finance and high-skill industries.[4] Panel (a) presents the average percentage change in employment due to AI, and Panel (b) presents the implied (weighted[5]) employment counts based on sectoral employment from the BLS Establishment Survey. In the latter, total headcount reduction for 2026 is expected to be 501,836 (less than 0.4% of the economy-wide workforce), with most of the reductions occurring in finance and high-skill services. Effects of this magnitude may be difficult to detect in aggregate statistics and are concentrated among large companies in certain sectors, rather than reflecting a widespread labor market disruption. The modest net employment effects we find are consistent with recent complementary analysis in Hampole et al. (2025); Gimbel et al. (2025); Yotzov et al. (2026); Massenkoff and McCrory (2026), though these papers do not document firm-size heterogeneity. Section 4.3 further examines potential mechanisms behind these limited aggregate employment effects in the near term.

3.2 Change in Labor Composition Due to AI: Shifting Tasks and Occupations

While the expected effects of AI on the broad labor market are modest to date, AI may affect the *composition* of employment within firms. AI adoption may induce companies to reallocate labor away from routine/clerical roles toward occupations that require higher skill, technical expertise, or managerial judgment. By design, many AI tools automate repetitive and standardized activities—such as data entry, transaction processing, or basic accounting—echoing patterns observed during earlier waves of IT adoption. However, unlike previous IT technologies, current AI systems are also capable of performing tasks traditionally considered high-skill or technical, including coding and data analysis. As a result, the types of tasks and occupations exposed to automation or augmentation may differ substantially from past technological transitions, and remain an open empirical question.

To address this question, we first examine firms' current and expected workforce composi-

[4]See Appendix Figure A12 for a more granular relationship between AI-driven employment changes and firm size.

[5]Both representativeness and importance weights are used for aggregation. Representativeness weights align the distribution of firms across sectors and size bins with the US economy; importance weights use each firm's 2024 employment share to give greater weight to economically larger firms.

Figure 2: Expected Employment Effects (2026) Attributable to AI Use

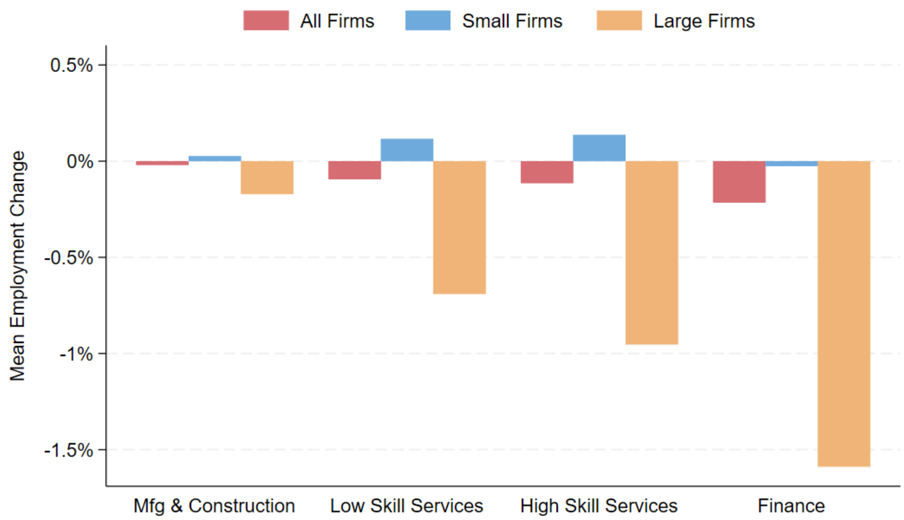

(a) AI-Driven Employment Change by Sector and Size

	Mfg & Construction	Low-Skill Services	High-Skill Services	Finance
Change in Employment	-57,665	-61,638	-271,290	-111,243
	[-124,403; 9,072]	[-209,484; 86,208]	[-443,692; -98,888]	[-148,024; -74,462]

	Aggregate
Change in Employment	-501,836
	[-741,394; -262,279]

(b) Implied AI-Driven Change in Aggregate Headcount

Notes: Panel (a) reports average (unweighted) percent change in employment by firm size and sector across all firms (survey question 9). Large firms are defined as those with 500 or more employees. Panel (b) reports sectoral and aggregate implied (weighted) changes in employment counts based on sectoral employment from the BLS Establishment Survey. Firm-level weights are the combined representativeness and importance weights. Representativeness weights are constructed to match the Census distribution of firms (count) by sector (four broad categories: manufacturing and construction, low-skill services, high-skill services, and finance) and firm size (1–99, 100–499, and 500+ employees). The importance weight is firm-level employment in 2024 (winsorized at 1st and 95th percentiles to reduce the impact of noise). In panel (b), 95% confidence intervals are shown in brackets.

tion in 2025, 2026, and 2028 (Survey Question 12). Survey respondents report the allocation of their workforce across routine/clerical roles (e.g., data entry, accounting), skilled technical roles (e.g., engineers, data analysts, scientists), creative and managerial roles (e.g., design, strategy, leadership), and other roles. These responses allow us to construct changes in the distribution of employment across occupation types relative to 2025.

Figure 3 summarizes expected changes in the composition of the workforce across sectors (Panel a) and firm size (Panel b) in 2026 and 2028. Panel (a) reveals an expected reduction by 2028 in the share of routine clerical workers of more than 2.5% in the finance, high-skill, and low-skill industries. These losses are offset somewhat by increases in technical hiring, especially in low-skill service industries. Interestingly, the near-term effects of AI on manufacturing and construction are relatively smaller, in contrast to earlier waves of robotics and automation. Overall, these findings imply that AI is expected to result in net white-collar job loss, with losses concentrated in routine clerical roles over the next few years.

Appendix Tables **??** and **??** corroborate that these shifts in employment composition are correlated with firms' AI investment. Firms with higher AI investment—both at the extensive and intensive margins—are significantly more likely to reduce the share of routine clerical workers. At the same time, AI investment is positively associated with an expansion of skilled technical employment, although the relationship is weaker and less precisely estimated at the intensive margin.

Building on the earlier evidence (Figure 2) of a strong firm-size gradient in AI-driven employment adjustments—where reductions in total employment are concentrated among larger firms—Panel (b) examines changes in workforce composition for small and large firms. While both groups reduce routine clerical employment on average, the declines are bigger among large companies. In contrast, both groups exhibit small, positive employment changes in skilled technical occupations. Appendix Figure **??** explores in more granular detail the heterogeneity across firm size, revealing meaningful labor reallocation *across firms*: the largest firms disproportionately shed routine jobs, while small firms exhibit relatively stronger growth in technical roles.

A second source of data on AI's impact on labor composition is from open-ended responses to questions on the roles and responsibilities expected to be *replaced* or *enhanced* by AI (Survey questions 10 and 11). These responses provide granular evidence on tasks and occupations exposed to AI—either negatively, through displacement, or positively, through augmentation. We map each response to BLS SOC groups.[6] For example, one response to

[6] Four research assistants coded responses into the 23 major SOC groups, and the research team reviewed and validated the classifications; these appear as Occupations in Panel (c) of Figure 3 and are the groupings on which the NEI index is based. The responses were further classified into 341 BLS job categorizations that appear as Sub-Occupations in the figure.

Figure 3: Changes in Workforce Composition

(a) Changes by Task and Sector

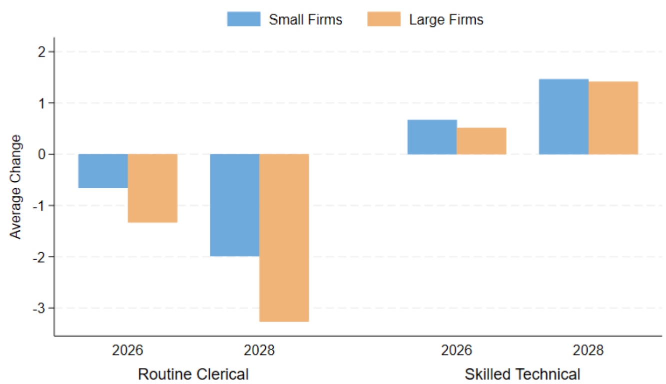

(b) Changes by Task and Firm Size

Occupation	Sub-Occupations	NEI (RM/EM)
Office and Administrative Support	Bookkeeping, Accounting, and Auditing Clerks; Office Clerks Customer Service Representatives	2.025
Business and Financial	Accountants, and Auditors; Financial Analysts; Management Analysts	0.829
Computer and Information Technology	Software Developers, Quality Assurance Analysts, & Testers; Computer Programmers; Computer Systems Analysts	0.596
Legal	Lawyers; Paralegals & Legal Assistants	0.471
Production	Assemblers & Fabricators; Metal & Plastic Machine Workers; Quality Control Inspectors; Machinists	0.308
Sales	Advertising Sales Agents; Wholesale & Manufacturing Sales Representatives; Sales Engineers	0.298
Management	Top Executives; Financial Managers; Advertising, Promotions, and Marketing Managers	0.138
Architecture and Engineering	Civil/Mechanical/Industrial Engineers; Architectural & Engineering Managers; Engineering Technologists & Technicians	0.100

(c) Negative Exposure Index by Occupation

Notes: Panels (a) and (b) shows the average expected changes in workforce composition in 2026 and 2028 by sector (a) and firm size (b). Changes are measured relative to 2025. Worker categories include routine/clerical roles (e.g., data entry, accounting), skilled technical roles (e.g., engineers, data analysts, scientists). Limited headcount effects are expected in creative or managerial roles (e.g., design, strategy, leadership), so these are not displayed in the figure, nor are "Other" workers not falling into any of the other job categories; therefore, employment changes across the displayed groups do not sum to 0%. Panel (c) summarizes AI exposure across major occupation groups based on firms' open-ended responses describing roles and responsibilities expected to be replaced or enhanced by AI (Survey Questions 10-11; see Appendix Section **??**). The Negative Exposure Index (NEI) is defined as the ratio of replacement mentions to enhancement mentions for which occupation group, aggregated across all firms.

the open-ended question on tasks enhanced by AI listed several tasks: "Accounting, Finance; Fraud Group; Customer Experience, Call Center." These responses were mapped to *Business and Financial Operations*, *Legal*, and *Office and Administrative Support*, respectively. This produces a firm×occupation dataset in which each pair is classified as *replacement*, *enhancement*, or *no mention*. The dataset includes 278 firms reporting replacement in various occupations and 536 reporting enhancement.[7]

To summarize exposure, we construct two measures: *replacement mentions* (RM) and *enhancement mentions* (EM), and define a *Negative Exposure Index* (NEI) as RM/EM. Values above one indicate that AI is more often described as replacing rather than enhancing work in that occupation. By construction, the NEI abstracts from employment shares and captures the *direction* of exposure, conditional on exposure (i.e., conditional on AI leading to either job replacement or enhancement).

Panel (c) of Figure 3 reports exposure across major occupation groups, ranked by the total number of mentions (restricting to those with at least 20 mentions). *Office and Administrative Support* occupations are most frequently mentioned and show clear negative exposure (NEI > 1), consistent with the automation of routine clerical tasks and the previously described decline in their employment share (Figure 2). In contrast, *Business and Financial Operations* exhibit roughly balanced replacement and enhancement, pointing to task reallocation rather than uniform displacement. Other groups—including sales, technical, and professional roles—are more often described as being enhanced by AI tools.[8]

To quantify how these replacement and enhancement mentions translate into employment outcomes, Appendix Table A6 relates firm-level AI-driven employment changes to measures of replacement and enhancement derived from open-ended responses. Firms reporting replacement across more occupations experience larger employment declines: an additional occupation described as replacement—controlling for enhancement—is associated with a 0.43 percentage point reduction in expected employment growth in 2026. This pattern indicates that qualitative assessments of AI's impact on roles and responsibilities align closely with realized employment adjustments.

Appendix Figure A15 presents word clouds from firms' open-ended responses, reinforcing these patterns: AI is most often described as complementing analytical, managerial, and

[7] If a firm reports both replacement and enhancement for the same occupation, we classify it as replacement. This reflects that replacement indicates declining demand for the occupation within the firm, whereas enhancement captures complementarity with AI.

[8] Appendix Figure A14 provides additional evidence on specific business tasks for which firms report AI to be most helpful. Among other findings, AI is currently less helpful in tasks requiring high-stakes, case-specific judgment such as M&A advisory or tax strategy. Appendix Table A5 reports sector-level exposure measures, showing that AI more often enhances than replaces workers in most sectors, though negative exposure varies with occupational composition and task structure.

information-processing tasks (e.g., marketing, accounting, and finance), while substitution is more associated with routine clerical and operational activities such as data entry, manual tasks, administrative support, and customer service. Some functions—notably marketing and accounting—appear in both categories, suggesting that AI may complement some tasks while replacing others within the same occupation, potentially leading to task reallocation within occupations.

Our findings on occupational exposure broadly align with existing AI exposure measures based on task content and estimated AI capabilities for augmenting or replacing those tasks, or on observed usage (e.g., Eloundou et al. (2024); Appel et al. (2026) and the Anthropic Economic Index), which identify the occupations most exposed to AI. However, those *exposure* measures alone cannot directly speak to actual changes in labor demand for certain occupations. Our approach complements these studies by capturing how firms themselves interpret AI's impact on tasks–specifically, whether it replaces or augments workers–thereby providing direct evidence on the expected direction of labor demand across occupations.

4 AI and Productivity

An active literature estimates the productivity effects of AI using both micro-level and aggregate approaches (Babina, 2026). Occupation-level and experimental studies document sizable gains in specific tasks and settings (Brynjolfsson et al., 2025; Cui et al., 2025), while aggregate frameworks predict more modest TFP effects over the medium run (Acemoglu, 2025; Aghion and Bunel, 2024). This section examines the labor productivity effects of AI at the firm level, focusing on the sources, heterogeneity across firms, and the mechanisms underlying productivity gains.

4.1 AI Paradox: Reported vs. Implied Labor Productivity Effects

To estimate firm-level labor productivity effects from AI, we analyze two complementary productivity measures based on survey questions 8 and 9 in the appendix, which ask firms about outcomes attributable to AI use.

First, we construct a measure of labor productivity growth implied by firms' reported effects of AI on revenues and employment. Specifically, the survey asks CFOs to report the percentage change in revenues and employment due to their company using AI. Let $\Delta \ln Y_{it}^{AI}$ denote firm i's self-reported AI-attributed revenue growth, and let $\Delta \ln L_{it}^{AI}$ denote the corresponding AI-attributed change in employment. We define *implied* AI-attributed labor productivity growth as

$$\Delta \ln(Y/L)_{it}^{AI} \equiv \Delta \ln Y_{it}^{AI} - \Delta \ln L_{it}^{AI}. \tag{1}$$

This measure serves as our primary survey-based measure of productivity outcomes and captures the implied labor productivity effect of AI based on firms' own assessments of how AI investment has affected—or is expected to affect—revenues and employment.

Second, we also use executives' direct responses on the percentage change in labor productivity (output per worker) attributable to AI over the past year (2025) and expected over the next year (2026). We refer to this measure as perceived, or *reported*, AI-related labor productivity growth, denoted by $\Delta \ln LP_{it}^{AI,CFO}$.

Panel (a) of Figure 4 shows average implied (hashed bars) and reported (solid bars) AI-related productivity growth for 2025 and 2026 across sectors. Two main patterns emerge.

First, productivity gains from using AI are positive and exhibit meaningful sectoral heterogeneity. For 2025, the mean implied labor productivity growth attributable to AI is 0.6% (0.6% weighted). Firms in high-skill services experience the largest gains, with implied labor productivity growth of about 0.8% in finance. Firms in low-skill services, manufacturing, and construction see smaller but still positive gains. These effects are expected to strengthen in 2026, with average implied labor productivity growth reaching 1.8% (1.9% weighted), with the largest anticipated increases again concentrated in high-skill services and finance, exceeding 2%. Though not shown in the figure, for the full sample, the median implied productivity growth in 2026 is 1.6% (2.4%) for small (large) companies.

While the revenue-based productivity gains we document are positive, they remain substantially smaller than those observed during the late-1990s IT-driven productivity resurgence, consistent with AI still being in an early phase of technological deployment (see Appendix Figure A16 for a comparison of IT and AI investment-in-technology timelines).

The second pattern from Panel (a) of Figure 4 is that firms consistently *report* larger productivity gains from AI than those *implied* by contemporaneous changes in revenue and employment; that is, $\Delta \ln LP^{AI,CFO}it > \Delta \ln(Y/L)^{AI}it$. This wedge likely reflects delayed output realization and quality improvements that are not yet captured in measured revenues. Firms might report improvements in task efficiency and production-time productivity when using AI, while a substantial share of time is still devoted to experimentation, implementation, and learning, whose benefits materialize only gradually in revenues. This interpretation aligns with the "productivity J-curve" in the IT literature (Brynjolfsson et al., 2021) and related evidence on delayed productivity effects of new technologies (Greenwood and Yorukoglu, 1997). More broadly, firms' conceptual notion of productivity appears to ex-

Figure 4: Labor Productivity Effects of AI: Reported, Implied, and Decomposition

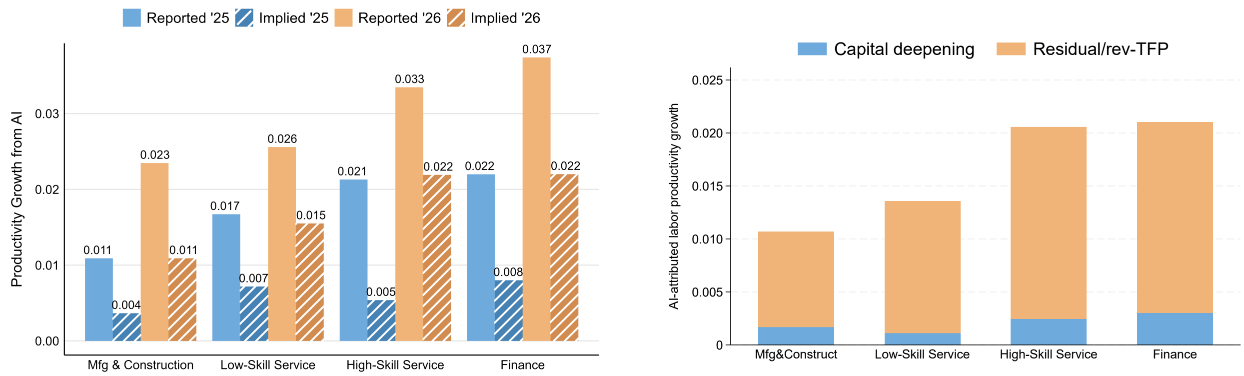

(a) Reported vs. Implied Productivity Effects of AI across Sectors

(b) Decomposing Implied Productivity Effects, 2026

Figure notes: Panel (a) shows mean reported (light-colored bars), $\Delta \ln LP_{it}^{AI,\,CFO}$, and implied (dark-shaded bars with hash marks), $\Delta \ln(Y/L)_{it}^{AI}$, labor productivity growth attributable to AI across sectors for 2025 and 2026. For the full sample, median implied productivity growth in 2026 is 1.62% (2.41%) for small (large) companies. Panel (b) decomposes expected implied AI-attributed labor productivity growth for 2026, $\Delta \ln(Y/L)_{it}^{AI}$, into capital deepening and the residual (revenue-based TFP), as defined in Equation (4). The 2025 decomposition corresponding to panel (b) appears in Appendix Figure **??**.

Labor Productivity Gains from AI: Mechanisms

	(1)	(2)	(3)	(4)
	2025		2026	
	Implied ΔLP	Implied ΔTFP	Implied ΔLP	Implied ΔTFP
Operational Efficiency Channel				
Production Efficiency	-0.002	-0.002	0.010**	0.009**
	(0.003)	(0.003)	(0.005)	(0.004)
Reduce Labor Costs	0.002	0.000	0.004	-0.000
	(0.004)	(0.003)	(0.005)	(0.004)
Reduce Other Costs	-0.002	-0.003	-0.002	-0.001
	(0.004)	(0.004)	(0.005)	(0.005)
Decision Making/Mgmt	-0.004	-0.005	-0.003	-0.004
	(0.003)	(0.003)	(0.004)	(0.004)
Innovation & Demand Channel				
Product Development/Improvement	0.008**	0.011***	0.012***	0.010***
	(0.003)	(0.003)	(0.004)	(0.004)
Reach/Serve Customers	0.008**	0.008**	0.011**	0.011***
	(0.003)	(0.003)	(0.004)	(0.004)
Factor Upgrading				
Upgrade Capital	-0.000	0.001	0.005	0.007
	(0.004)	(0.003)	(0.005)	(0.005)
Workforce Development	-0.002	-0.001	-0.006	-0.004
	(0.003)	(0.003)	(0.005)	(0.004)
Sector FEs	Yes	Yes	Yes	Yes
R^2	0.06	0.07	0.13	0.13
Observations	356	342	369	348

Table notes: This table reports firm-level regressions of changes in implied revenue-based labor productivity (ΔLP), and implied revenue-based total factor productivity (ΔTFP) in 2025 and 2026 on indicators for the importance of different motivations for AI investment. Each motivation is coded as a dummy equal to one if the firm rates it as *very important* or *extremely important*, and zero otherwise (*not at all*, *slightly*, or *moderately* important). All regressions include broad sector fixed effects (manufacturing and construction, low-skill services, high-skill services, and finance) and report heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

tend beyond contemporaneous revenue-per-worker calculations, encompassing improvements in workflows, task efficiency, and organizational capacity whose revenue effects accrue over time.

This pattern echoes classic "productivity paradox" arguments, whereby transformative technologies are widely perceived as important well before their effects are fully reflected in measured productivity.[9] In this vein, it is notable that *reported* productivity gains in 2025 closely align with *implied* gains for 2026 in all four sectors, consistent with a lag between AI adoption and realized revenue effects–as reflected in CFO expectations.

Appendix Figures A18–A20 confirm that our results are robust to alternative mappings of categorical survey responses into continuous measures and show heterogeneity by firm size and AI investment status.

Finally, Appendix Section A4.2 examines more closely the quantitative effects of AI investment—along both the extensive and intensive margins—on implied labor productivity growth and on the revenue-based productivity residual (formally defined in the next section). Firms that invest in AI experience larger productivity gains, while the elasticity of productivity with respect to investment intensity is modest, suggesting that much of the near-term labor productivity differences arise from whether firms adopt AI rather than from marginal increases in AI spending.

4.2 Decomposing Labor Productivity Growth from AI

Part of the labor productivity gains attributed to AI may reflect capital deepening—the possibility that AI-related investments increase the amount of capital used per worker, thereby raising output per worker. To clarify the sources of the implied AI-attributed labor productivity gains, we apply a standard growth-accounting decomposition (Jorgenson et al., 2008).

Under a Cobb–Douglas production function with constant returns to scale,

$$Y = AK^{\alpha}L^{1-\alpha}, \tag{2}$$

where Y denotes output, K capital, L labor, and A residual (total factor) productivity. In practice, because our data come from CFO surveys, we measure Y using revenue rather than value added. The decomposition below should therefore be interpreted as a revenue-based productivity accounting exercise rather than a strict value-added growth-accounting

[9]See Robert M. Solow, "We'd Better Watch Out," *New York Times Book Review*, July 12, 1987, for the original formulation of the productivity paradox: "You can see the computer age everywhere but in the productivity statistics."

decomposition.

Labor productivity can be written as

$$\ln(Y/L) = \ln A + \alpha \ln(K/L). \tag{3}$$

Taking differences and focusing on AI-attributed changes yields the following decomposition of implied AI-related labor productivity growth:

$$\Delta \ln(Y/L)_{it}^{AI} = \underbrace{\alpha\, \Delta \ln(K/L)_{it}^{AI}}_{\text{Capital deepening}} \quad + \quad \underbrace{\Delta \ln A_{it}^{AI}}_{\text{Residual / revenue-based TFP}}. \tag{4}$$

The first term captures capital deepening–any increase in capital per worker attributable to AI adoption. Using the identity

$$\Delta \ln(K/L)_{it}^{AI} \equiv \Delta \ln(K)_{it}^{AI} - \Delta \ln(L)_{it}^{AI},$$

we construct the capital-deepening contribution from AI-attributed changes in capital and employment. In our baseline specification, both $\Delta \ln(K)_{it}^{AI}$ and $\Delta \ln(L)_{it}^{AI}$ are taken directly from survey responses regarding AI-attributed changes in capital and employment. Changes in capital are proxied using survey-reported AI-attributed changes in the book value of property, plant, and equipment (PP&E) and intangible assets.[10]

The second term on the right side of Equation (4) captures residual *revenue*-based productivity changes not accounted for by capital deepening. This component reflects improvements in production efficiency, product or service quality, and pricing power, as well as AI-attributed changes in the quality, effectiveness, or use of intermediate inputs and labor. Because our measure of output is revenue rather than value added, this residual should be interpreted as revenue-based TFP and may therefore combine true efficiency gains with changes in intermediate input use and markups.[11]

Panel (b) of Figure 4 and Appendix Figure A17 decompose implied AI-attributed labor productivity growth into capital deepening and the residual (revenue-based TFP) using our baseline methodology. On average, capital deepening accounts for a moderate share—about 15%—of overall productivity growth among surveyed firms. Appendix Section A4.3

[10] We set α based on the standard aggregate value of 0.30, allowing for sectoral differences in capital intensity: $\alpha = 0.40$ for manufacturing and construction, $\alpha = 0.20$ for low-skill services, and $\alpha = 0.30$ for high-skill services and finance. These values are broadly consistent with sectoral capital income shares in standard growth accounting (e.g., BLS industry productivity statistics and EU KLEMS).

[11] Revenue-based TFP can exceed value-added TFP when AI reduces intermediate input costs or raises output prices, and can understate it when AI-driven revenue gains are accompanied by rising input expenditures.

shows that an alternative measure of capital deepening based on partial capitalization of AI expenditures yields quantitatively similar results, reinforcing the robustness of the baseline decomposition.

The relatively modest role of capital deepening for the average firm is not surprising. First, most firms remain in the early stages of AI adoption: while many report positive investment, spending remains small relative to firm size (Appendix Figures A8 and A9). Second, a substantial share of AI spending—especially for small firms—takes the form of operating expenses rather than capitalized investment (Figure A10). Subscriptions to AI tools (e.g., large language models), cloud computing, and SaaS support tasks such as coding, drafting, and analysis are expensed rather than recorded as capital. Nearly two-thirds of AI spending for small firms and more than half for large firms is subscription-based. As a result, much of AI's productivity impact operates through channels not captured by the average firm's capital deepening.

It is useful to compare this to the IT boom of the 1990s, when capital deepening was central to productivity growth: firms invested heavily in IT hardware and software, and *aggregate* labor productivity gains reflected both IT capital accumulation and TFP growth in roughly comparable measure (Jorgenson et al., 2008). More recently, Rubinton and Patro (2026) show that AI-related investment already contributes sizably to GDP growth in 2025, with magnitudes comparable to—or exceeding—IT's role during the dot-com boom. However, this investment is highly concentrated in a small set of large technology firms building data centers and cloud infrastructure. Consistent with this, our firm-level evidence suggests that for the *average firm*, current AI-driven productivity gains are only modestly driven by capital deepening. Many firms improve productivity without large firm-specific investment, instead accessing AI through cloud services, subscriptions to large language models, and other rented digital infrastructure. This shift—from firm-owned IT capital to centralized cloud and AI services—enables broad productivity gains without corresponding balance-sheet capital accumulation at the typical firm.

4.3 Productivity Gains from AI: Mechanisms

To better understand the channels through which companies experience revenue gains from AI, we leverage the survey questions on companies' motivations for AI investment, which include improving production efficiency, developing new or improving existing products and services, and reaching or serving customers more effectively (see Figure A11 for the full list). We encode each motivation as a dummy variable equal to one if the firm rates it as "very important" or "extremely important," and zero otherwise ("not at all," "slightly," or "moderately" important).

In the bottom portion of Figure 4, we regress implied labor productivity growth, and the revenue-based productivity residual, on these motivation dummies, allowing us to assess which channels are most strongly associated with observed productivity gains. Since these questions are only asked conditional on positive AI investment, the sample is restricted to firms that report positive AI investment. Across specifications, the motivations most strongly and consistently correlated with contemporaneous and expected future revenue gains are those related to developing new or improved products and services and reaching or serving customers more effectively. These demand-side and innovation-oriented motivations exhibit positive coefficients throughout and represent the strongest correlates of firm-level revenue gains among the mechanisms considered.[12] Corroborating this finding, Appendix Figure A14 shows that when firms are asked in a separate survey question which business tasks AI is most helpful for, they rank marketing and product development highest.

These findings suggest that AI-related productivity gains in our sample operate mainly through innovation and demand-side channels, consistent with a growing literature on general-purpose technologies, particularly AI, in fostering new products and services (Cockburn et al., 2018; Babina et al., 2024; Aghion et al., 2017) At the same time, recent studies show that advances in digital and information technologies improve matching between firms and customers, and that more effective customer acquisition is closely linked to higher revenues and expanded product creation (Gourio and Rudanko, 2014; Baslandze et al., 2023; Argente et al., 2025).

In contrast, in Table 4, motivations related to cost reduction, workforce development, capital upgrading, and enhanced decision-making or management do not show robust positive associations with revenue-per-worker gains, exhibiting statistically imprecise and often negative coefficients, although cost reduction gains more prominence for larger firms (see Table A8). These patterns are consistent with the view that such channels operate through deeper organizational and human-capital adjustments—such as training, task reallocation, and complementary capital investment—that involve adjustment costs and learning frictions. As a result, their productivity effects likely materialize gradually and are not fully captured in short-run revenue-based measures. This interpretation aligns with a large literature showing that gains from general-purpose technologies emerge with delay as firms upgrade skills, reorganize production, and accumulate complementary inputs (Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002).

[12] Appendix Tables A7 and A8 show similar patterns using weighted regressions. However, the importance of reaching or serving customers weakens when applying firm-size weights, suggesting that these gains are driven primarily by smaller firms.

Overall, the evidence points to innovation and market expansion as key short-run channels through which AI raises productivity, while cost-based and organizational channels may operate more slowly. This distinction also helps explain the limited employment effects in Section 3. Appendix Table A9 shows that innovation- and demand-oriented motives are associated with higher revenue, but not sizable employment reductions. Employment declines are more evident among firms emphasizing efficiency and labor-cost reduction, especially in routine-clerical roles. Thus, current AI use appears to support labor demand through revenue expansion, while labor-saving applications remain more limited, early-stage, occupation-specific, and in larger firms.

5 Conclusion

Parallels and Contrasts with Previous Technology Waves

Our evidence allows us to place AI adoption within the diffusion of a broader set of general-purpose technologies. Several patterns echo earlier technology waves: adoption is uneven and led by larger firms; measured productivity effects lag perceived usefulness as firms learn, reorganize workflows, and build complementary skills; and early labor-market adjustments are concentrated in routine clerical work. At the same time, AI differs from the 1990s IT wave in two important ways. First, adoption at the typical firm can occur with limited capital deepening: spending is broad but shallow and concentrated in subscriptions, services, and training rather than hardware or internal software development. This cloud- and subscription-based access may speed diffusion and make AI useful to a wider set of firms, even as aggregate AI capital deepening is concentrated among a small set of technology firms building data centers and cloud infrastructure. Second, the initial-stage productivity gains at typical firms are not yet strongly tied to labor-cost cutting (especially among small firms), but instead to innovation and demand-side uses such as product development and customer reach. Overall, AI resembles earlier general-purpose technologies in its diffusion frictions, productivity lags, and task reallocation, but differs in where investment occurs: much of the capital-intensive buildout is upstream, while user firms access AI through operational spending and realize gains through demand-enhancing channels.

Caveats and Future Research

There are several caveats to interpreting our findings as predictions of the eventual impact of AI. First, our survey focuses mostly on the initial stage of AI adoption – realized (2025) and expected (2026) – plus a limited look as far ahead as 2028. We think this short-term outlook is appropriate, given how difficult it is for companies to forecast more than one year ahead even in normal times (Graham, 2022). The explosive nature of new AI technology makes forecasting beyond the very near-term a formidable task. A related consideration is that non-AI-related shocks to the economy not fully anticipated at the time of our survey might affect the realized speed of AI adoption. As AI is integrated into the economy, future research should repeat the task of projecting the effects of AI on the corporate sector.

A second caveat relates to the nature of survey data. A strength of survey data is that we can ask questions designed to isolate causal mechanisms (e.g., employment changes *due to AI*), better understand company motivations and channels driving the results, and we can gather data in real-time. While in the appendix we document internal consistency and validate the quality of the responses along several dimensions, survey data have limitations. These include potential ambiguity in how respondents interpret questions, the influence of individual perceptions or optimism (e.g., Meyer and Weitz, 2025), and noise in reported answers. More broadly, as AI technologies continue to evolve, firms' understanding of their capabilities and implications may change, potentially altering their assessments and expectations. Future research should revisit these issues as the technology matures.

A third caveat is that we examine the effects of AI at the typical company, largely avoiding the important macroeconomic impact of the development of large-scale AI infrastructure and AI technology occurring in a small number of very large companies. To determine the economy-wide effects of AI, additional research is needed to complement the aggregate implications from our paper.

Lastly, our analyses focus mainly on gains and reallocation of resources among incumbent firms. Accordingly, a final caveat is that our survey does not explore the possibility that AI will lead to the creation of new companies that may be very productive in their use of AI, increasing aggregate economic growth and potentially affecting aggregate labor implications. Our data contain a hint that this may be an important consideration in that we find that innovation in products and services is a source of productivity gains among incumbent companies. Future research should explore the innovation channel in the context of the creation of brand-new, AI-productive firms.

References

Acemoglu, D. (2025). The simple macroeconomics of ai. *Economic Policy*, 40(121):13–58.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4B(none):1043–1171.

Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in u.s. wage inequality. *Econometrica*, 90(5):1973–2016.

Aghion, P. and Bunel, S. (2024). Ai and growth: Where do we stand? *Working Paper*.

Aghion, P., Jones, B. F., and Jones, C. I. (2017). Artificial intelligence and economic growth. NBER Working Paper 23928, National Bureau of Economic Research.

Appel, R., Massenkoff, M., McCrory, P., McCain, M., Heller, R., Neylon, T., and Tamkin, A. (2026). The anthropic economic index: Economic primitives. Technical report, Anthropic. Lead authors marked with * in original report.

Argente, D., Fitzgerald, D., Moreira, S., and Priolo, A. (2025). How do entrants build market share? the role of demand frictions. *American Economic Review: Insights*, 7(2):160–76.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration*. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Babina, T. (2026). Understanding firms' ai efforts and their economic impact. Working Paper 35123, National Bureau of Economic Research.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Baslandze, S., Greenwood, J., Marto, R., and Moreira, S. (2023). The expansion of varieties in the new age of advertising. *Review of Economic Dynamics*, 50:171–210.

Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, 117(1):339–376.

Brynjolfsson, E. and Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative ai at work. *Quarterly Journal of Economics*, 140(2):889–942.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–72.

Cockburn, I. M., Henderson, R., and Stern, S. (2018). The impact of artificial intelligence on innovation. Working Paper 24449, National Bureau of Economic Research.

Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2025). The effects of generative ai on high-skilled work: Evidence from three field experiments with software developers. Working paper.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Gimbel, M., Kinder, M., Kendall, J., and Lee, M. (2025). Evaluating the impact of AI on the Labor Market: Current state of affairs. The Budget Lab at Yale.

Gourio, F. and Rudanko, L. (2014). Customer capital. *The Review of Economic Studies*, 81(3):1102–1136.

Graham, J. R. (2022). Presidential address: Corporate finance and reality. *The Journal of Finance*, 77(4):1975–2049.

Greenwood, J. and Yorukoglu, M. (1997). 1974. *Carnegie-Rochester Conference Series on Public Policy*, 46(1):49–95.

Hampole, M., Papanikolaou, D., Schmidt, L., and Seegmiller, B. (2025). Artificial intelligence and the labor market. NBER Working Paper 33509, National Bureau of Economic Research.

Jorgenson, D. W., Ho, M. S., and Stiroh, K. J. (2008). A retrospective look at the u.s. productivity growth resurgence. *Journal of Economic Perspectives*, 22(1):3–24.

Liu, H., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Technology and labor markets: Past, present, and future; evidence from two centuries of innovation. NBER Working Papers 34386, National Bureau of Economic Research, Inc.

Massenkoff, M. and McCrory, P. (2026). Labor market impacts of ai: A new measure and early evidence.

Menkhoff, M. (2025). Belief updating and ai adoption: Experimental evidence from firms. CESifo Working Paper 12291, CESifo. Available at SSRN.

Meyer, B. and Weitz, D. J. (2025). In optimism we trust? explaining the disconnect between post-election optimism and own-firm expectations. *Policy Hub*, 2025(2):1–10.

Rubinton, H. and Patro, B. A. (2026). Tracking AI's contribution to GDP growth. St. Louis Fed On the Economy. January 12, 2026.

Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., Jalca, A., Meyer, B. H., Mizen, P., Navarrete, M. A., Smietanka, P., Thwaites, G., and Wang, B. Z. (2026). Firm data on ai. Working Paper 34836, National Bureau of Economic Research.