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, and Daniel Weitz

Working Paper 2026-4
March 2026

Abstract: We use novel data from a survey of nearly 750 corporate executives to study the effects of artificial intelligence (AI) on productivity and the workforce. We document substantial heterogeneity in AI adoption across firms, with more than half having already invested, though many smaller firms are only beginning to do so. Labor productivity gains are positive, vary across sectors, and are expected to strengthen in 2026, with the largest effects concentrated in high-skill services and finance. These gains are not primarily driven by firms' capital deepening but instead reflect increases in revenue-based total factor productivity, closely associated with innovation- and demand-oriented channels. We document a productivity paradox, in which perceived productivity gains are larger than measured productivity gains, likely reflecting a delay in revenue realizations. In labor markets, we find little evidence of near-term aggregate employment declines due to AI, though larger companies anticipate AI-driven workforce reductions, while smaller firms expect modest gains. We also find evidence of compositional reallocation of labor both within and across firms, with routine clerical roles declining and a relative demand for skilled technical roles increasing. We develop an index that ranks job functions most negatively affected by AI.

JEL classification: O33, D22, J24

Key words: artificial intelligence, productivity, technological change, labor markets, occupations

https://doi.org/10.29338/wp2026-04

The authors appreciate the excellent research assistance by Kathleen Barrow, Akshara Bassi, Mac Gilliam, and Ayush Vats. They thank participants in the Federal Reserve Bank of Atlanta brown bag seminar and the Duke University brown bag seminar for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Atlanta, Federal Reserve Bank of Richmond, and the Federal Reserve System.

Salomé Baslandze is with the Federal Reserve Bank of Atlanta and CEPR. Ty McClure, Brent Meyer, Michael Sparks, and Daniel Weitz are with the Federal Reserve Bank of Atlanta. Zachary Edwards and Sonya Ravindranath Waddell are with the Federal Reserve Bank of Richmond. John R. Graham is with the Fuqua School of Business at Duke University and NBER.

1 Introduction

The rapid acceleration of artificial intelligence (AI) technologies and their deployment in business operations has drawn increasing attention from researchers and policymakers. The pace of AI development and adoption has surprised many observers and prompted comparisons to earlier waves of technological diffusion, such as the spread of personal computers.[1] 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 question. An ongoing discussion 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 economic growth.[2]

Companies drive AI investment, and many AI-driven productivity effects will take place in the workplace. It follows that financial decision-makers are well-suited to provide insight into corporate investment, utilization, and expected outcomes for AI. In this paper, we survey nearly 750 financial executives ("CFOs") to gather novel data about the use of AI at their firms, which allows us to document new findings about how AI usage has affected and is expected to affect productivity and the size and composition of the workforce for the typical firm adopting AI. An advantage of using survey data is that we can directly ask corporate decision-makers for timely historic data and their expectations of the future. Moreover, we can specifically ask for their AI-attributed changes in workforce and productivity outcomes (rather than trying to deduce 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 the early diffusion of AI across firms and its implications for productivity and workforce outcomes, rather than on the relatively small set of firms responsible for developing AI technologies or building large-scale AI infrastructure.

We find that more than half of companies have already invested in AI, but adoption varies widely, with many smaller firms only beginning to invest in 2026. Obstacles to investing in AI include lack of workforce training, a belief that AI technology is not yet sufficiently mature to be beneficial, and concerns about privacy. Among companies that invest in AI, small firms invest somewhat more intensively per employee than do large firms. Small companies predominantly incur AI-related operating expenses (such as subscriptions to use AI software); large companies also have substantial operating expenditures, but also spend on hardware and internal development to tailor AI tools to their own company needs.

Our analysis of the impact of AI focuses on two key dimensions: labor productivity, and workforce size and composition.

[1]See Appendix Figure A4 for a comparison of the pace of adoption across earlier IT diffusion episodes and the recent AI investment cycle.

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

The productivity analysis yields several new findings. First, CFO-reported (that is, perceived) improvements in labor productivity due to AI are substantially larger than the revenue-based productivity gains implied by observed changes in revenue and employment due to AI. This wedge—likely reflecting delayed output realization and quality improvements that are not yet captured in measured revenues—aligns with the classic "productivity paradox," whereby transformative technologies are widely viewed as important well before their effects are fully reflected in measured productivity.

Second, we find that implied revenue-based labor productivity effects in 2025 are positive, display meaningful sectoral heterogeneity, and are expected to roughly double in 2026. In 2025, firms in high-skill services—particularly finance—experience the largest gains, with implied annual labor productivity growth of about 0.8 percent. Firms in low-skill services, manufacturing, and construction see smaller but still positive gains of roughly 0.4 percent. These effects are expected to strengthen in 2026, with the largest anticipated increases again concentrated in high-skill services and finance, with implied productivity gains exceeding 2 percent.

Third, we decompose implied labor productivity gains to assess the role of capital deepening (i.e., increased capital/labor) versus other mechanisms driving AI-related revenue growth per employee. We show that at the typical firm, only a small portion of near-term labor productivity gains are driven by capital deepening, reflecting the early stage of AI adoption and the fact that much AI spending—particularly among smaller firms—takes the form of operating expenses rather than capitalized investment. As a result, residual labor productivity, which we refer to as revenue-based TFP (total factor productivity), accounts for the bulk of observed productivity gains and captures potential improvements in efficiency and product quality, as well as changes in intermediate input use or markups.

To understand the mechanisms underlying these gains, we relate implied revenue-based productivity gains to firms' stated motivations for AI investment. We find that productivity gains are most closely associated with innovation- and demand-oriented channels rather than cost reduction alone. In particular, motivations related to developing new or improved products and services, and reaching or serving customers more effectively, are the strongest and most consistent correlates of both contemporaneous and expected future revenue productivity gains.

In the second part of our analysis, we examine the effects of AI on the labor force, focusing on both the size of the workforce and its composition across tasks and occupations. Our central finding is that, in the near term, AI has not led—and is not expected to lead—to meaningful reductions in aggregate employment. The overall employment effects are small. In particular, firm-size-and-sector-weighted aggregate employment is expected to decline by less than 0.4% due to AI in 2026. At the same time, employment responses are heterogeneous across firm size: large companies expect to shed workers due to AI adoption, whereas smaller firms anticipate modest employment growth associated with AI.

Despite limited effects on total employment, CFOs expect a shift in the composition of their workers in terms of the tasks employees perform: away from routine clerical work and more towards skilled-technical roles. Over the next three years, the share of routine clerical employment is expected to decline by more than 2 percentage points (mostly among large firms), with partially offsetting increases in skilled technical roles (e.g., engineers, data analysts, or scientists) and other positions (mostly among small firms). Thus, the reallocation of the labor force 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 of employees 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. The results reveal 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 more frequently described as being enhanced by AI tools, suggesting complementarity with analytical and decision-oriented tasks. Overall, these patterns 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.

The paper proceeds as follows. Section 2 reviews the related literature and discusses how our findings contribute to existing work on the productivity and workforce effects of technological waves. Section 3 describes the survey design and provides descriptive evidence on AI adoption across firms. We then examine the current and expected future implications of AI in two parts: productivity effects in Section 4 and labor market effects in Section 5. Section 6 concludes with caveats and directions for future research.

2 Literature Review

This paper builds on two literatures related to technology in the economy: one seeks to understand the implications of technological developments on productivity, and the other explores the implications of new technologies for the labor market. At the end of the 1980s, just as computers were becoming ubiquitous, there was a slowdown in U.S. productivity growth, prompting Nobel Laureate Robert Solow to comment that "we can see computers everywhere but in the productivity statistics." Explaining this 'Solow Paradox' inspired work that sought to ascertain the impact of information and communication technologies (ICTs) on productivity. Gordon (2000) famously argued that the ICT economy did not measure up to the great productivity-enhancing inventions of the past, but Fernald and Ramnath (2004), among others, document the acceleration in total factor productivity (TFP) in the 1990s and argue that the benefits of ICT may be subtle, affecting measured TFP in sectors where firms reorganize production in order to take advantage of new technologies.

In 2017, Brynjolfsson et al. (2017) concluded that AI had not yet shown up in the productivity statistics—a "redux of the Solow Paradox"—because it takes time to sufficiently diffuse and harness the most impressive capabilities of new technologies. More recently, some have attributed the step up in U.S. productivity in 2024 and 2025 at least partly to the rising incorporation of AI, but it is not entirely clear what role AI is playing in this productivity boost. It is even less clear what the U.S. economy can expect from AI in the future.

Nonetheless, there is a growing literature estimating the impact of AI on productivity. In a study by Brynjolfsson et al. (2025), the authors use data on customer support agents and find that access to AI assistance significantly increased worker productivity in this setting. As another occupation-specific example, Cui et al. (2025) evaluate the impact of generative AI on software developer productivity via randomized controlled trials and find a significant increase in completed tasks, albeit with considerable heterogeneity. Multiple researchers and institutions have continued in this vein, looking at the productivity value of AI tools such as Claude, Copilot, and ChatGPT. More broadly, Acemoglu (2025) employs the task-based framework developed in earlier papers and predicts modest TFP gains of at most 0.66% over the next 10 years. Aghion and Bunel (2024) use a similar framework and estimate that AI should increase aggregate TFP growth by between 0.07pp and 1.24pp per year over the next decade. They also use a different approach that exploits the parallel between the AI revolution and past technological revolutions; in this framework, they estimate that the AI revolution should increase aggregate productivity growth between 0.8 and 1.3pp per year over the next decade.

The literature on the role technology plays in the workforce is similarly deep and often focused on how changing demand for skills impacts relative wages. Davis and Haltiwanger (1991) argue that the increased skill intensity required in the manufacturing sector (skill-biased technical change) was the driving force behind rising wage inequality from 1963 to 1982. This finding is corroborated by Katz and Murphy (1992) where they emphasize the rapid secular growth in the demand for more skilled workers, more educated workers, and women. Later studies (e.g., Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018)) use the task-based framework to analyze the role of technological change on the labor market; Acemoglu and Restrepo (2022) find that the displacement of tasks induced by automation in industries facing rapid automation account for most of the changes in the U.S. wage structure since 1980.

There is some recent evidence that AI is impacting the workforce but that, like in the case of productivity and in the case of other technological change, it will take time. Gimbel et al. (2025) do not (yet) find a discernible shift in occupational mix since ChatGPT's release, arguing that widespread effects will take longer to materialize. Based on a model by Hampole et al. (2025) that links task-specific technological advances to overall labor demand for an occupation, Liu et al. (2025) find that in past decades, technology-led innovation led to economically meaningful shifts in labor demand across occupations, with increased demand for occupations with higher education requirements, occupations that pay higher wages, and occupations with a greater fraction of female workers. On the other hand, they find that this time might be different: in contrast to the past two centuries, their framework suggests that AI—by substituting primarily for cognitive tasks—will shift relative demand toward occupations with lower education, lower wages, and a greater share of male workers.

Even with a growing body of research exploring the AI revolution, the future is still unclear. This paper adds to our understanding of AI adoption and its implications using a set of questions introduced into The CFO Survey to understand AI adoption among surveyed firms and the implications for productivity and the labor force. This adds to the body of work trying to understand AI through surveys that have been preliminarily documented by Crane et al. (2025). Yotzov et al. (2026) conduct a large, international survey of corporate managers about AI usage and impact during the same year-end-2025 time frame as our study. For the few overlapping questions between their survey and ours, their findings generally corroborate ours: small past changes in employment and productivity, with moderate changes expected in the near future.

3 Data From a Survey of Corporate Executives

We collect data on AI adoption and its implications for labor, capital, and productivity through two waves of surveys of senior financial executives, most of whom are CFOs. The primary survey was in the field from November 11 to December 16, 2025, and produced 603 responses; it targets panel members of The CFO Survey, a quarterly survey conducted jointly by Duke University and the Federal Reserve Banks of Richmond and Atlanta. We supplement this sample with surveys sent to senior financial decision-makers at firms that are members of Financial Executives International (FEI) and/or NASDAQ, as well as Duke University alumni; these supplemental surveys ran from mid-December 2025 to mid-January 2026 and yielded 145 additional responses.

A key advantage of our survey design is that executives are asked directly about *AI-attributed* changes in corporate outcomes—both current and expected—allowing us to isolate the effect of AI rather than possibly conflating it with other concurrent drivers of productivity and employment that are difficult to disentangle in standard econometric settings. Combined with the timeliness of the data, which captures real-time and near-term expectations of corporate decision-makers, these questions form the empirical foundation for our analysis of the productivity and labor market effects of AI adoption.

Table 1 presents summary statistics for key variables used in our analysis.[3] Like the US economy broadly, the sample contains a small number of very large firms and a large number of small-to-medium-sized firms. The median (mean) firm had revenue of \$46 million (\$3.5 billion) and 118 (2,715) employees, with 22.0% of companies having at least 500 employees. The median (mean) firm expects to invest \$35,000 (\$8.3 million) in AI in 2026. The mean reported increase in labor productivity (ΔLP) attributable to AI investment was 1.8% in 2025 and is expected to reach 3.0% in 2026.

We validate the quality of our sample along three dimensions: sample coverage and forecast accuracy relative to both same-firm realized outcomes and aggregate economy outcomes. (See

[3] Appendix A5 presents variable definitions and the survey questions underlying each variable.

Appendix A1 for further details.) First, the CFO Survey sample spans all 50 states, a wide range of firm sizes, and every major nonfarm industry,[4] with the four most frequently represented industries being finance and insurance, manufacturing, professional and business services, and information technology. Because the sample is not perfectly representative of the US economy by construction, we re-weight our aggregate analysis to align the sample's industry and size distribution with the US Census. Specifically, our baseline results are unweighted and should be interpreted as applying to the average firm in our sample. When making statements about aggregate implied effects, we apply two sets of weights: *representativeness* weights that align the distribution of firms in our sample with the US Census by broad sector (4 bins) and firm size (3 bins), and *importance* weights that aggregate firm-level outcomes to economy-wide totals using each firm's 2024 employment share.[5]

Second, as shown in the appendix, year-ahead forecasts by CFO Survey respondents align closely with same-firm realizations for sales revenue growth, output price growth, input cost growth, and wage growth. This within-firm forecast accuracy lends credibility to the forward-looking expectations we use throughout the paper.

Third, aggregate sales growth constructed from CFO Survey forecasts tracks closely with economy-wide sales from the Bureau of Economic Analysis. This alignment, corroborated by external research using the long history of the CFO Survey (Gennaioli et al., 2016), supports the use of our panel to draw inferences about the broader US economy.

3.1 AI Investment

To assess the prevalence, scale, and composition of AI investment across firms, we begin with descriptive evidence about firms' AI expenditures. We ask CFOs whether and how much their companies had invested in AI in 2025, and how much they expect to invest in 2026. Panel A of Figure 1 reports the proportion of firms investing, by sector. Companies in high-skill sectors have invested (and plan to continue investing) the most in AI. Also, across all sectors, the percentage of companies investing in AI is expected to grow substantially from 2025 into 2026. Appendix Table A5 shows that many small firms are coming off the sidelines in 2026, with fewer than half having invested in AI by 2025, but a marked increase in investment is expected in 2026.

While AI adoption is widespread, there have been barriers to AI adoption in that more than 40% of companies did not invest in 2025. Panel B lists the main reasons for companies not investing in AI. As of year-end 2025, 42% of firms indicated that they view AI technology as too immature to have invested in, while 36% indicate that their workforce is not yet trained to use AI in the workplace, and 36% say that they have privacy concerns.

The survey respondents indicate the dollar amount they spent on AI in 2025 and 2026 (see Figure 2, which presents results separately for small firms vs. large firms, where large is defined as companies with at least 500 employees). The figure shows a wide range of spending, with 30% of

[4]See `https://www.richmondfed.org/research/national_economy/cfo_survey/survey_participation`.

[5]Since our variables of interest are percent changes in employment, productivity, and related outcomes, the appropriate importance weights are prior-year (2024) revenue or employment shares.

Table 1: Summary Statistics

		N	Mean	Median	Std. Dev.	Pctl. 10	Pctl. 25	Pctl. 75	Pctl. 90
Panel A: General Demographics									
	Large Firm (dummy)	722	0.220	0.000	0.415	0.000	0.000	0.000	1.000
	No. Employees	722	2,715	118	16,049	5	29	418	1,800
	Revenue ('25)	599	$ 3,461,042	$46,000	$ 25,154,316	$1,771	$ 10,300	$252,000	$ 1,743,775
	Book Value PP&E	660	$633,936	$6,000	$3,961,175	$ 10	$300	$56,703	$520,000
	Routine Job Share	637	0.190	0.100	0.205	0.020	0.050	0.250	0.500
	Technical Job Share	636	0.247	0.120	0.276	0.000	0.010	0.435	0.700
	Creative Job Share	637	0.191	0.110	0.203	0.020	0.050	0.250	0.470
	Other Job Share	637	0.370	0.260	0.354	0.000	0.000	0.730	0.880
Panel B: AI Variables									
2025	AI Invest (dummy)	723	0.585	1.000	0.493	0.000	0.000	1.000	1.000
2025	AI Investment	691	$5,762	$2	$115,147	$0	$0	$75	$300
2025	Reported Δ LP	699	0.018	0.000	0.032	0.000	0.000	0.030	0.075
2025	Implied Δ LP	678	0.006	0.000	0.022	0.000	0.000	0.000	0.030
2025	Implied Δ TFP	657	0.005	0.000	0.020	0.000	0.000	0.000	0.030
2026	AI Invest (dummy)	700	0.854	1.000	0.353	0.000	1.000	1.000	1.000
2026	AI Investment	700	$8,294	$35	$153,507	$0	$2	$300	$750
2026	Reported Δ LP	672	0.030	0.030	0.034	0.000	0.000	0.030	0.075
2026	Implied Δ LP	644	0.018	0.000	0.033	0.000	0.000	0.030	0.060
2026	Implied Δ TFP	611	0.015	0.000	0.028	0.000	0.000	0.030	0.051

Notes: All dollar-denominated values are in thousands of USD. Large Firm is a dummy variable indicating whether a firm has at least 500 employees. Reported ΔLP is change in labor productivity reported by CFOs on the survey; Implied ΔLP is the calculated change in labor productivity based on revenue growth reported by CFOs on the survey; Implied TFP is the residual change in total factor productivity (i.e., change in labor productivity not captured by change in capital deepening, which itself is measured as the change in AI-attributed K/L). N indicates the number of CFOs who responded to a given question or an item about which we have sufficient information to include in our analysis. For example, we have employment data, and therefore can determine firm size, for 722 companies. Of these 722, 22% are defined as large (500 or more employees).

Figure 1: AI Investment and Reasons for Not Investing

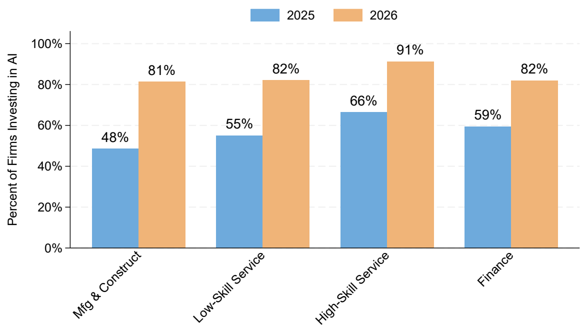

(a) Percent of Firms Investing in AI

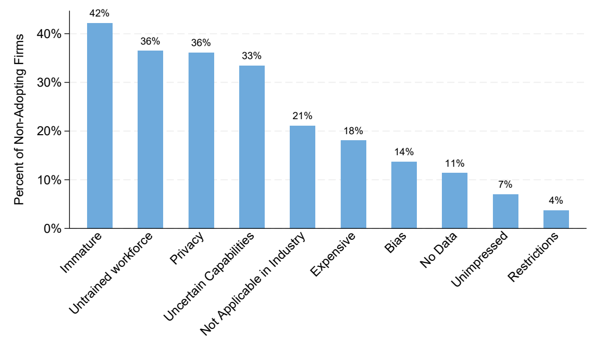

(b) Reasons for Not Investing

Notes: Panel A shows the extensive margin percentage of firms that have invested in 2025 (or expect to invest 2026) in AI technology, which in 2025 (2026) is determined by a Yes answer to a question asking whether the firm has invested in AI (a positive dollar amount for expected to AI investment in 2026). Mnfg&Constr includes "construction", "manufacturing", and "mining and utilities"; high-skill services include "educational services", "health care and social assistance", "information", "professional and business services", and "real estate and rental and leasing"; low-skill services include "leisure and hospitality", "retail and wholesale trade", "transportation and warehousing", and "other services except government". High-skill services are selected as those services for which the share of employees (from the 2024 CPS) with at least bachelor's degrees is greater than the cross-sector average. Low-skill services are the remaining services with lower-than-average bachelor degree shares. For the subsample of companies that indicated that they had not invested in AI in 2025, Panel B presents the percentage of respondents who selected a given reason for not investing in AI.

large firms expecting to spend over $1,000,000 on AI in 2026; in contrast, more than half of small companies plan to spend $20,000 or less on AI. In 2026, the median large (small) firm expects to spend $300,000 ($12,500). Mean spending levels (not reported) are considerably higher due to a small number of firms reporting very large expenditures. Although larger firms spend more on AI in absolute terms, AI investment intensity—measured per employee or relative to capital expenditures—is somewhat higher among smaller firms (Appendix Figures A6 and A7).

One of the supplemental surveys (which gathered 183 valid responses) asked about the *type* of AI spending (see Figure 3). In 2026, large (small) companies are expected to allocate on average 55% (64%) of AI spending on operations (AI subscriptions, services, and training), 31% (24%) on internal development of customized systems, and 13% (7%) on hardware, GPUs, and servers.

Figure 4 summarizes the factors that motivate companies to invest in AI. Responses to this question are on a scale from 0 (the factor was not a motivation to invest in AI) to 4 (an extremely important motivation). Many of the responses range from 2 (moderate motivation) to 3 (very large motivation). Two of the top motivations reflect a desire to improve productivity: Improve production efficiency (average rating of 2.9) and improve labor productivity (2.7). The third most important motive is to enhance decision-making and management (2.7). At the other end of the spectrum, while still moderately important, two of the least important motives to invest in AI are to achieve cost savings: reducing labor costs (2.0) and reducing non-labor costs (1.8). Overall, therefore, in the near-term, the goal of investing in AI is not to reduce workforce or costs but rather to improve productivity, the latter of which is a central focus of our paper.

Figure 2: AI Investment by Firm Size

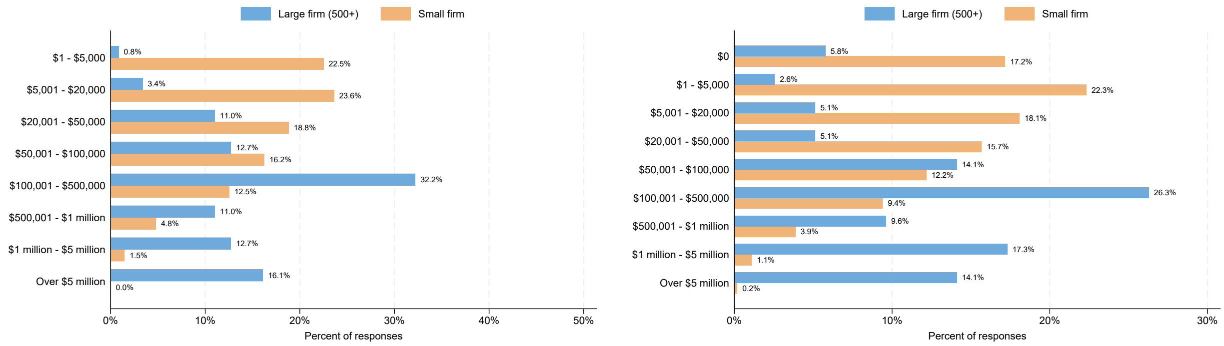

(a) 2025 (b) 2026

Notes: This figure depicts the amount of realized (expected) investment in AI technologies in 2025 in the left panel (2026 in right panel) for both large (blue, upper bars) and small (orange) firms. Large firms are defined as having more than 500 employees.

Figure 3: AI Investment by Investment Category

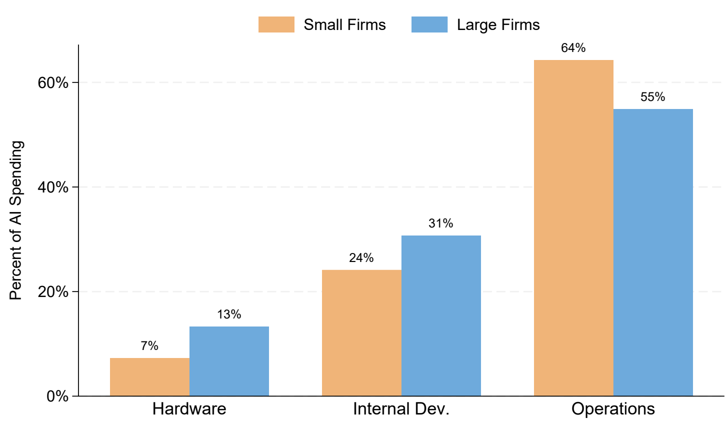

Notes: This figure shows the composition of expected 2026 AI expenditure by type and firm size (large firms have at least 500 employees). The question was drawn from a supplemental survey question on the 2026Q1 CFO Survey; the question produced 183 2026Q1 responses from CFOs who also participated in the main 2025Q4 sample we study in this paper. Operational expenses include AI subscriptions, services, and training; hardware includes servers, GPUs, and devices; internal development includes developing or customizing proprietary AI systems. Allocations do not necessarily sum to 100% due to an “Other” category that is not shown.

Figure 4: Motivations for AI Investment

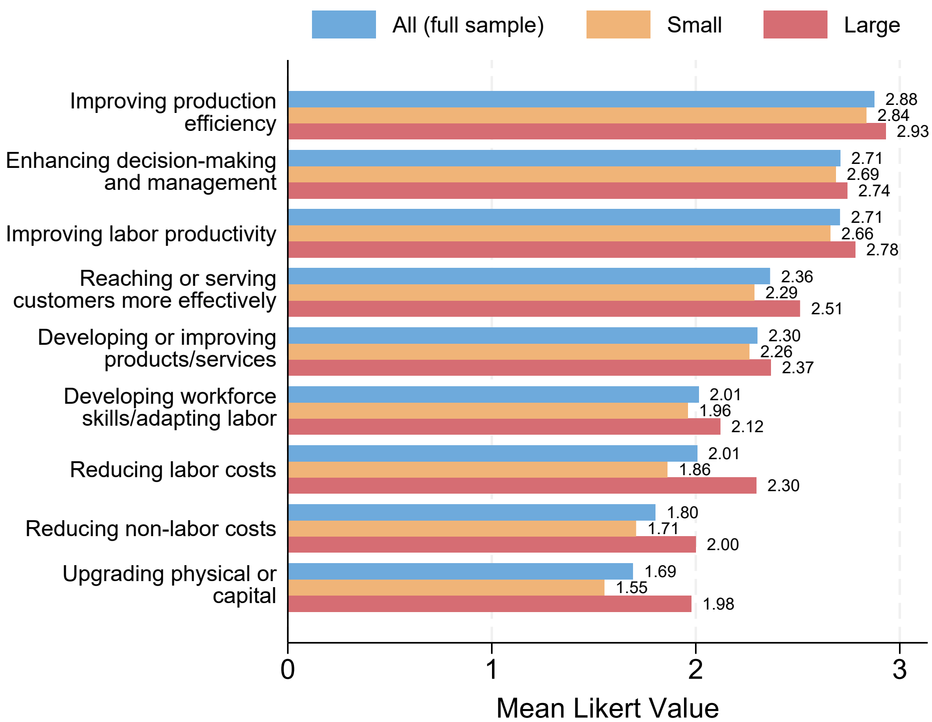

Notes: This figure depicts the motivations CFOs state for why their firm invested in AI. The results are broken down by firm size, with large firms having at least 500 employees. The Mean Likert Value for each category is calculated by averaging the following importance in choosing to invest in AI:'Not' as 0, 'Slightly' as 1, 'Moderately' as 2, 'Very' as 3, and 'Extremely' as 4.

The previous figure summarizes motivations to spend company money in 2025 on AI resources. On a limited basis, some free AI services are available (e.g., ChatGPT, Gemini); therefore, in Figure 5 we explore the realized (2025) and expected (2026) benefits of *using* AI – all responding companies are asked this question, regardless of whether they explicitly invest dollars in AI. Several patterns stand out. First, both small and large companies generally expect larger effects from AI in 2026 than they experienced in 2026. Second, the primary outcomes from using AI relate to labor productivity, quality of output, decision speed, and shifting employee focus to tasks with higher value-add. For example, the second cluster in the figure indicates that, due to the use of AI, large firms expect output per worker to be 3.06% higher in 2026 than it otherwise would have been. Third, at the bottom of the figure, there is little evidence that AI has had, or in the near-term will have, large effects on the total number of employees or on costs. One moderate exception to this statement is that large firms expect to reduce headcount by 0.7% in 2026 due to using AI.

4 AI and Productivity

This section examines the labor productivity effects of AI, their sources, heterogeneity across firms, and the mechanisms underlying both realized and expected productivity gains.

4.1 Reported vs. Implied Labor Productivity Effects

To estimate firm-level labor productivity effects from AI and to understand differences between perceived and realized returns to AI investment, we construct two complementary measures of productivity from our survey.

First, the survey asks firms to report the percentage change in labor productivity (output per worker) attributable to their use of AI over the past year (2025) and expected over the next year (2026). Responses are provided in ranges, which we convert to numeric midpoints; "no change" responses are coded as zero, while "not sure" responses are treated as missing. We refer to this measure as perceived, or *reported*, AI-related labor productivity growth, denoted by $\Delta \ln LP_{it}^{AI,CFO}$.

Second, 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.

For 2025, the mean reported labor productivity growth attributable to AI is 1.8% (1.7%

Figure 5: Effects of AI Usage

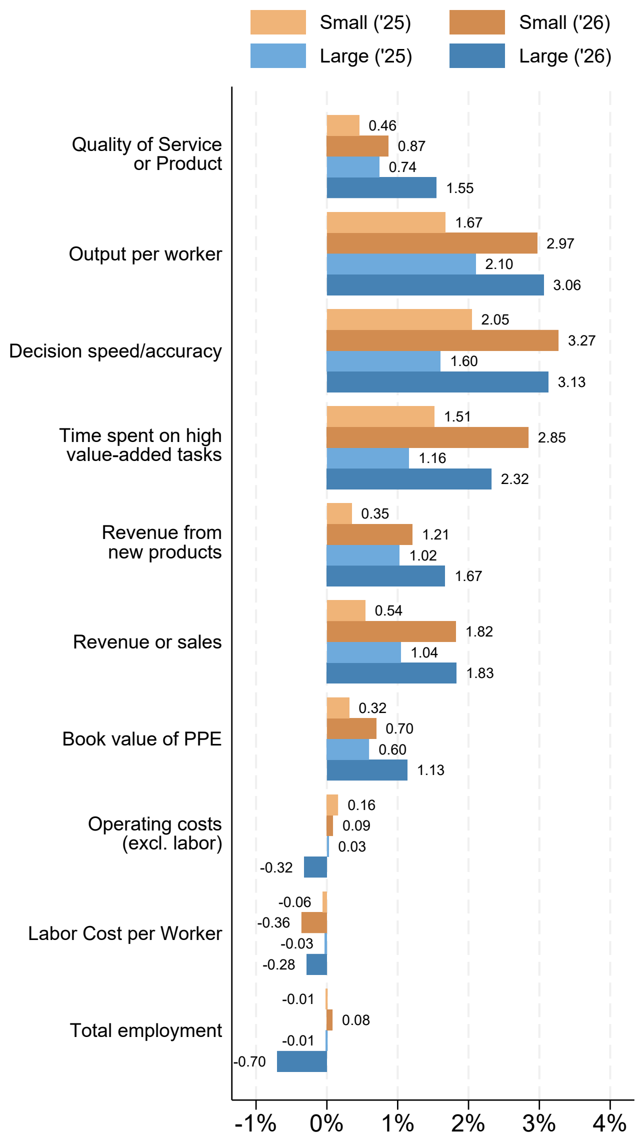

Notes: This figure depicts the effects of AI usage on company performance and cost metrics. The responses are categorical, with choices being "increased significantly, more than 10%", "increased moderately, 5.1 to 10%", "increased somewhat, 1 to 5%", "little to no change", "decreased somewhat, -1 to -5%", "decreased moderately, -5.1 to -10%", "decreased significantly, more than -10%"). This chart uses the midpoint of each range, and $+/-5\%$ above/below the highest/lowest option to calculate numeric averages across respondents.

Figure 6: Reported vs. Implied Productivity Effects of AI across Sectors

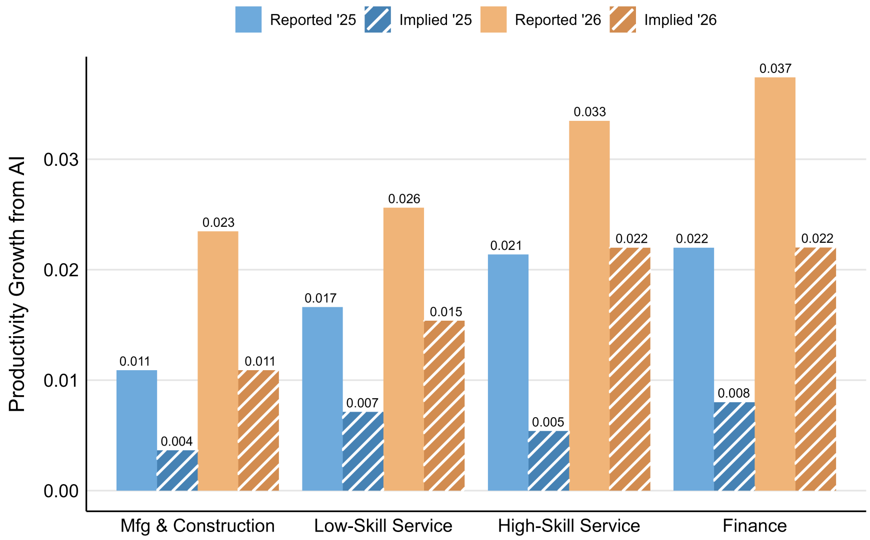

Notes: Bars show mean reported (light-colored bars), $\Delta \ln LP_{it}^{AI,\,CFO}$, and implied (dark-shaded bars), $\Delta \ln(Y/L)_{it}^{AI}$ labor productivity growth attributable to AI for 2025 and 2026 across sectors.

weighted[6]), while implied labor productivity growth is substantially smaller at 0.6% (0.6% weighted). For 2026, productivity gains from AI are expected to grow considerably: mean reported labor productivity growth rises to 3.0% (3.0% weighted), while the implied measure reaches 1.8% (1.9% weighted).

Figure 6 shows average reported (light-colored bars) and implied (dark-shaded bars) AI-related productivity growth for 2025 and 2026 across sectors. Two main patterns emerge.

First, firms consistently report larger productivity gains from AI than those implied by contemporaneous changes in revenue and employment; that is, $\Delta \ln LP_{it}^{AI,CFO} > \Delta \ln(Y/L)_{it}^{AI}$. This wedge likely reflects delayed output realization and quality improvements that are not yet captured in measured revenues. More broadly, firms' conceptual notion of productivity appears to extend beyond mechanical revenue-per-worker calculations, encompassing improvements in workflows, task efficiency, and organizational capacity whose revenue effects materialize only gradually. 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.[7] It is interesting to note that the 2025 reported productivity for all four sectors approximately equals the 2026 implied number for the same sector. This suggests a possible resolution to the productivity paradox, namely that productivity gains due to 2025 investment in AI will increase sales revenue the following year–at least according to CFO expectations.

Appendix Figure A8 provides a sensitivity check to the mapping of categorical survey responses

[6]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.

[7]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."

into continuous outcomes. Instead of using midpoint values, these checks assign lower and upper bounds to each response range. The results are similar, confirming that the gap between reported and implied productivity effects is not driven by the coding choice.

The second pattern from Figure 6 is that even implied productivity effects are positive and exhibit meaningful sectoral heterogeneity. In 2025, firms in high-skill services experience the largest overall 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 the largest anticipated increases again concentrated in high-skill services and finance, exceeding 2%. While revenue-based productivity gains are positive, they remain substantially smaller than those observed during the late-1990s productivity resurgence from IT, consistent with AI still being in an early phase of technological deployment (see Appendix Figure A4 for a comparison of the IT and AI technology timelines).

The numbers in Figure 6 report average effects across all companies, including those that do not directly invest in AI.[8] Appendix Figure A9 restricts the sample to firms with positive AI investment. Conditional on investment in 2025, firms experience higher realized returns to AI investment, with implied labor productivity gains ranging from 0.8% to 1.5%. Conditional on investing in 2025 and/or 2026, substantially larger gains are expected, ranging from 1.8% in manufacturing to 3.3% in finance.

Finally, Appendix Figure A10 reports results separately by firm size. While splitting the sample by both sector and size leads to smaller cell sizes, the core findings remain unchanged across firm types. Companies consistently report productivity gains that exceed those implied by revenue and employment changes, realize positive but modest gains in 2025, and expect these gains to increase markedly in 2026.

4.2 Decomposing Labor Productivity Growth from AI

Part of the labor productivity gains attributed to AI may reflect capital deepening—the fact 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 firm surveys, we measure Y using firm 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 decomposition.

[8]These firms still answer questions about AI-related outcomes, as they may benefit from free or embedded AI services.

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:

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

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

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

we construct the capital-deepening contribution from firm-reported 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, while the capital share α is calibrated from the literature.[9] Changes in capital are proxied using reported AI-attributed changes in the book value of property, plant, and equipment (PP&E) and intangible assets.

The second term in 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.[10]

Figure 7 decomposes 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 A4 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 contribution of capital deepening for the average firm is not surprising. First, most companies remain in the early stages of broad-based AI adoption: although many firms report positive AI investment, these expenditures are still small relative to firm size (see

[9] We set $\alpha = 0.40$ for manufacturing and construction, $\alpha = 0.20$ for low-skill services, and $\alpha = 0.30$ for high-skill services and finance. These values are consistent with sectoral capital income shares used in standard growth-accounting studies and official productivity accounts (e.g., BLS industry productivity statistics and EU KLEMS) and reflect well-documented differences in capital intensity across sectors.

[10] 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.

Figure 7: Decomposing Labor Productivity Growth from AI

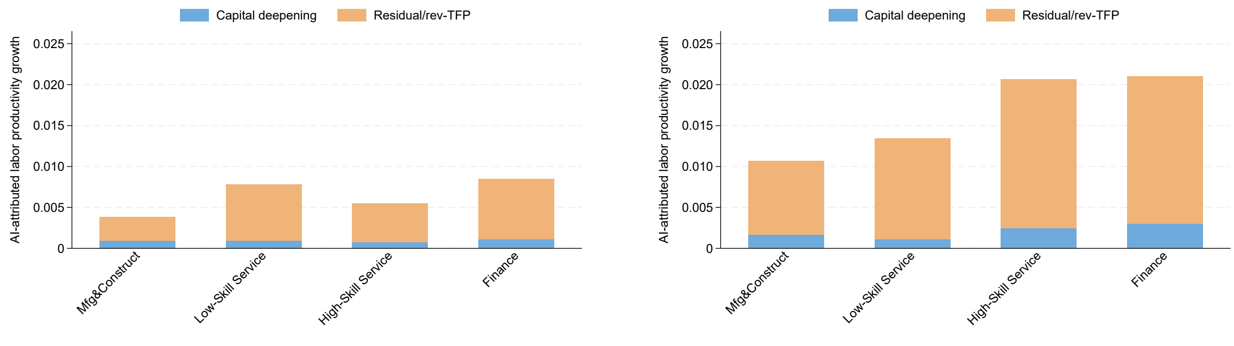

(a) 2025 (b) 2026

Notes: Bars show the decomposition of mean implied AI-attributed labor productivity growth, $\Delta \ln(Y/L)_{it}^{AI}$, across sectors into capital deepening and the residual (revenue-based TFP), as defined in Equation (4). Panel (a) decomposes productivity gains for 2025, while panel (b) decomposes expected gains for 2026.

Appendix Figures A6 and A7). Second, a substantial share of AI spending—especially among small firms—takes the form of operating expenses rather than capitalized investment (Figure 3). Expenditures on subscriptions to AI tools (e.g., large language models), cloud computing, or SaaS typically support activities such as coding, drafting, analysis, and decision-making, but are expensed rather than recorded as capital. As a result, a significant portion of AI's productivity impact operates through channels not captured by capital deepening.

It is instructive to compare these findings to the previous wave of general-purpose technologies associated with the information technology (IT) boom of the 1990s. During that period, capital deepening was a central driver of productivity growth: firms across the economy invested heavily in IT hardware and software, and aggregate labor productivity gains reflected both rapid IT capital accumulation and TFP growth in roughly comparable measure (Jorgenson et al., 2008). More recently, Rubinton and Patro (2026) show that aggregate AI-related investment categories have already made a sizable contribution to aggregate GDP growth in 2025, with magnitudes comparable to—or even exceeding—the contribution of IT components during the dot-com boom. However, much of the current AI investment surge is highly concentrated in a small group of large technology firms that build and operate data centers and cloud infrastructure. As a result, our firm-level evidence suggests that, for the *average firm in the economy*, current AI-related productivity gains are only modestly driven by capital deepening. Many firms can realize productivity improvements without undertaking large firm-specific capital investments, instead accessing AI capabilities through cloud services, subscriptions to large language models, and other rented digital infrastructure. This shift—from on-premise servers and firm-owned IT capital toward centralized cloud computing and AI services—allows productivity gains to diffuse broadly across firms without corresponding capital accumulation on their balance sheets.

4.3 Extensive and Intensive Margins of AI Investment

We examine more closely the quantitative effects of AI investment—along both the extensive and intensive margins—on reported and implied labor productivity growth and on the revenue-based productivity residual (see Table 2).[11]

Panel A presents results for the extensive margin of AI investment. Consistent with the discussion above, there is a gap between reported and implied productivity gains from AI. Relative to firms that do not invest in AI, investing firms report AI-related labor productivity growth that is 2.4 percentage points higher in 2025 and expect gains that are 3.3 percentage points higher in 2026.[12] The corresponding implied productivity gains are smaller, at 1.0 and 1.8 percentage points in 2025 and 2026, respectively. Most of these gains are accounted for by increases in implied revenue-based TFP (the productivity residual); for example, the TFP coefficient in column (3) is nearly as large as the column (2) coefficient, and likewise for columns (5) and (6).

The next panels examine the intensive margin of AI investment, including zero-investment firms in Panel B and restricting to firms with positive AI investment in Panel C. The estimated semi-elasticities are modest, even for 2026. In Panel B columns (1) and (2), a 10% increase in AI investment is associated with a 0.02 percentage-point increase in reported labor productivity growth and a 0.01 percentage-point increase in implied productivity growth in 2025, with effects that are roughly twice as large in 2026. To aid interpretation, a one-standard deviation increase in log AI investment is associated with a 1.1 percentage-point increase in reported labor productivity growth (5.85×0.002) and a 0.5 percentage-point increase in implied labor productivity growth. Restricting the sample to firms with positive AI investment (Panel C) shows that these semi-elasticities are statistically insignificant for 2025, but remain statistically significant and economically larger for expected productivity growth in 2026.

4.4 Productivity Gains from AI: Mechanisms

Finally, to better understand the channels through which companies experience revenue gains from AI, we leverage the survey questions on companies' motivations for AI investment (see Figure 4). Companies are asked to rate the importance of the following motivations for their AI expenditures: improving production efficiency (e.g., faster processes, automation or optimization of internal operations, logistics, or maintenance); reducing labor costs; reducing non-labor costs; enhancing decision-making and management (e.g., data analysis, forecasting, workflow or HR optimization); developing or improving products and services (e.g., new or higher-quality offerings, personalization, testing, faster R&D cycles); reaching or serving customers more effectively (e.g., marketing, customer interaction, after-sales support); upgrading physical or digital capital (e.g., hardware, data infrastructure, or cloud systems); developing workforce skills or adapting labor (e.g., training, hiring, or reorganizing teams for AI use); and other motivations. We encode each motivation as

[11] Appendix Tables A1 and A2 report similar results using weighted regressions.

[12] Recall that the survey question is asked of all firms, irrespective of whether they invest in AI – even firms that do not explicitly invest company resources in AI solutions mayexperience changes in productivity.

Table 2: Revenue-Productivity Growth and Extensive and Intensive Margins of AI Investment

	(1)	(2)	(3)	(4)	(5)	(6)
		2025			2026	
	Reported ΔLP	Implied ΔLP	Implied ΔTFP	Reported ΔLP	Implied ΔLP	Implied ΔTFP
Panel A: AI invest dummy						
AI Adopt (‘25)	0.024***	0.010***	0.008***			
	(0.002)	(0.002)	(0.001)			
AI Adopt (‘25/‘26)				0.033***	0.018***	0.014***
				(0.002)	(0.002)	(0.002)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.15	0.05	0.04	0.15	0.06	0.06
Observations	684	663	649	667	639	611
Panel B: AI investment (all firms, unconditional)						
Log AI Inv (‘25)	0.002***	0.001***	0.001***			
	(0.000)	(0.000)	(0.000)			
Log AI inv (‘25/‘26)				0.003***	0.002***	0.002***
				(0.000)	(0.000)	(0.000)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.15	0.06	0.05	0.17	0.12	0.10
Observations	660	642	630	625	601	575
Panel C: AI investment (only positives, conditional)						
Log AI Inv (‘25)	0.000	0.001	0.001			
	(0.001)	(0.001)	(0.001)			
Log AI inv (‘25/‘26)				0.002***	0.004***	0.004***
				(0.001)	(0.001)	(0.001)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.02	0.02	0.02	0.05	0.13	0.10
Observations	364	334	334	528	504	478

Notes: The table reports firm-level regressions of changes in reported labor productivity, implied revenue-based labor productivity (Δ LP), and implied revenue-based total factor productivity (Δ TFP) for 2025 and 2026 on extensive and intensive measures of AI investment. Growth in reported labor productivity and implied revenue-based labor productivity are defined in Section 4.1, and revenue-based TFP (the productivity residual) is defined in equation 4. Panel A studies the extensive margin of AI investment. *AI Adopt'25* equals one if the firm reports any AI investment in 2025, and *AI Adopt'25/'26* equals one if the firm reports any AI investment in either 2025 or 2026. Panel B studies the intensive margin for all firms. *Log AI Inv'25* is the logarithm of AI investment reported in 2025 (plus one), and *Log AI Inv'25/'26* is the logarithm of total AI investment reported over 2025 and 2026 (plus one). Panel C restricts the sample to firms with positive AI investment. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 Table 3, we perform regressions of reported and implied labor productivity growth and the revenue-based productivity residual for 2025 and 2026 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.

Table 3: Labor Productivity Gains from AI: Mechanisms

	(1)	(2) 2025	(3)	(4)	(5) 2026	(6)
	Reported ΔLP	Implied ΔLP	Implied ΔTFP	Reported ΔLP	Implied ΔLP	Implied ΔTFP
Operational Efficiency Channel						
Production Efficiency	0.011**	-0.002	-0.002	0.013***	0.010**	0.009**
	(0.004)	(0.003)	(0.003)	(0.005)	(0.005)	(0.004)
Reduce Labor Costs	0.004	0.002	0.000	0.004	0.004	-0.000
	(0.005)	(0.004)	(0.003)	(0.005)	(0.005)	(0.004)
Reduce Other Costs	-0.011**	-0.002	-0.003	-0.009**	-0.002	-0.001
	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Decision Making/Mgmt	0.011**	-0.004	-0.005	0.001	-0.003	-0.004
	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)
Innovation & Demand Channel						
Product Development/Improvement	0.009**	0.008**	0.006**	0.011***	0.012***	0.010***
	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)
Reach/Serve Customers	0.000	0.008**	0.008**	0.007*	0.011**	0.011***
	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)
Factor Upgrading						
Upgrade Capital	-0.003	-0.000	0.001	-0.008	0.005	0.007
	(0.005)	(0.004)	(0.003)	(0.005)	(0.005)	(0.005)
Workforce Development	-0.002	-0.002	-0.001	0.003	-0.006	-0.004
	(0.005)	(0.003)	(0.003)	(0.004)	(0.005)	(0.004)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.09	0.06	0.07	0.10	0.13	0.13
Observations	375	356	342	387	369	348

Notes: This table reports firm-level regressions of changes in reported labor productivity, 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results in Table 3 indicate the channels through which AI investment is associated with revenue-based productivity gains. Across specifications and productivity measures, 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.[13] Corroborating this finding, Appendix Figure A13 shows that when firms

[13] Appendix Tables A3 and A4 show similar patterns using weighted regressions. However, the importance of reach-

are asked in a separate survey question which business tasks AI is most helpful for, they most frequently cite marketing and product development.

These findings suggest that for the firms in our sample, the revenue productivity gains associated with AI investment operate largely through innovation, product-market and demand-side channels. This interpretation is consistent with a growing literature that emphasizes the role of general-purpose technologies, and AI in particular, in fostering innovation and the development of 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, motivations related to cost reduction, workforce development, capital upgrading, and enhanced decision-making or management do not display robust positive associations with gains in revenue per labor, often exhibiting statistically imprecise and often negative coefficients. These patterns are consistent with the view that such channels operate through deeper organizational and human-capital adjustments—such as workforce training, task reallocation, and complementary capital investments—that involve substantial adjustment costs and learning frictions. As a result, their productivity effects may materialize only gradually and are unlikely to be fully reflected in short-run revenue-based productivity measures. This interpretation aligns with a large literature emphasizing that the gains from new general-purpose technologies typically emerge with a delay, as firms upgrade skills, reorganize production, and accumulate complementary inputs (Greenwood and Yorukoglu, 1997; Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002).

Overall, the evidence underscores the importance of innovation and market-expansion channels in shaping the short-run productivity payoffs to AI adoption, while leaving open the possibility that cost-based and organizational channels may operate over longer horizons not fully captured in our sample.

5 AI and the Labor Force

Technological change alters production processes and, in doing so, reshapes the set of tasks performed within firms and the roles required to carry them out. 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).

Artificial intelligence represents a new and potentially distinct technological shift. However, because AI adoption remains at an early stage, its implications for the labor force are not yet well understood. In this section, we provide early evidence on how firms report that AI is already

ing or serving customers weakens when applying firm-size weights, suggesting that these gains are driven primarily by smaller firms.

affecting employment, and how they expect AI to affect it in the future. We begin with employment effects, then turn to the composition, tasks, and roles of the workforce.

5.1 Employment Effects of AI

We examine AI-driven employment changes using two complementary sources of 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 in Appendix Section A5). Second, firms report the expected and realized changes in total employment attributed to AI, based on the survey questions used in our previous analysis (survey question 8).

A first finding is that, at the extensive margin, most firms do not anticipate job replacement from AI. Based on the CFO's open-ended responses to a question asking which roles and responsibilities will be replaced by AI, 50% of firms 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. AI-job effects vary greatly conditional on whether a company has invested in AI. Among firms investing in AI, roughly half anticipate some role replacement; among non-investors, one-in-five firms report expected job replacement due to AI—suggesting that AI-related labor reallocation is not confined to those currently investing in AI.

These qualitative responses align with firms' reported numerical employment changes attributed to AI. Companies indicating role replacement in the open-ended responses also report lower employment growth due to AI in the quantitative question. Table 4 presents average AI-driven employment changes in 2025 and expected changes in 2026, separately for firms that (in the open-ended questions) anticipate some role replacement and those that do not. Firms expecting job replacement experience modest AI-related employment declines, whereas firms reporting no replacement show small positive employment effects. Although the magnitudes are modest, the sign and timing of these changes closely match the qualitative responses, providing reassurance about the internal consistency of the survey data.

Table 4: Employment Changes by AI Role Replacement Status

	Firms reporting no replacement of roles	Firms reporting some replacement of roles
Employment change in 2025, %	0.08	-0.14
Employment change in 2026, %	0.22	-0.55

Notes: Table reports average employment changes due to AI in 2025 and expected changes in 2026. Firms are grouped by whether they anticipate AI replacing any roles or responsibilities ("some replacement") or not ("no replacement").

To further unpack this heterogeneity in employment effects, we examine AI-driven employment changes across firm size. Figure 8 reveals a clear negative gradient with respect to firm size: larger firms are substantially more likely to report employment reductions due to AI, while smaller firms are more likely to report employment gains. This gradient is already visible in 2025 and becomes more pronounced in expectations for 2026. Importantly, the relation to company size is driven entirely by firms that invest in AI; non-investors show neither systematic employment declines nor a size gradient.[14]

The stronger employment reduction among larger companies is intuitive. Large firms employ more heterogeneous workforces with a wider variety of job occupations, and more workers performing a narrow set of tasks, increasing the likelihood that at least some jobs are vulnerable to automation. In addition, larger firms face higher aggregate labor costs and may therefore have stronger incentives to use AI explicitly to cut down employment.[15] Consistent with this interpretation, Figure 4 shows that larger firms are more likely to cite labor cost reduction as a motivation for AI adoption relative to small firms.

Figure 8: AI-Driven Employment Changes by Firm Size and Investment Status

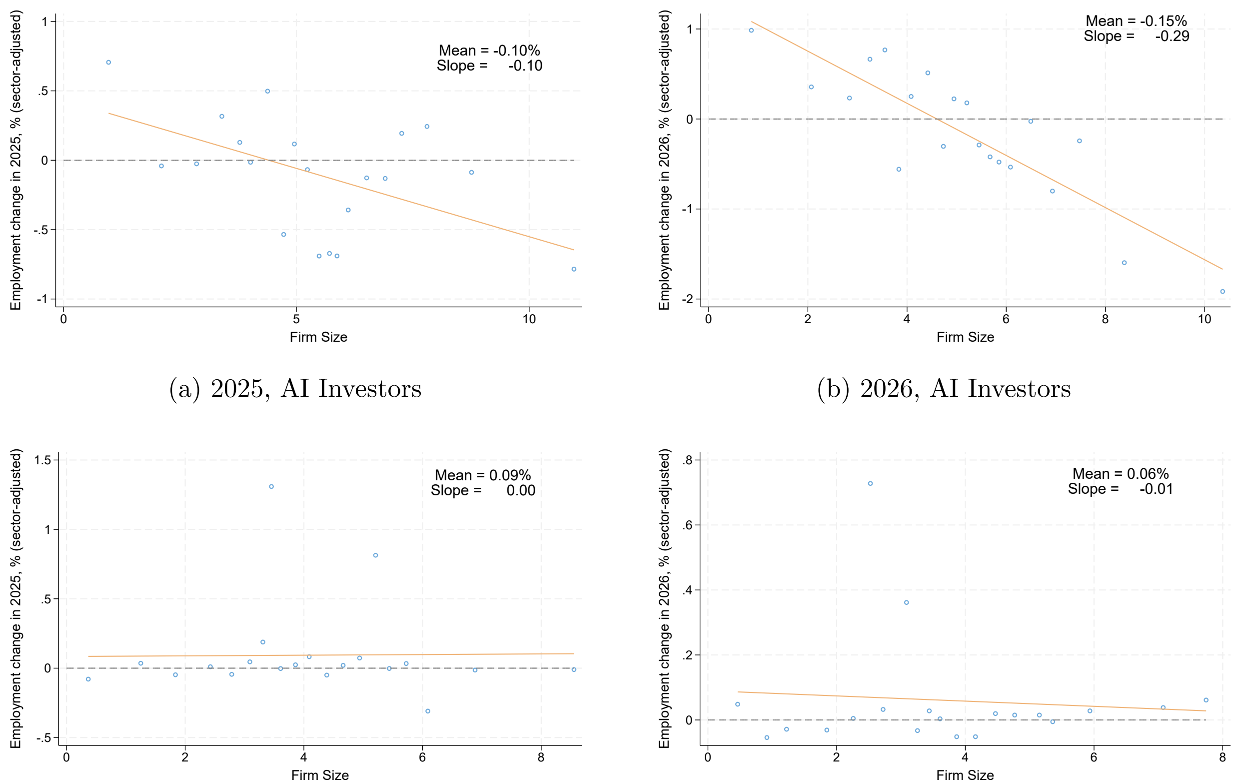

(a) 2025, AI Investors

(b) 2026, AI Investors

(c) 2025, Non-Investors

(d) 2026, Non-Investors

Notes: Binscatter plots of employment changes due to AI (realized for 2025, expected for 2026) firm size as measured by the firms log number of employees in 2024. The top panels show firms with positive AI investment; the bottom panels show firms with no AI investment in the respective period. Binscatters are residualized for sector fixed effects. Red lines indicate linear regression fits.

[14] A supplemental module in the 2026Q1 CFO survey asks about AI-driven employment changes expected for 2026 and 2028. Differences in sample and question framing prevent direct integration with our main analysis. Nonetheless, responses suggest that overall employment effects are expected to remain small through 2028, with approximately half of CFOs saying that they expect no job loss at their companies in either 2026 and 2028 due to AI usage. Among companies that expect job loss (gain) due to AI in 2026, they expect a slightly larger magnitude job loss (gain) to occur in 2028.

[15] A body of literature shows that larger firms are more likely to engage in cost-cutting process innovation (Klepper, 1996; Cohen and Klepper, 1996; Tham et al., 2025).

In aggregate, combining these heterogeneous responses implies essentially no AI-driven employment effect in 2025 (both weighted and unweighted means), but a small negative effect in 2026 of 0.1% for sample firms and 0.4% for the aggregate economy (weighted). To translate these estimates into numeric headcounts and highlight sectoral heterogeneity, Table 5 reports weighted average AI-driven employment changes by sector, along with implied employment counts based on total sectoral employment from the BLS Establishment Survey. The implied aggregate effect for 2026 is a decline of about 502,000 workers, with half of these losses coming from firms in high-skill services. These estimates are based on responses from *incumbent* firms; if new technologies also lead to entry of new firms and reallocation of labor toward younger firms not captured in current surveys, the net employment effects could be smaller.

Overall, our analysis suggests that AI is expected to displace some workers, but the aggregate effects remain very small and will likely be hard to detect in aggregate statistics. The employment consequences of AI, at least at this early stage, are concentrated among larger firms and AI adopters, and vary somewhat across sectors, but do not yet reflect a major labor market disruption.

Table 5: Sectoral and Aggregate AI-Driven Employment Changes, 2026

	Mfg& Construct	Low-Skill Services	High-Skill Services	Finance
Δ Empl, %	-.259%	-.12%	-.514%	-1.204%
	(-.559%, .041%)	(-.406%, .167%)	(-.84%, -.187%)	(-1.602%, -.806%)
Δ Empl, count (thousands)	-58	-62	-272	-111
	(-124, 9)	(-209, 86)	(-444, -99)	(-148, -74)

	Count	% of Labor Force
Aggregate Δ Empl	−502	−0.370%
	(-925, -78)	(-0.70%, -0.01%)

Notes: The table reports expected employment changes due to AI in 2026 by sector and in the aggregate. Mean effects are weighted using representativeness and importance weights. Employment counts are reported in thousands and are based on the share of the labor force in the covered sectors. 95% confidence intervals are shown in parentheses.

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

The labor analysis above indicates that the overall impact of AI adoption on total workforce size has been minimal to date and is expected to remain limited through 2026. In this section, we examine whether AI will reshape the *composition* of employment within firms.

AI adoption may induce firms 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 remains an open empirical question.

To address this issue, we first examine firms' current and expected workforce composition in 2026 and 2028. 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. These responses allow us to construct changes in the distribution of employment across occupation types relative to 2025 and relate them to AI investment.

Figure 9 summarizes expected changes in the composition of the workforce in 2026 (top graphs) and 2028 (bottom graphs). On average, CFOs expect there to be a 0.76% reduction in 2026 in the proportion of their workforce doing routine clerical work, and a 2.19% reduction by 2028. This will be partly offset by a 0.62% increase in skilled technical workers in 2026, and 1.35% by 2028.[16] Appendix Tables A5 and A6 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.

Given the earlier evidence (Figure 8) of a strong firm-size gradient in AI-driven employment adjustments—with reductions in total employment concentrated among larger companies—Figure 9 examines changes in workforce composition by firm size, measured by (log) employment in 2024. Relative to small companies, large companies are expected to shed routine clerical employees more, while they hold steady creative and technical employment; while small firms are more likely to hold steady in routine jobs and increase skilled technical jobs. This suggests measurable labor reallocation *across firms* attributable to workforce composition changes. This complements evidence presented above on overall employment responses that showed that some AI-driven labor adjustments are expected to also occur within firms.

Second, we analyze firms' responses to open-ended survey questions describing the roles and responsibilities expected to be *replaced* or *enhanced* by AI (Questions 10 and 11 in Appendix Section A5). These responses provide more granular evidence on the specific tasks and occupations most exposed to AI tools—either negatively, through potential displacement, or positively, through task augmentation. To systematically interpret these answers, we first map the text of each replacement and enhancement response to occupation groups based on the BLS Standard Occupational Classification (SOC) categories.[17] For example, an enhancement response of "Fraud Group, Accounting, Finance, Customer Experience, Call Center" is mapped to several occupation groups, including *Business and Financial Operations*, *Legal*, and *Office and Administrative Support*. This procedure

[16] Workers not categorized as routine/clerical, creative, or skilled-technical are not presented in the graphs, so the changes do not net to 0%.

[17] A team of four research assistants coded the open-ended responses into the 23 major BLS SOC occupation groups. The research assistants were business graduate students who had work experience with AI and its effects on the workforce. Then, our research team reviewed and confirmed the RAs' work.

Figure 9: Expected Change in Types of Jobs Conditional on Firm Size

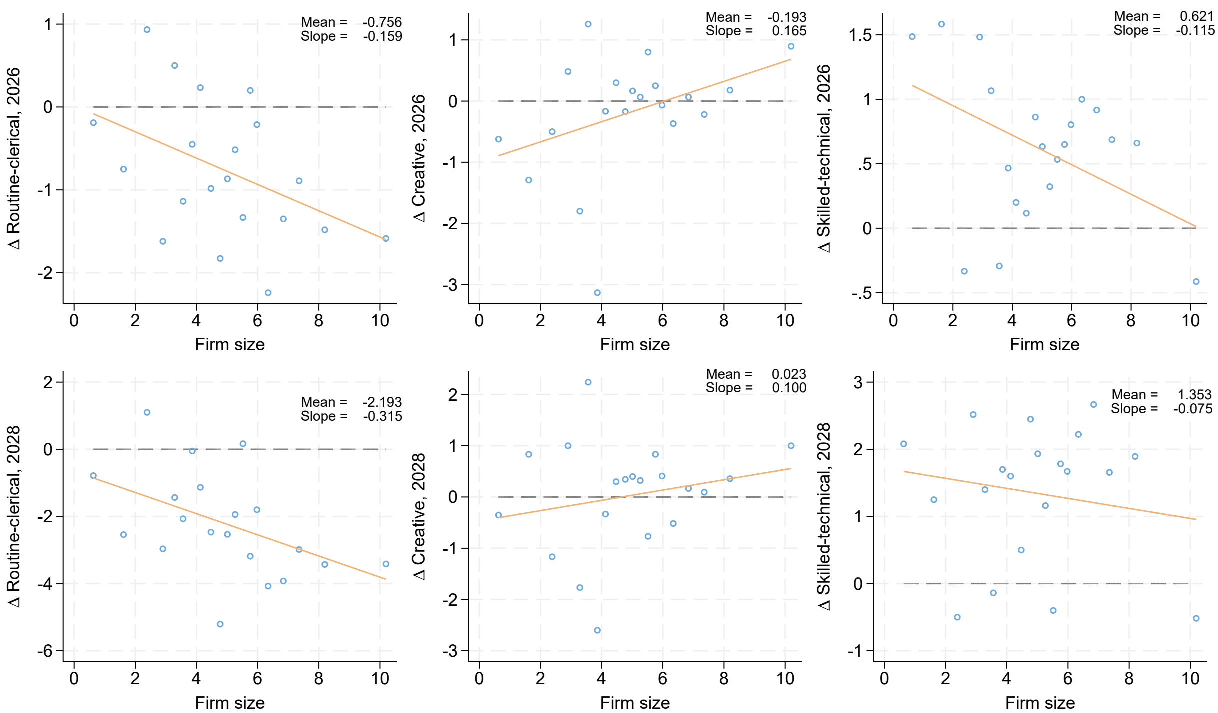

Notes: Figures show binscatter plots of expected changes in workforce composition in 2026 (top panels) and 2028 (bottom panels) as a function of firm size. Changes are measured relative to 2025. Firm size is defined as the log number of employees in 2024. Worker categories include routine/clerical roles (e.g., data entry, accounting), skilled technical roles (e.g., engineers, data analysts, scientists), and creative or managerial roles (e.g., design, strategy, leadership). Workers not falling into these categories are classified as “Other” and are not shown in the graphs; therefore, employment changes across the displayed groups do not sum to 0%.

produces a firm×occupation dataset in which each firm–occupation pair is tagged as either *replacement*, *enhancement*, or *no mention*. The resulting dataset contains 278 distinct firms (each mapped to multiple occupation codes) describing roles being replaced by AI and 536 describing roles whose tasks are enhanced.[18]

To summarize the balance between replacement and enhancement across sectors and occupations, we construct two simple measures. First, we count *replacement mentions* (RM), defined as instances in which firms report AI replacing roles or tasks within a given occupation. Second, we count *enhancement mentions* (EM), defined as cases in which AI is described as complementing or augmenting workers in that occupation. Using these counts, we construct a *Negative Exposure Index* (NEI),

$$\text{Negative Exposure Index (NEI)} = \frac{\text{Replacement Mentions (RM)}}{\text{Enhancement Mentions (EM)}}.$$

An index greater than one indicates that AI is more frequently described as replacing rather than enhancing work in that occupation or sector. By construction, the NEI abstracts from employment shares and instead captures the *direction* of AI exposure conditional on being mentioned in survey responses.

Table 6: Negative Exposure Index by Occupation Group

Occupation	Job Titles	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

Notes: The table 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 A5). The second column lists examples of more detailed occupations within each SOC occupation group. The Negative Exposure Index (NEI) is defined as the ratio of replacement mentions to enhancement mentions for each occupation group, aggregated across all firms. An index greater than one indicates that AI is more frequently described as replacing rather than enhancing tasks in that occupation group, while values below one indicate that enhancement mentions dominate. The table reports only occupation groups with at least 20 total mentions of either replacement or enhancement.

[18] In cases where a firm reports both replacement and enhancement for the same occupation (for example, reporting outsourced accounting under replacement while internal accounting under enhancement), we classify the occupation as replacement. This approach aligns with the interpretation that replacement reflects declining demand for that occupation within the firm, whereas enhancement reflects complementarity between AI and the occupation's tasks.

Table 6 summarizes AI exposure across major occupation groups, restricting attention to groups with at least 20 total mentions of either replacement or enhancement. The table also lists representative sub-occupations within each SOC group. Several patterns emerge. *Office and Administrative Support* occupations exhibit clear negative exposure (NEI greater than one), indicating that AI is more frequently described as substituting for workers in these roles. This pattern is consistent with the automation of routine clerical activities such as data entry, transaction processing, and scheduling, and aligns with the earlier evidence in Figure 9, which showed declining shares of routine clerical employment within firms adopting AI. In contrast, *Business and Financial Operations* occupations exhibit roughly balanced replacement and enhancement mentions, suggesting substantial task reallocation within these jobs rather than uniform displacement. Other occupation groups—including sales, technical, and professional roles—are more frequently described as being enhanced by AI tools.[19]

Figure 10: Word Clouds of Tasks and Roles Enhanced or Replaced by AI

(a) Enhancement (b) Replacement

Notes: The figure displays word clouds summarizing roles and responsibilities described by firms as being enhanced (left) or replaced (right) by AI based on open-ended survey responses (Survey Questions 10–11; see Appendix Section A5). Text responses are processed as follows. We first manually identify fragments of each response corresponding to meaningful occupations or tasks. These fragments are decomposed into bigrams and clustered into semantically similar groups using agglomerative clustering based on cosine similarity of 384-dimensional dense vector embeddings generated by the `all-MiniLM-L6-v2` sentence transformer. Based on these clusters, we manually compile a list of 48 task categories frequently appearing in the responses and reclassify the original text into these categories. We then compute embeddings for the full text responses, their bigrams, and the 48 task categories and assign tasks when the cosine similarity exceeds 30% for full-text embeddings or 50% for bigram embeddings. This procedure yields 5,165 distinct firm–task observations in the enhancement category and 2,130 in the replacement category. Word sizes reflect the frequency with which tasks appear in the corresponding category.

To further illustrate the types of tasks affected by AI, Figure 10 presents word clouds constructed from the open-ended survey responses describing roles and responsibilities that firms expect AI to enhance or replace. The enhancement cloud reveals several clear "winners," with the largest and most frequent mentions concentrated in areas such as marketing, accounting, finance, and management, alongside supporting analytical roles such as data analysis and software-related tasks. These activities involve information processing, analysis, and decision support—areas where AI tools appear to complement workers by improving productivity and expanding capabilities. In contrast,

[19] Appendix Figure A13 provides additional survey evidence on the specific business tasks for which firms report AI tools to be most helpful.

the replacement cloud is more diffuse, with mentions spread across a wider range of tasks including manual tasks, administrative staff, data entry, customer services, routine tasks, and accounting support roles. This broader dispersion suggests that AI-driven substitution is not concentrated in a single dominant function but instead affects a variety of routine and operational activities. Overall, these patterns reinforce the occupation-level evidence above: AI appears most likely to substitute for routine clerical and operational tasks while complementing higher-skill analytical and decision-oriented work. Similar patterns emerge in large-scale analyses of real-world AI usage, such as the Anthropic Economic Index, which finds that most occupations exhibit more AI-assisted task augmentation than automation, while clerical and administrative roles show relatively higher rates of automation.[20]

Table 7 reports exposure patterns across sectors. The first column measures the overall prevalence of AI exposure, defined as the share of firms in a sector that mention either replacement or enhancement. The second column reports the Negative Exposure Index conditional on being exposed. Overall, AI is mentioned more frequently as enhancing jobs than replacing them: the NEI is below one for most sectors. This suggests that, at least at present, firms more often view AI as augmenting worker productivity than eliminating positions. At the same time, there is substantial heterogeneity across sectors, reflecting differences in occupational composition and task structure. Sectors such as real estate and rental and leasing, other services (except government), finance, leisure and hospitality, and information exhibit relatively higher NEI values, indicating that AI is more frequently discussed as substituting for workers in these sectors.

Taken together, these results suggest that while AI adoption has not yet led to meaningful changes in total employment, it is already beginning to reshape the allocation of tasks and occupations within firms. In particular, routine clerical and administrative activities appear more exposed to substitution, while analytical, technical, and managerial tasks are more often described as being complemented by AI. These patterns are visible both in the reported changes in workforce composition and in firms' open-ended descriptions of tasks and responsibilities affected by AI.

6 Conclusion

This paper provides new firm-level evidence on how artificial intelligence is affecting productivity and the workforce by leveraging timely survey data from corporate financial decision-makers. We show that AI adoption is already widespread and associated with measurable revenue-based labor productivity gains, and these gains operate primarily through innovation- and demand-oriented channels rather than capital deepening. At the same time, our evidence suggests that AI's near-term labor market effects are characterized less by aggregate job losses and more by shifts in tasks and occupational exposure across workers.

There are several caveats to interpreting our findings as predictions of the eventual impact of AI. First, our survey questions focus mostly on the very near-term effects of AI – realized (2025)

[20]See the Anthropic Economic Index job explorer: `https://www.anthropic.com/economic-index#job-explorer`.

Table 7: Firm Exposure and Negative Exposure Index by Sector

Sector	Share of Firms Exposed	NEI (RM/EM)
Information	0.704	0.929
Transportation and Warehousing	0.677	0.765
Professional and Business Services	0.636	0.731
Health Care and Social Assistance	0.625	0.688
Other Services Except Government	0.622	1.308
Construction	0.609	0.520
Leisure and Hospitality	0.609	1.000
Real Estate and Rental and Leasing	0.603	1.227
Manufacturing	0.596	0.580
Finance and Insurance	0.582	1.100
Retail and Wholesale Trade	0.575	0.758
Mining and Utilities	0.524	0.800
Educational Services	0.409	0.714

Notes: The table summarizes AI exposure across sectors 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 A5). The first column reports the share of firms within each sector that mention either replacement or enhancement due to AI. The second column reports the Negative Exposure Index (NEI), defined as the ratio of replacement mentions to enhancement mentions within a sector. An index greater than one indicates that AI is more frequently described as replacing rather than enhancing tasks in that sector, while values below one indicate that enhancement mentions dominate. Mentions are aggregated across all firms in the sector.

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. 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 we document internal consistency and validate the quality of the responses along several dimensions, survey data also 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 analysis focuses mainly on gains at incumbent firms and reallocation of resources among existing firms. Accordingly, a third 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. 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. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

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.

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., 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. (2017). Artificial intelligence and the modern productivity paradox: A clas of expectations and statistics.

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

Cohen, W. M. and Klepper, S. (1996). Firm size and the nature of innovation within industries: The case of process and product R&D. *The Review of Economics and Statistics*, 78(2):232–43.

Crane, L. D., Green, M., and Soto, P. (2025). Measuring AI uptake in the workplace. FEDS Notes, Board of Governors of the Federal Reserve System. February 5, 2025.

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. *None*, None(None):None.

Davis, S. J. and Haltiwanger, J. (1991). Wage dispersion between and within u.s. manufacturing plants, 1963-86. *Brookings Papers on Economic Activity*, 22(1991 Micr):115–200.

Fernald, J. and Ramnath, S. (2004). The acceleration in u.s. total factor productivity after 1995: The role of information technology. *Economic Perspectives*, 28(1):52–67.

Gennaioli, N., Ma, Y., and Shleifer, A. (2016). Expectations and investment. *NBER Macroeconomics Annual*, 30:379–431.

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.

Gordon, R. J. (2000). Does the 'new economy' measure up to the great inventions of the past? *Journal of Economic Perspectives*, 14(4):49–74.

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.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Klepper, S. (1996). Entry, exit, growth, and innovation over the product life cycle. *American Economic Review*, 86(3):562–83.

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.

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.

Tham, W. W., Baslandze, S., Sojli, E., and Liu, L. (2025). Foundational processes and growth. CEPR Discussion Paper 19858, CEPR Press, Paris & London.

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.

Appendix

A1 Additional Data Details

Most of the AI survey sample (603 observations) used in this paper are responses from financial executives (mostly CFOs) who are part of the panel of approximately 2500 executives who participate in The CFO Survey, conducted jointly each quarter since June 2020 by Duke University and the Federal Reserve Banks of Atlanta and Richmond. Prior to working with the Federal Reserve starting in June 2020, Duke University conducted the survey each quarter starting in Summer 1996. Many of the Duke CFO participants joined the Duke-Fed joint survey project in June 2020, and additional CFOs have been invited to join the sample in each quarter since. This section summarizes information on the full CFO panel and the responses from the 603 respondents who completed our AI questionnaire.

The full CFO panel contains approximately 2500 financial executives, and therefore the 603 respondents to the 2025Q4 AI survey reflect a 24% response rate. Of course, some CFOs may have left their positions or stopped responding to the survey, so 2500 is likely too large a denominator to measure the response rate among *active* CFOs. Approximately 1330 unique CFOs participated in at least one quarterly CFO survey during 2024-2025. Relative to this denominator, the AI project response rate is 45%.

Since we rely on firms' expectations in our examination of AI investment and its impact, we assess whether these expectations are accurate in key business areas. Figure A1 plots year-ahead forecasts of four key variables against the self-reported realizations that the same companies experience for these same four variables over the same timeframes. The four variables are sales revenue growth, price growth for the goods and services that companies produce, unit cost growth for the price of inputs that companies use, and wage growth. Each graph contains 56 points that represent the average for the variable in each of 14 industry sectors for each year 2022-2025 (e.g., one dot represents the average price change forecast vs. realization for the finance sector in 2022, another dot represents the same in 2023, etc.) Overall, the aggregate forecast for each of the four variables aligns well with its realization, lending credibility to the aggregate projections produced by The CFO Survey.

We also perform several analyses to determine the extent to which the full panel of CFOs is reflective of the U.S. economy. Figure A2 shows the industry breakdown of the CFO panel. The sample covers a wide set of industries, with the four most frequently represented industries being finance and insurance, manufacturing, professional and business services, and information technology. To account for differences between the sectoral composition of our sample and that of the broader US economy, we repeat the analysis of our key findings using weights that align the AI response sample with the sectoral distribution of firms as measured by the US Census.

We next show that aggregate sales revenue from CFO Survey respondents track closely with economy-wide realized sales growth, illustrating the reasonable forecasting accuracy of our respondents and the close alignment between our panel and the broader US economy. Figure A3 shows that year-ahead sales revenue growth forecasts for CFOs align closely with the realized revenue growth for the full US economy over the same period, as reported by the Bureau of Economic Analysis. The CFO forecasts missed the strong recovery in 2021 revenue coming out of the COVID downturn but perform reasonably well during 2022-2025.

Lastly, the accuracy of forecasts from the long history of The CFO Survey has also been verified by external sources. Gennaioli et al. (2016) provide evidence that CFO forecasts of earnings and corporate investment, respectively, closely track realized earnings and investment over the following 12 months, during both up and down cycles.

Figure A1: Expectations against Realizations for CFO Firms

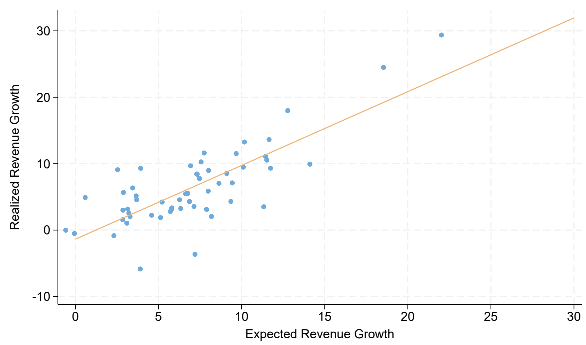

(a) Expected Revenue Growth

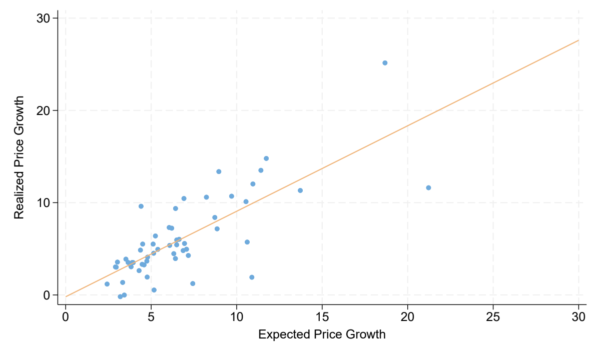

(b) Expected Price Growth

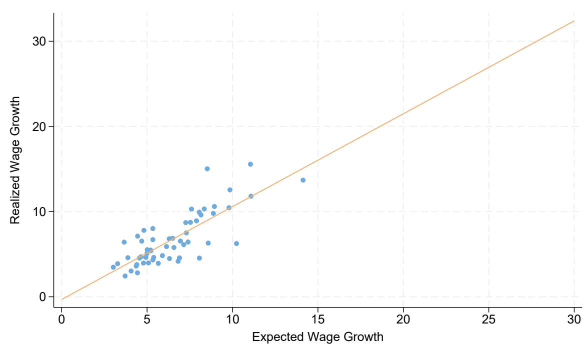

(c) Expected Wage Growth

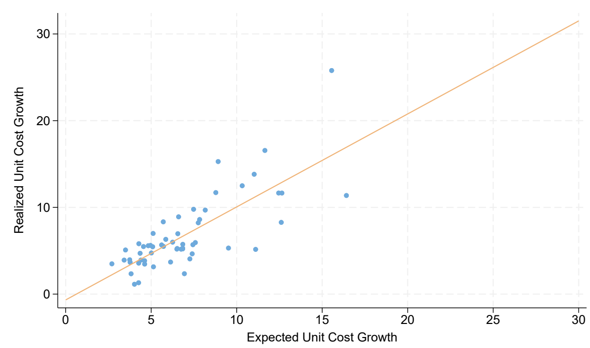

(d) Expected Unit Cost Growth

Notes: CFO's expectations are presented for revenue growth (a), price growth (b), wage growth (c) and unit cost growth (d), against average same-firm realizations of the four variables. Each dot represents a sector for a given calendar year. X-axis values reflects the average expectations for variable growth in a single calendar year, calculated as the average across respondents in each quarter within that calendar year. Y-axis value reflects the average realized growth in that variable (by sector) for that same calendar year, averaged across all responding firms. All calculations are unweighted. Line of best fit is plotted; A linear 45-degree line through the origin would represent average variable growth equaling averaged realize growth in a calendar year for a given sector.

Figure A2: Sector Employment Share

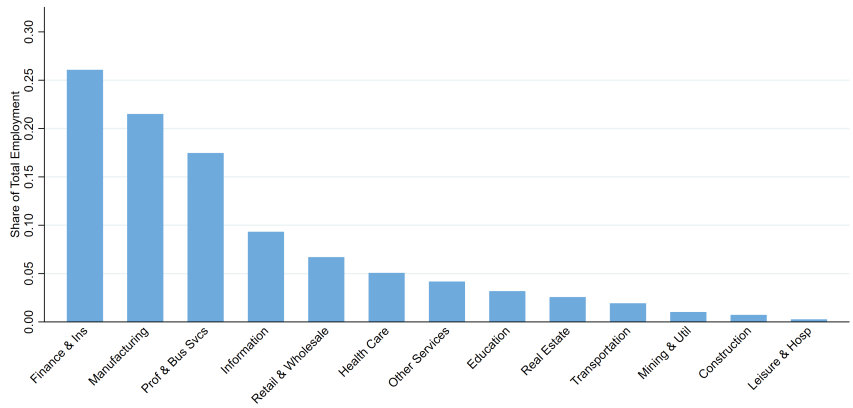

Notes: This figure shows share of total employment by sector among firms in the full AI survey sample.

Figure A3: Firm Sales Growth vs. Aggregate Sales Growth

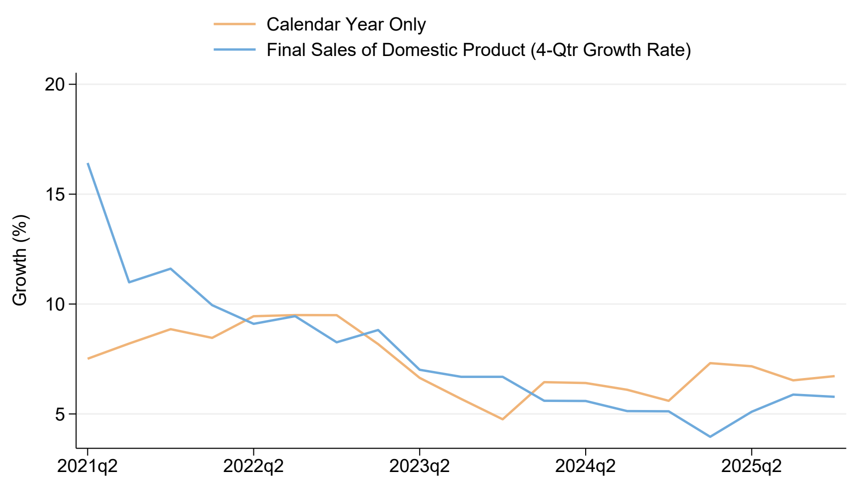

Notes: Calculations based on respondents to The CFO Survey for year-ahead forecasts make during 2020Q2 to 2024Q4, compared to realizations over 12-month periods ending in 2021Q2 to 2025Q4. The orange line is the forecast of sales revenue based on forecasts made in the four quarters ending in a given quarter, and the blue line is the realization of the preceding 12 months up to the given quarter. Responses are weighted by firm revenue. Realized data from the Bureau of Economic Analysis, Haver Analytics; forecast data are from the CFO Survey.

A2 Additional Figures

Figure A4: IT vs. AI Investment Waves

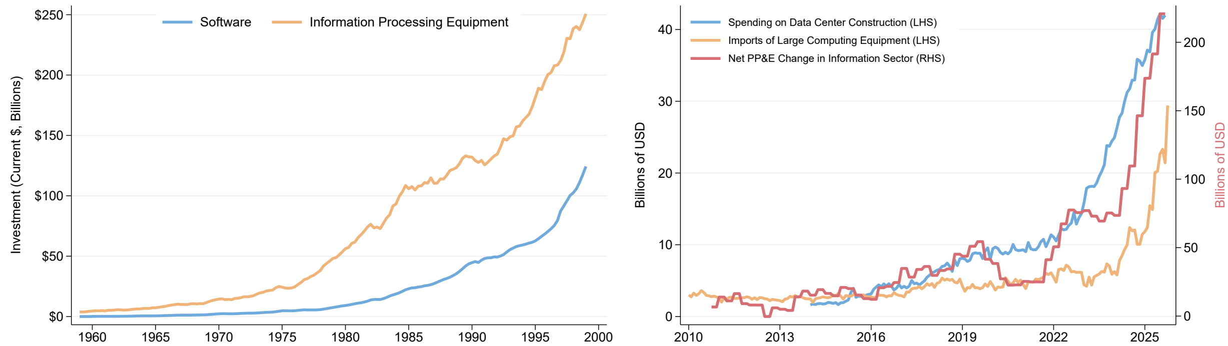

Notes: Panel A compares private fixed investment in current dollars between Software (BEA Series: B985RC) and Information Processing Equipment (BEA Series: B935RC). Information Processing Equipment contains categories for computers, peripheral equipment, photocopy machines, etc. Panel B compares spending on data center construction (US Census Bureau's Construction Spending Series) with imports of large computing equipment (HS Codes 847150 and 847330) and net year-over-year change in Property, Plant, and Equipment expenditures by US corporations from the US Census Bureau's Quarterly Financial Report (QFR).

Figure A5: Share of Firms Investing in AI by Sector and Firm Size.

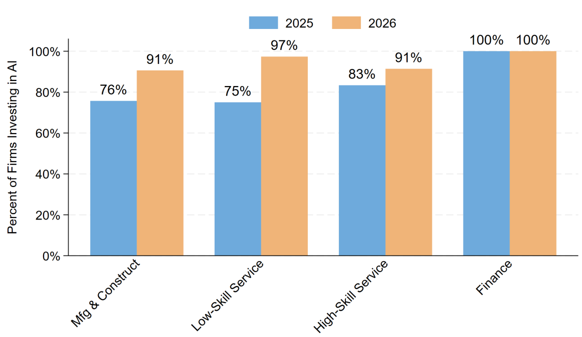

(a) Large Firms

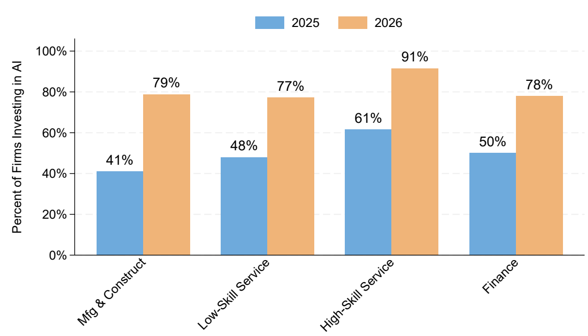

(b) Small Firms

Notes: Figure shows the extensive margin percentage of firms that have invested in 2025 (or expect to invest 2026) in AI technology, which in 2025 (2026) is determined by a Yes answer to a question asking whether the firm has invested in AI (a positive dollar amount for expected to AI investment in 2026). Mnfg&Constr includes "construction", "manufacturing", and "mining and utilities"; high-skill services include "educational services", "health care and social assistance", "information", "professional and business services", and "real estate and rental and leasing"; low-skill services include "leisure and hospitality", "retail and wholesale trade", "transportation and warehousing", and "other services except government". High-skill services are selected as those services for which the share of employees (from the 2024 CPS) with at least bachelor's degrees is greater than the cross-sector average. Low-skill services are the remaining services with lower-than-average bachelor's degree shares.

Figure A6: AI Investment over Employment, by Firm Size

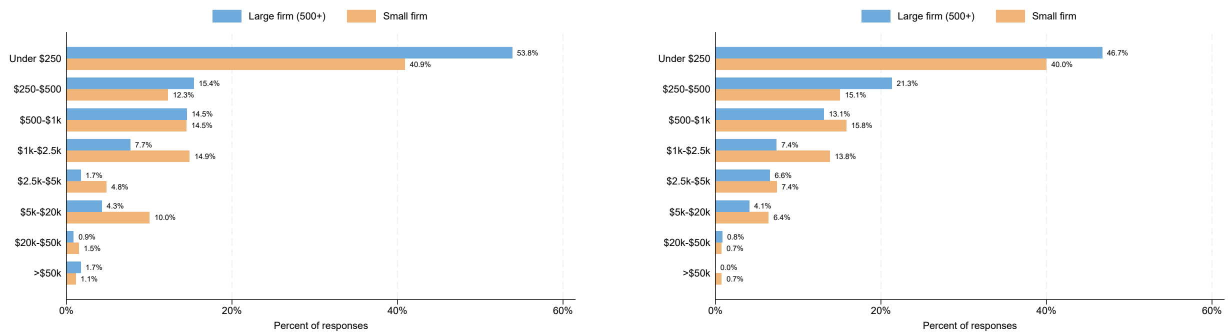

(a) 2025 (b) 2026

Notes: This figure depicts the differences in the amount of realized (expected) investment in AI technologies over 2025 (2026), divided by the firm's number of employees, between large and small firms. Blue bars represent 'large' firms with more than 500 employees. Red bars represent firms with fewer than 500 employees.

Figure A7: AI Investment over CapEx, by Firm Size

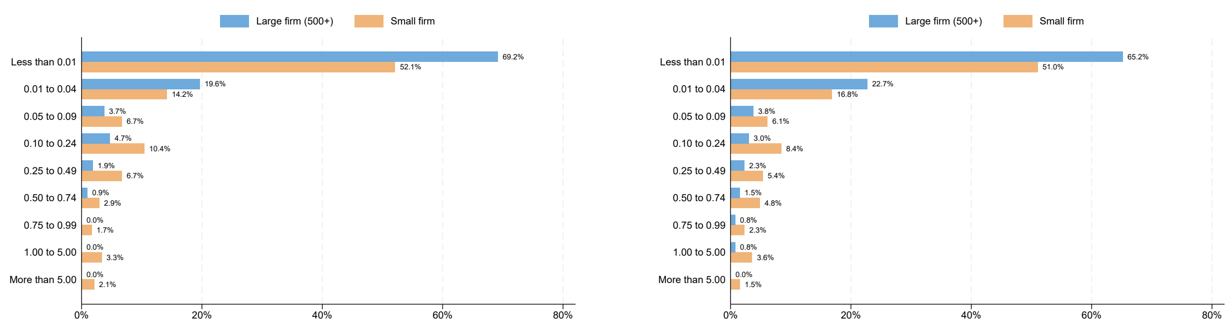

(a) 2025 (b) 2026

Notes: This figure depicts the differences in the amount of realized (expected) investment in AI technologies over 2025 (2026), divided by the firm's capital expenditures, between large and small firms. Blue bars represent 'large' firms with more than 500 employees. Red bars represent firms with fewer than 500 employees.

Figure A8: Reported vs. Implied Productivity Effects of AI. Robustness to Midpoint-Measurement

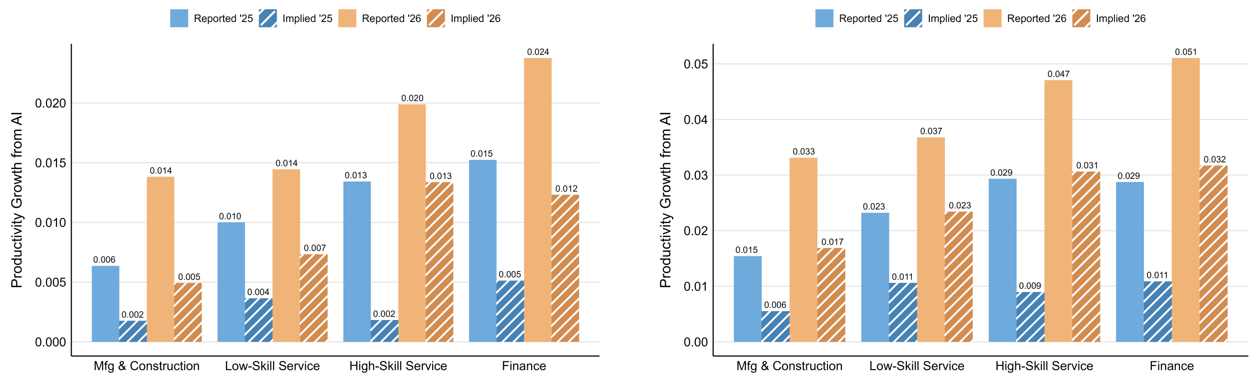

(a) Midpoint → lower end

(b) Midpoint → upper end

Notes: The figure is similar to Figure 6 in the main text but explores robustness to alternative codings of the survey responses. Bars show the mean reported (light-colored bars), $\Delta \ln LP_{it}^{AI,\,CFO}$, and implied (dark-shaded bars), $\Delta \ln(Y/L)_{it}^{AI}$, productivity growth attributable to AI for 2025 and 2026 across sectors. Panel (a) assigns the lower bound of each categorical response when mapping survey ranges to continuous outcomes, while panel (b) assigns the upper bound.

Figure A9: Reported vs. Implied Productivity Effects of AI. Firms with Positive Investment

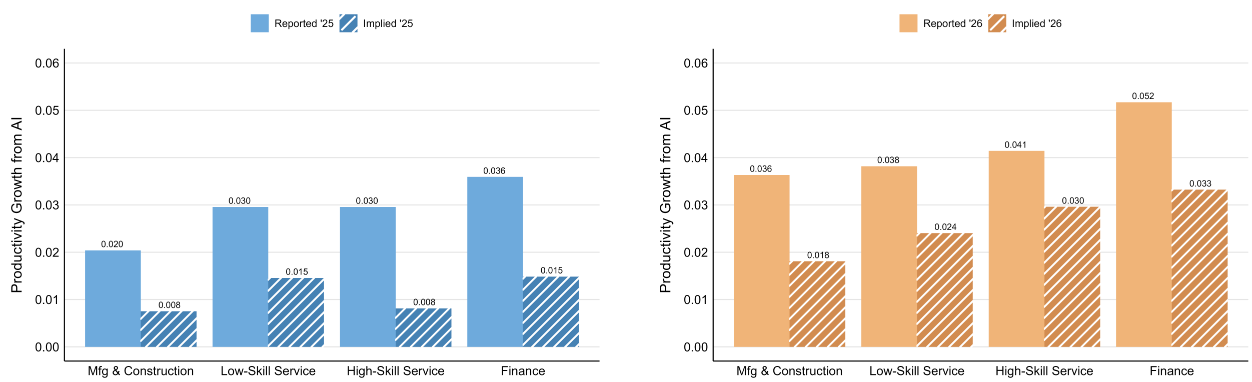

(a) Companies that Invested in AI in '25

(b) Companies that Invested in AI in '25 or '26

Notes: Bars show mean reported (light-colored bars), $\Delta \ln LP_{it}^{AI,\,CFO}$, and implied (dark-shaded bars), $\Delta \ln(Y/L)_{it}^{AI}$, productivity growth attributable to AI for 2025 and 2026 for firms with positive AI investment. Panel (a) reports results for firms with positive AI investment in 2025, while panel (b) reports results for firms with positive AI investment in 2025 and/or positive expected investment in 2026.

Figure A10: Reported vs. Implied Productivity Effects of AI: by Firm Size

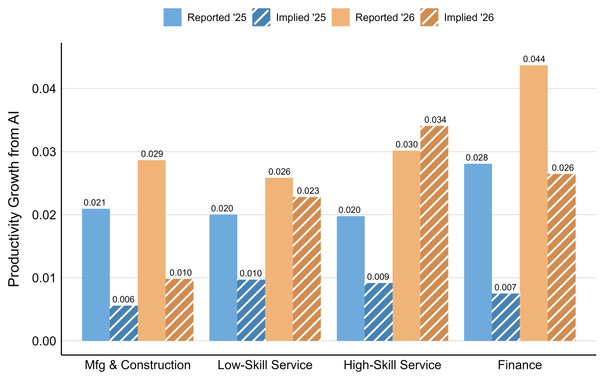

(a) Large Firms

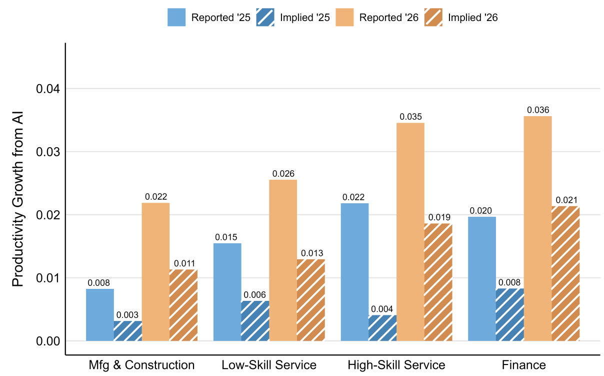

(b) Small Firms

Notes: Bars show mean reported (light-colored bars), $\Delta \ln LP_{it}^{AI,\, CFO}$, and implied (dark-shaded bars), $\Delta \ln(Y/L)_{it}^{AI}$, productivity growth attributable to AI for 2025 and 2026 across sectors. Panel (a) reports results for large firms with more than 500 employees, while panel (b) reports results for small firms.

Figure A11: Decomposing Labor Productivity Growth from AI: Firms with Positive Investment

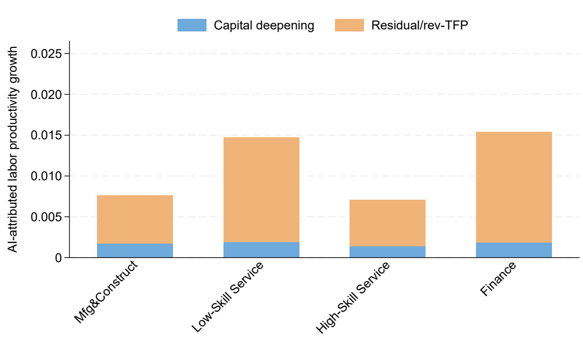

(a) 2025

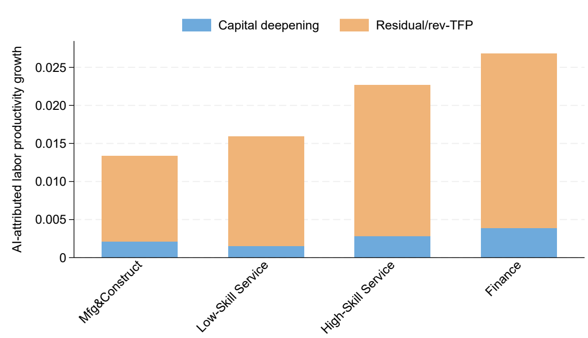

(b) 2026

Notes: Bars show the decomposition of mean implied AI-attributed labor productivity growth, $\Delta \ln(Y/L)_{it}^{AI}$, across sectors into capital deepening and the residual (revenue-based TFP), as defined in Equation (4) for firms with positive AI investment. Panel (a) reports results for firms with positive AI investment in 2025, while panel (b) reports results for firms with positive AI investment in 2025 and/or positive expected investment in 2026.

Figure A12: Decomposing Labor Productivity Growth from AI: by Firm Size, 2025

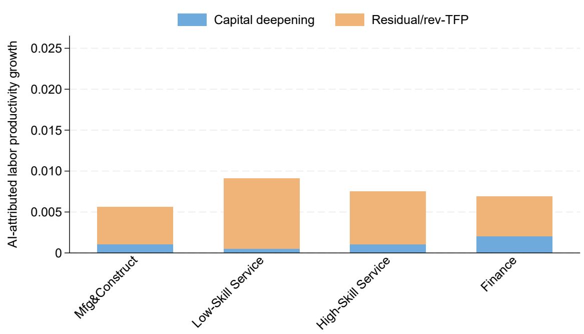

(a) Large Firms

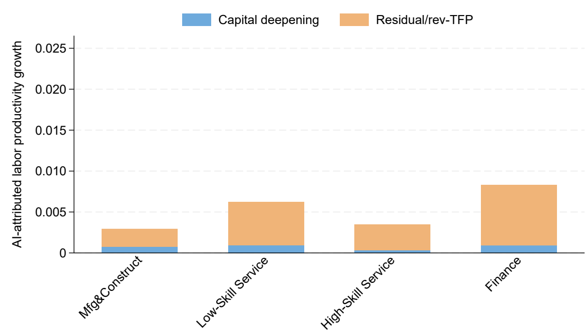

(b) Small Firms

Notes: Bars show the decomposition of mean implied AI-attributed labor productivity growth, $\Delta \ln(Y/L)^{AI}_{i2005}$, across sectors into capital deepening and the residual (revenue-based TFP), as defined in Equation (4). Panel (a) reports results for large firms with more than 500 employees, while panel (b) reports results for small firms.

Figure A13: AI Helpfulness for Various Tasks

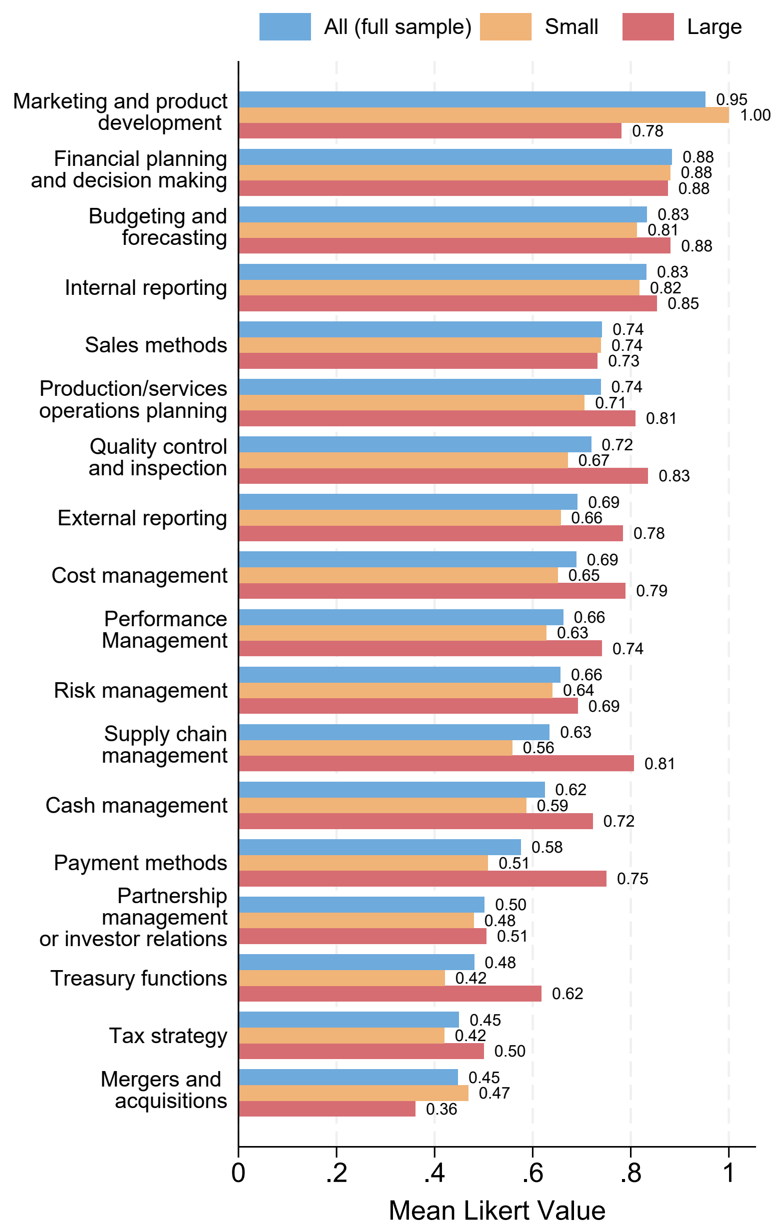

Notes: Based on Survey Question 13 (see Appendix Section A5), CFOs were asked to evaluate how helpful AI tools have been for accomplishing a range of business tasks. Responses are recorded on a Likert scale from 0 to 2 (0 = Not helpful, 1 = Moderately helpful, 2 = Very helpful). The figure reports mean responses by task category. Firms report that AI is most helpful in marketing and product development, financial planning and decision-making, budgeting and forecasting, and internal reporting, indicating strong perceived usefulness in analytical and information-processing tasks. In contrast, AI is reported to be less helpful for treasury management, tax strategy, and mergers and acquisitions, which tend to involve greater reliance on strategic judgment, institutional knowledge, and relationship-based activities. The figure also reports responses by firm size (large firms have at least 500 employees; small firms have fewer than 500 employees).

A3 Additional Tables

Appendix Table A1: Productivity Growth and Extensive and Intensive Margins of AI Investment. Weighted Estimation (representativeness)

	(1)	(2)	(3)	(4)	(5)	(6)
		2025			2026	
	Reported ΔLP	Implied ΔLP	Implied ΔTFP	Reported ΔLP	Implied ΔLP	Implied ΔTFP
Panel A: AI invest dummy						
AI Adopt ('25)	0.035***	0.008**	0.006*			
	(0.005)	(0.004)	(0.004)			
AI adopt ('25/'26)				0.035***	0.014***	0.012***
				(0.004)	(0.003)	(0.003)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.26	0.04	0.03	0.21	0.08	0.09
Observations	671	650	636	656	628	600
Panel B: AI investment (all firms, uncond)						
Log AI Inv ('25)	0.004***	0.001	0.000			
	(0.001)	(0.001)	(0.001)			
Log AI inv ('25/'26)				0.004***	0.002***	0.001***
				(0.000)	(0.000)	(0.000)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.27	0.03	0.02	0.27	0.09	0.10
Observations	647	629	617	615	591	565
Panel C: AI investment (only positives, cond)						
Log AI Inv ('25)	0.002	-0.005	-0.004			
	(0.004)	(0.004)	(0.004)			
Log AI inv ('25/'26)				0.007***	0.003	0.002
				(0.002)	(0.002)	(0.002)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.02	0.12	0.11	0.12	0.05	0.06
Observations	359	329	329	520	496	470

Notes: The table is similar to Table 2 in the main text but applies firm-level weights in the regressions. Firm-level 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).

Appendix Table A2: Productivity Growth and Extensive and Intensive Margins of AI Investment. Weighted Estimation (importance)

	(1)	(2)	(3)	(4)	(5)	(6)
		2025			2026	
	Reported ΔLP	Implied ΔLP	Implied ΔTFP	Reported ΔLP	Implied ΔLP	Implied ΔTFP
Panel A: AI invest dummy						
AI Adopt ('25)	0.022***	0.011***	0.009***			
	(0.004)	(0.003)	(0.003)			
AI adopt ('25/'26)				0.031***	0.020***	0.017***
				(0.003)	(0.003)	(0.003)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.13	0.05	0.04	0.11	0.05	0.05
Observations	671	650	636	656	628	600
Panel B: AI investment (all firms, uncond)						
Log AI Inv ('25)	0.002***	0.001**	0.001*			
	(0.000)	(0.000)	(0.000)			
Log AI inv ('25/'26)				0.002***	0.003***	0.003***
				(0.001)	(0.001)	(0.001)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.14	0.07	0.06	0.11	0.21	0.19
Observations	647	629	617	615	591	565
Panel C: AI investment (only positives, cond)						
Log AI Inv ('25)	0.002	0.001	0.001			
	(0.002)	(0.002)	(0.002)			
Log AI inv ('25/'26)				0.001	0.006***	0.005***
				(0.001)	(0.001)	(0.001)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.03	0.04	0.04	0.03	0.25	0.24
Observations	359	329	329	520	496	470

Notes: The table is similar to Table 2 in the main text but applies firm-level weights in the regressions. 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 1/95 level to reduce the impact of noise).

Appendix Table A3: Labor Productivity Gains from AI: Mechanisms. Weighted Estimation (representativeness)

	(1)	(2)	(3)	(4)	(5)	(6)
		2025			2026	
	Reported ΔLP	Implied ΔLP	Implied ΔTFP	Reported ΔLP	Implied ΔLP	Implied ΔTFP
Operational Efficiency Channel						
Production Efficiency	0.005	-0.001	-0.004	0.006	0.002	-0.002
	(0.009)	(0.006)	(0.005)	(0.009)	(0.006)	(0.006)
Reduce Labor Costs	0.012	0.004	0.009	0.013	0.007	0.007
	(0.011)	(0.007)	(0.008)	(0.009)	(0.009)	(0.008)
Reduce Other Costs	-0.013	-0.011	-0.018	-0.030***	-0.012	-0.012
	(0.013)	(0.012)	(0.013)	(0.012)	(0.010)	(0.010)
Decision Making/Mgmt	0.027***	-0.003	-0.006	0.018**	0.010	0.005
	(0.009)	(0.006)	(0.006)	(0.007)	(0.007)	(0.006)
Innovation & Demand Channel						
Product Development/Improvement	0.005	-0.001	-0.003	0.008	0.014**	0.013**
	(0.008)	(0.009)	(0.008)	(0.007)	(0.006)	(0.006)
Reach/Serve Customers	0.015	0.018***	0.013**	0.026***	0.019**	0.019***
	(0.011)	(0.006)	(0.006)	(0.008)	(0.008)	(0.006)
Factor Upgrading						
Upgrade Capital	0.002	-0.019	-0.013	0.009	0.006	0.005
	(0.014)	(0.013)	(0.013)	(0.013)	(0.012)	(0.011)
Workforce Development						
	-0.026**	0.000	0.006	-0.020**	-0.028***	-0.020***
	(0.010)	(0.007)	(0.007)	(0.008)	(0.009)	(0.007)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.21	0.13	0.13	0.29	0.28	0.28
Observations	370	351	337	382	364	343

Notes: The table is similar to Table 3 in the main text but applies firm-level weights in the regressions. Firm-level 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).

Appendix Table A4: Labor Productivity Gains from AI: Mechanisms. Weighted Estimation (importance)

	(1)	(2)	(3)	(4)	(5)	(6)
		2025			2026	
	Reported ΔLP	Implied ΔLP	Implied ΔTFP	Reported ΔLP	Implied ΔLP	Implied ΔTFP
Operational Efficiency Channel						
Production Efficiency	0.001	-0.008	-0.009	0.012	-0.004	-0.004
	(0.008)	(0.006)	(0.006)	(0.008)	(0.009)	(0.008)
Reduce Labor Costs	0.010	0.003	-0.000	0.005	0.022**	0.016*
	(0.008)	(0.008)	(0.006)	(0.008)	(0.010)	(0.009)
Reduce Other Costs	-0.013*	-0.002	-0.002	-0.016*	-0.007	-0.005
	(0.007)	(0.007)	(0.006)	(0.009)	(0.010)	(0.010)
Decision Making/Mgmt	0.016**	-0.006	-0.011	0.011*	0.010	0.003
	(0.007)	(0.008)	(0.008)	(0.006)	(0.008)	(0.007)
Innovation & Demand Channel						
Product Development/Improvement	0.015**	0.017*	0.013*	0.013*	0.026***	0.022***
	(0.006)	(0.009)	(0.008)	(0.007)	(0.008)	(0.007)
Reach/Serve Customers	0.006	0.010	0.013	0.013*	-0.006	-0.002
	(0.007)	(0.010)	(0.009)	(0.007)	(0.008)	(0.007)
Factor Upgrading						
Upgrade Capital	-0.004	0.000	0.001	-0.010	0.012*	0.013**
	(0.006)	(0.007)	(0.007)	(0.006)	(0.007)	(0.006)
Workforce Development	-0.013	-0.005	0.003	-0.004	-0.021**	-0.014*
	(0.008)	(0.006)	(0.006)	(0.008)	(0.009)	(0.008)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.13	0.12	0.15	0.15	0.21	0.19
Observations	370	351	337	382	364	343

Notes: The table is similar to Table 2 in the main text but applies firm-level weights in the regressions. 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 1/95 level to reduce the impact of noise).

Appendix Table A5: Expected Change in Types of Jobs and AI Investment

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Routine-Clerical		Δ Creative		Δ Skilled-Technical	
	2026	2028	2026	2028	2026	2028
Panel A: AI invest dummy						
AI Adopt ('25)	-0.726**		-0.230		0.429**	
	(0.302)		(0.361)		(0.217)	
AI adopt ('25/'26)		-3.196***		-0.033		1.203***
		(0.728)		(0.417)		(0.348)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.02	0.02	0.01	0.01	0.01	0.01
Observations	623	630	623	630	623	630
Panel B: AI investment (all firms, uncond)						
Log AI Inv (25)	-0.087***		-0.013		0.036	
	(0.024)		(0.028)		(0.023)	
Log AI inv (25/26)		-0.346***		0.033		0.118***
		(0.059)		(0.037)		(0.043)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.03	0.04	0.01	0.01	0.01	0.02
Observations	597	591	597	591	597	591
Panel C: AI investment (only positives, cond)						
Log AI Inv (25)	-0.268***		0.036		0.004	
	(0.097)		(0.101)		(0.111)	
Log AI inv (25/26)		-0.512***		0.132		0.149
		(0.130)		(0.122)		(0.135)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.04	0.02	0.01	0.01	0.01	0.01
Observations	381	517	381	517	381	517

Notes: The table reports firm-level regressions of changes in the share of routine, creative, and technical workforce in 2026 and 2028 relative to 2025 on extensive and intensive measures of AI investment. Panel A studies the extensive margin of AI investment. *AI Adopt'25* equals one if the firm reports any AI investment in 2025, and *AI Adopt'25/'26* equals one if the firm reports any AI investment in either 2025 or 2026. Panel B studies the intensive margin for all firms. *Log AI Inv'25* is the logarithm of AI investment reported in 2025 (plus one), and *Log AI Inv'25/'26* is the logarithm of total AI investment reported over 2025 and 2026 (plus one). Panel C restricts the sample to firms with positive AI investment. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table A6: Expected Change in Types of Jobs and AI Investment. Weighted Estimation

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Routine-Clerical		Δ Creative		Δ Skilled-Technical	
	2026	2028	2026	2028	2026	2028
Panel A: AI invest dummy						
AI Adopt ('25)	-0.789***		-0.621		0.390	
	(0.270)		(0.849)		(0.267)	
AI adopt ('25/'26)		-2.896***		0.419		0.774
		(0.655)		(1.028)		(0.492)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.02	0.03	0.01	0.02	0.02	0.03
Observations	611	618	611	618	611	618
Panel B: AI investment (all firms, uncond)						
Log AI Inv ('25)	-0.081**		-0.068		0.051	
	(0.033)		(0.058)		(0.043)	
Log AI inv ('25/'26)		-0.265***		-0.015		0.100
		(0.076)		(0.083)		(0.083)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.03	0.05	0.01	0.02	0.02	0.03
Observations	585	580	585	580	585	580
Panel C: AI investment (only positives, cond)						
Log AI Inv ('25)	-0.178		-0.047		0.070	
	(0.132)		(0.148)		(0.185)	
Log AI inv ('25/'26)		-0.297*		-0.043		0.111
		(0.157)		(0.132)		(0.189)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.03	0.03	0.02	0.02	0.01	0.03
Observations	375	508	375	508	375	508

Notes: The table is similar to Table A5 but applies firm-level weights in the regressions. 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 1/95 level to reduce the impact of noise).

A4 Calculating Capital Deepening Using AI Expenditures

This appendix develops an alternative proxy for AI-attributed capital deepening that leverages firms' reported AI expenditures, rather than self-reported changes in the book value of capital.

Recall that implied AI-attributed labor productivity growth decomposes as

$$\Delta \ln(Y/L)_{it}^{AI} = \alpha \, \Delta \ln(K/L)_{it}^{AI} + \Delta \ln A_{it}^{AI}, \qquad \Delta \ln(K/L)_{it}^{AI} = \Delta \ln K_{it}^{AI} - \Delta \ln L_{it}^{AI}.$$

As in the baseline specification, $\Delta \ln L_{it}^{AI}$ is taken directly from company responses. The key difference lies in how we proxy $\Delta \ln K_{it}^{AI}$—the component of the change in the firm's capital input attributable to AI.

$\Delta \ln K_{it}^{AI}$ is a counterfactual object defined as the difference between the firm's end-of-period capital stock in the observed economy and the capital stock that would have prevailed absent AI investment, holding fixed all non-AI investment decisions (denoted as K_{it}^{0}):

$$\Delta \ln K_{it}^{AI} \equiv \ln K_{it} - \ln K_{it}^{0}. \tag{5}$$

Let total capital evolve according to the standard law of motion,

$$K_{it} = (1-\delta) K_{i,t-1} + I_{it}^{nonAI} + I_{it}^{AI}, \tag{6}$$

where δ is depreciation, and I_{it}^{nonAI} and I_{it}^{AI} are non-AI and AI investments, respectively. Then the no-AI counterfactual is obtained by setting $I_{it}^{AI} = 0$:

$$K_{it}^{0} = (1-\delta) K_{i,t-1} + I_{it}^{nonAI}. \tag{7}$$

Then, using a first-order approximation,

$$\Delta \ln K_{it}^{AI} = \ln\left(1 + \frac{I_{it}^{AI}}{K_{it}^{0}}\right) \approx \frac{I_{it}^{AI}}{K_{it}^{0}}. \tag{8}$$

Since K_{it}^{0} is unobserved, we scale AI investment by the beginning-of-period capital stock:[21]

$$\Delta \ln K_{it}^{AI} \approx \frac{I_{it}^{AI}}{K_{i,t-1}}. \tag{9}$$

This expression says that an AI investment equal to x percent of last year's capital stock raises the firm's capital input by approximately x percent due to AI.

To measure I^{AI}, we partially capitalize the AI expenditure figures reported by the firms. In particular, for a subset of firms (N=183), we asked about the allocation of AI expenditures into capitalizable outlays (e.g., AI-related hardware and internally developed software) and operating

[21] For the 2025 calculations, we require $K_{i,2024}$, which is unavailable for most firms. We therefore construct $K_{i,2024}$ from the 2025 book value of PP&E and intangible assets, assuming a 20% depreciation rate.

expenses (e.g., software subscriptions and cloud computing services). As shown earlier in Figure 3, a large share of AI expenditures is attributable to operating costs, especially for large firms. Following Bureau of Labor Statistics and national accounts practices whereby software subscriptions and cloud computing services (SaaS, IaaS, PaaS) are generally treated as intermediate inputs rather than capital investment, while expenditures on hardware and internally developed software are included in the capital stock, we capitalize only the portion of AI expenditures corresponding to hardware and internal development. For firms for which we do not have the AI expenditure split, we extrapolate the capitalization shares using the allocation from the known sample. In particular, for 2025, a median large firm reports 55% of capitalizable outlays, while the median small firm reports 0%; the same figures for 2026 are 40% and 20%, respectively.[22]

Figure A14 presents the decomposition obtained using this alternative, expenditure-based measure of AI-attributed capital growth. Consistent with the baseline results, capital deepening accounts for a modest share of AI-attributed revenue labor productivity growth (about 16% on average), yielding a decomposition that is nearly identical to the baseline. For 2026, the expenditure-based approach implies a somewhat larger contribution of capital deepening—particularly in high-skill services—but the average contribution remains moderate at 24%.

The close alignment between the baseline decomposition and this alternative approach reinforces our conclusions regarding the sources of AI-related productivity growth. In particular, the similarity between the decomposition based on firm-reported AI-attributed changes in PP&E and that obtained by capitalizing AI expenditures related to hardware and internal development suggests that respondents interpreted the PP&E question as intended and appropriately distinguished capital investment from operating AI expenses.

Figure A14: Decomposing Labor Productivity Growth from AI: Capitalizing AI Expenditures

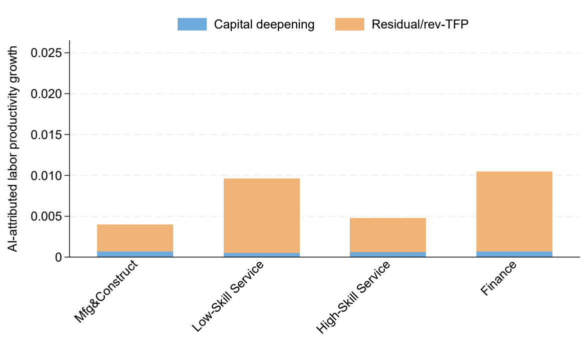

(a) 2025

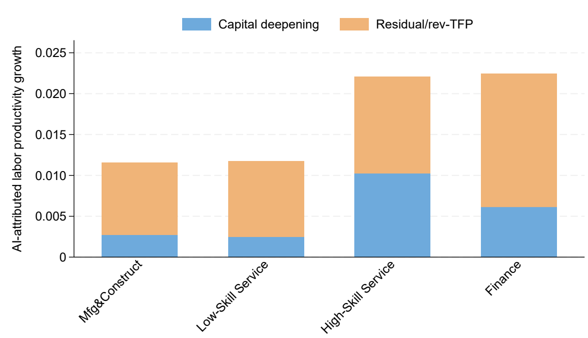

(b) 2026

Notes: Bars show the decomposition of mean implied AI-attributed labor productivity growth, $\Delta \ln(Y/L)_{it}^{AI}$, across sectors into capital deepening and the residual (revenue-based TFP), as defined in Equation (4). Capital growth due to AI is constructed using partial capitalization of firm-reported AI expenditures as described in Section A4. Panel (a) decomposes productivity gains for 2025, while panel (b) decomposes expected gains for 2026.

[22] For 2025 calculations, we need to use K_{i2024}, which is not available for most of the firms in the data. We construct the value of K_{i2024} using the 2025 book value of PP&E and intangible assets and a 20% depreciation rate.

A5 Survey Questions on Artificial Intelligence

Answer options are in italic, separated by a semicolon.

1. Artificial intelligence (AI) refers to the broad field of machines capable of replicating human behavior and intelligence.

   Over the last 12 months, has your firm made any expenditures or financial investments in AI technology or solutions?

   This includes:

   - AI applications (e.g., large language models, machine learning, speech/voice recognition, data/text analytics, virtual agents/chatbots, visual content creation, robotics, mechanization, etc.)
   - Infrastructure to support these applications - including equipment (e.g., semiconductors and information processing equipment), structures (e.g., data centers or power production), and intellectual property (e.g., spending on model RD and software).

   Yes; No; Not Sure

2. (If "Yes" to 1) Please rate the importance of the following possible motivations behind your firm's expenditure/investment in AI:

   - Improving production efficiency
     - e.g., improving speed of processes, automating or optimizing internal processes, logistics, or maintenance
   - Improving Labor Productivity
     - Increasing revenue per worker or output per worker
   - Reducing Labor Costs
   - Reducing non-labor costs
   - Enhancing Decision-making and Management
     - e.g., data analytics, forecasting, workflow/HR optimization
   - Developing or Improving Products/Services
     - e.g., new or higher-quality offerings, personalization, testing, faster RD cycles
   - Reaching or Serving Customers More Effectively
     - e.g., marketing, customer interaction, after-sales support
   - Upgrading Physical or Digital Capital
     - e.g., investments in hardware, data infrastructure, or cloud systems
   - Developing Workforce Skills/Adapting Labor

– e.g., training, hiring, or reorganizing teams for AI use

- Other (Please Specify)

Not at all important; Slightly important; Moderately important; Very important; Extremely important

3. (If "Yes" to 1) What was your company's total expenditure and financial investment in AI technology/solutions <u>over the last 12 months</u>?

   $0; $1-$5,000; $5,001-$20,000; $20,001-$50,000; $50,001-$100,000; $100,001-$500,000; $500,001-$1 million; $1 million - 5 million; Over $5 million; Prefer not to say

4. (If "Over $5 million " in 3) Roughly, what was your company's total expenditure and financial investment in AI technology/solutions <u>over the last 12 months</u>? *(estimates are acceptable)*

5. (If "No" to 1) Why did you firm <u>not</u> make any expenditures or financial investments in AI technology or solutions over the last 12 months? *(Please select all that apply)*

   Too expensive; AI is not a mature enough technology yet; Lack of knowledge on the capabilities of AI; Concerns about privacy/security; Concerns about bias; Our workers are not yet adequately trained on AI; Lack of required data; Laws or regulations prevent or restrict use of AI; Previous or current use of AI did not meet expectations; AI is not applicable to this business; Other (Please explain)

6. What do you expect your company's total expenditure and financial investment in AI technology/solutions will be <u>over the next 12 months</u>?

   $0; $1-$5,000; $5,001-$20,000; $20,001-$50,000; $50,001-$100,000; $100,001-$500,000; $500,001-$1 million; $1 million - 5 million; Over $5 million; Prefer not to say

7. (If "Over $5 million " in 6) Roughly, what do you expect your company's total expenditure and financial investment in AI technology/solutions will be <u>over the next 12 months</u>? *(Estimates are acceptable)*

8. Over the <u>last 12 months</u>, how has your firm's use of AI affected the following outcomes fo your firm?

   - Labor Productivity (output per worker)
   - Total employment
   - Book value of PP&E and intangible assets (e.g. software)
   - Operating costs (exuding labor)
   - Labor costs per worker
   - Revenue or sales
   - Revenue from new products

- Decision-making speed or accuracy
- Customer satisfaction or retention
- Time spent on high value-add tasks
- Other

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%); Unsure/Not Applicable

9. Over the next 12 months, how has your firm's use of AI affected the following outcomes fo your firm?

    - Labor Productivity (output per worker)
    - Total employment
    - Book value of PP&E and intangible assets (e.g. software)
    - Operating costs (exuding labor)
    - Labor costs per worker
    - Revenue or sales
    - Revenue from new products
    - Decision-making speed or accuracy
    - Customer satisfaction or retention
    - Time spent on high value-add tasks
    - Other

    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%); Unsure/Not Applicable

10. The questions on this screen pertain to AI tools and your workforce. Please describe the roles/responsibilities of the employees that were (or you expect to be) replaced by AI tools. *[Open-ended response.]*

11. Please describe the roles/responsibilities of employees whose roles were (or you expect to be) complemented or enhanced by AI tools. *[Open-ended response.]*

12. Please provide the percent of your firm's total full-time employees that fall into each of the below categories. Please also provide your expectations for the percent of your total headcount that will fall into these categories 12 and 36 months from now. *(Results for each column should sum to 100%)*

- Routine/clerical (e.g. data entry, accounting)
- Skilled Technical (e.g. engineers, data analysts/scientists)
- Creative/managerial (e.g. design/strategy/leadership)
- All other

Current; 12 Months from Now; 36 Months from Now

13. How helpful are AI tools or solutions in performing the following responsibilities?

    - Budgeting and forecasting
    - Cash flow management
    - Cost management
    - Financial planning & decision-making
    - External reporting (e.g. financial statements)
    - Internal reporting (e.g. board reports)
    - Marketing and product development
    - Mergers and acquisitions
    - Partnership management or investor relations
    - Payment methods
    - Performance management
    - Production/services operations planning
    - Quality control and inspection
    - Risk management
    - Sales methods
    - Supply chain management
    - Tax strategy
    - Treasury functions
    - Other (Please describe)

    Very helpful; Moderately helpful; Not helpful; Unsure; Not Applicable/Do Not Know

14. Please provide specific examples of AI usage or initiatives that provide value for your firm. *[Open-ended response.]*

15. Roughly, how much did your firm spend on capital expenditures (including structures, land, equipment, software, and AI investment) over the last 12 months? *(Estimates are acceptable)*

16. Roughly, what is your company's current book value of PP&E and intangible assets (e.g. software)? *(Estimates are acceptable)*
17. Roughly, what was your firm's total labor costs (including wages, salaries, and benefits) over the last 12 months? *(Estimates are acceptable)*
18. Roughly, what were your firm's costs for intermediate inputs (COGS + materials/energy/hired professional services) over the last 12 months? *(Estimates are acceptable)*
19. Please indicate what percent of your firm's spending on AI technology/solutions over the next 12 months was allocated to each of the below categories. *(Results should sum to %100)*
    - Hardware for AI (e.g. servers/GPUs, devices)
    - Developing or customizing internal AI systems
    - Operational expenses: AI subscriptions, services, and training (cloud-based tools, software fees, consultants, employee training, data prep)
    - Other (Please explain)