

# A New Look at AI's Impact on Jobs: Firm-Level AI Spending and Workforce Adjustment\*

Ara Kharazian<sup>†</sup>

*Ramp*

Lisa Simon<sup>‡</sup>

*Revelio Labs*

Ryan Stevens<sup>§</sup>

*Ramp*

June 30, 2026

## Abstract

We study how employment changes when firms adopt generative AI using observed AI spending from Ramp card and bill pay data linked to Revelio Labs workforce records for 21,559 firms in the United States. We find that companies that adopt AI tend to grow faster following adoption, but the relationship is driven almost entirely by high-intensity adopters. Firms making the largest AI investments grow employment by roughly 10% following adoption, while low-intensity adopters see no statistically significant change. Entry-level headcount rises 12% for high-intensity adopters. Gains emerge gradually and are broad across roles, including engineering, sales, administration, and customer service. They are also uneven: adopters are already larger, more technical, faster-growing firms, and sector-level gains are concentrated in Information. The results counter predictions that AI adoption will lead to broad job loss.

---

\*We thank the following individuals, whose thoughts and feedback supported our work, and who could never be replaced by AI: Nick Bloom, Scott Cunningham, Martha Gimbel, Rafael Hajjar, Jon Hartley, Jeremy Huang, Nikki Neumann, Torsten Sløk, Noah Smith, and Ernie Tedeschi. The latest version of this paper is available online at [ramp.com/data/ai-jobs-impact](https://ramp.com/data/ai-jobs-impact).

<sup>†</sup>Lead Economist, Ramp. Corresponding author. [ara.kharazian@ramp.com](mailto:ara.kharazian@ramp.com)

<sup>‡</sup>Chief Economist, Revelio Labs. [lisa@reveliolabs.com](mailto:lisa@reveliolabs.com)

<sup>§</sup>Director of Applied Science, Ramp. [rstevens@ramp.com](mailto:rstevens@ramp.com)

# Contents

- 1 Introduction** **6**
  
- 2 Related Literature** **7**
  
- 3 Data** **11**
  - 3.1 Ramp Business Spend Data . . . . . 11
  - 3.2 Benchmarking AI Adoption Rates . . . . . 12
  - 3.3 Revelio Workforce Data . . . . . 14
  - 3.4 Matching and Analytical Sample . . . . . 15
  
- 4 What Drives AI Adoption?** **16**
  - 4.1 Pre-Adoption Growth and the Preferred Comparison . . . . . 19
  
- 5 Empirical Methodology** **20**
  - 5.1 Treatment Intensity . . . . . 21
  - 5.2 Estimator and Identification . . . . . 22
  - 5.3 Outcomes and Diagnostics . . . . . 23
  - 5.4 Comparison-Group Checks . . . . . 23
  
- 6 Results** **24**
  - 6.1 High-Intensity AI Adopters Grew Headcount by Ten Percent . . . . . 24
  - 6.2 High-Intensity Employment Gains Compounded Over Time . . . . . 26
  - 6.3 Entry-Level Employment Also Rose Among High-Intensity Adopters . . . . . 26
  - 6.4 Employment Gains Were Broad Across Roles . . . . . 29
  - 6.5 Early Employment Gains Are Concentrated in Information Firms . . . . . 30
  - 6.6 Gains Are Even Higher Against Never-Adopters . . . . . 31
  
- 7 Conclusion** **33**

**8 Event-Study Figures**

**35**

**References**

**45**

# List of Tables

- 1 Sample Summary by AI Intensity . . . . . 16
- 2 Adopter Characteristics . . . . . 17
- 3 Headcount by AI Intensity . . . . . 25
- 4 Workforce Shares by AI Intensity . . . . . 28
- 5 Headcount by Sector . . . . . 31
- 6 Never-Treated by AI Intensity . . . . . 32

# List of Figures

- 1 AI Adoption Benchmarks . . . . . 13
- 2 AI Adoption by Sector . . . . . 18
- 3 AI Adoption by Size and Engineering Share . . . . . 18
- 4 Indexed Headcount and Pre-Trends: AI Adopters vs. Never Adopters . . . 19
- 5 Indexed Headcount and Pre-Trends: AI Adopters vs. Not-Yet Adopters . . 20
- 6 Total Headcount . . . . . 36
- 7 Entry-Level Headcount . . . . . 36
- 8 Non-Entry Headcount . . . . . 37
- 9 Manager-Plus Headcount . . . . . 37
- 10 Engineering Headcount . . . . . 38
- 11 Entry-Level Engineering . . . . . 38
- 12 Sales Headcount . . . . . 39
- 13 Marketing Headcount . . . . . 39
- 14 Admin Headcount . . . . . 40
- 15 Finance Headcount . . . . . 40
- 16 Scientist Headcount . . . . . 41
- 17 Operations Headcount . . . . . 41
- 18 Customer Service Headcount . . . . . 42
- 19 Bachelor’s Headcount . . . . . 42
- 20 MBA Headcount . . . . . 43
- 21 JD Headcount . . . . . 43
- 22 PhD Headcount . . . . . 44

# 1 Introduction

The most important economic question of the next decade will answer how artificial intelligence will impact labor markets. When firms adopt, do they hire fewer workers or reorganize and expand? AI can automate tasks that workers previously performed, but it can also raise productivity, create new demand, and make complementary work more valuable.

Meanwhile, recent public statements by company executives have linked layoffs to AI adoption. Developers of AI itself oscillate between warning of mass job loss and unprecedented economic wealth.

Existing research has provided important early evidence on AI's labor market effects, but has largely relied on occupational exposure measures, surveys, or indirect proxies for adoption because direct firm-level adoption data have been difficult to obtain. This paper studies actual adoption: what's happening inside of real firms? We link Ramp line-item spending data, which identify payments to AI vendors, to Revelio Labs workforce records, which measure monthly employment at the same firms.

We find AI adoption is associated with hiring growth, especially for the firms that use AI a lot. Using AI spend per employee, we find that total headcount is not detectably higher for low AI-intensity adopters and is 10.2 percent higher for high AI-intensity adopters over the first 24 months after adoption. Entry-level headcount rises 12.0 percent for high-intensity adopters. High-intensity adopters also expand in AI-exposed work such as engineering, sales, customer service, and finance.

Descriptive evidence shows that adopters are highly selected: they are larger, more technical, and faster growing before adoption. The preferred design therefore compares adopters to later adopters in the same AI-intensity group while those later adopters have not yet adopted, with sector fixed effects. The post-adoption paths then widen most clearly for high-intensity adopters, consistent with a "learning curve" as firms and employees identify use cases, establish best practices, and integrate AI tools into workflows.

To our knowledge, this is the first paper to combine observed firm-level AI spending with workforce records at scale. It uses a new dataset, grounded in real firm spend, to replace messy surveys and exposure measurements. It provides early evidence on how firms adopting AI are changing their workforce composition and employment trajectories. It provides an early look at how work will change over the next decade. Finally, it's a look into the future: how early adopters of a new technology are changing their behavior, and how that may shape our own jobs in the future.

## 2 Related Literature

This paper relates to three literatures: the labor market effects of automation and AI, the measurement of firm-level technology adoption, and the econometrics of staggered treatment timing.

**Automation, AI, and labor demand.** The task-based framework of [Autor et al. \(2003\)](#) and [Acemoglu and Restrepo \(2018, 2019\)](#) distinguishes between technologies that substitute for workers in routine tasks and those that complement workers in non-routine tasks. Generative AI disrupts this distinction. Large language models perform non-routine cognitive work—drafting text, writing code, synthesizing information—that prior waves of automation left largely unaffected ([Eloundou et al., 2024](#)). The theoretical prediction is therefore ambiguous, and the empirical evidence to date is limited in scope.

At the micro level, studies of individual AI deployments find productivity gains without contemporaneous job losses. [Brynjolfsson et al. \(2025\)](#) document a 14 percent increase in resolution rates among customer-service agents at a single firm following the introduction of a generative AI assistant. [Noy and Zhang \(2023\)](#) report a 40 percent reduction in task completion time for professional writing in an experimental setting. These results speak to the intensive margin of AI's effects within a given firm or task, but they cannot address equilibrium employment adjustments across firms and occupations.

At the aggregate level, [Brynjolfsson et al. \(2025\)](#) provide the most direct evidence to date. Using payroll records from ADP and occupational AI exposure scores derived from [Eloundou et al. \(2024\)](#), they estimate an approximately 16 percent decline in employment for workers ages 22–25 in the highest-exposure occupations following the release of ChatGPT. Their identification relies on cross-occupation variation in predicted exposure within firms, and the authors note the absence of firm-level adoption data as a primary limitation. Our paper addresses this limitation. In a longer historical perspective, [Autor et al. \(2024\)](#) document that automation has historically coincided with the creation of new job categories that partially absorb displaced labor, though whether this pattern will hold for generative AI remains an open question.

**Measuring AI adoption.** A central challenge in this literature is that researchers have lacked direct measures of which firms adopt AI and when. The dominant approach has been to construct occupational exposure scores that vary across jobs but not across firms or over time. [Eloundou et al. \(2024\)](#) use GPT-4 to evaluate the share of O\*NET tasks performable by large language models. [Webb \(2020\)](#) matches AI patent text to occupational task descriptions. [Felten et al. \(2021\)](#) link AI capability benchmarks to occupational ability requirements. These indices have been widely used—[Brynjolfsson et al. \(2025\)](#) assign treatment by exposure quintile, and [Acemoglu et al. \(2022\)](#) use AI-related job postings as a proxy for establishment-level adoption—but they share a fundamental limitation: they cannot distinguish a firm that has adopted AI from one that has not.

[Massenkoff and McCrory \(2026\)](#) advance the measurement frontier by constructing an index of observed exposure from actual Claude usage data collected through the Anthropic Economic Index. Their measure captures which tasks workers actually delegate to AI, rather than which tasks AI could theoretically perform, and reveals a substantial gap between the two: only 33 percent of Computer and Mathematical tasks show meaningful observed usage despite 94 percent being theoretically feasible. The unit of analysis, how-

ever, remains the occupation. The measure cannot identify firm-level adoption decisions or their within-firm consequences.

The limitations of exposure-based approaches are well documented. [Gimbel et al. \(2026\)](#) compare seven exposure metrics and find different measures disagree: computer programmers rank in the 99th percentile under one index and the 88th under another. The authors argue that relying on any single score misrepresents AI's likely effects. The deeper problem is that exposure varies only across occupations, not across firms within an occupation. Two firms employing identical workers may differ sharply in AI adoption, and exposure indices cannot separate them. Despite these limitations, exposure measures have persisted in the literature because firm-level AI adoption data have not been available in standard economic datasets. Generative AI is adopted through software subscriptions and API calls, not through the highly visible and trackable capital equipment purchases that allowed researchers to benchmark earlier technologies.

Several papers construct firm-level proxies for AI investment, but each is indirect. [Babina et al. \(2024\)](#) measure the share of a firm's workforce with AI-related skills listed on their resumes and find that AI-investing firms experience higher sales growth and employment growth. [Hampole et al. \(2025\)](#) construct a time-varying, firm-specific AI exposure measure from patent-to-task similarity and instrument for it using historical university hiring networks. These approaches provide useful variation but do not observe the actual spending decisions that constitute adoption.

Corporate payment records offer a more direct measure. [Stevens \(2026\)](#) uses transaction data from Ramp to study substitution between contracted online labor and generative AI, exploiting ChatGPT's release as a shock in a difference-in-differences framework. He finds that firms with high baseline spending on freelance labor marketplaces subsequently adopted AI more intensively while reducing contracted labor spending, with a one-dollar decline in online labor spending associated with approximately three cents of additional AI spending. This is the first micro-level evidence of direct labor-AI substi-

tution in firm production, but the analysis is limited to a single expenditure margin and does not examine broader employment changes.

We build on [Stevens \(2026\)](#) by combining corporate transaction data from Ramp with Revelio Labs workforce data. Ramp records every payment a customer makes to AI vendors—the amount, timing, and vendor identity. Revelio Labs provides employee-level position histories derived from online public profiles, including job titles, O\*NET codes, seniority, and estimated salaries. Linking these datasets at the firm level allows us to observe both the adoption decision and its workforce consequences, defining treatment by actual AI spending rather than predicted occupational exposure. We classify adoption intensity using AI spend per employee over the first three months after adoption begins, testing whether employment dynamics differ across firms that adopt AI lightly or intensively soon after treatment.

**Econometric framework.** Firms in our sample begin paying AI vendors at different points between late 2022 and the present, creating a staggered adoption design. A substantial recent literature has shown that conventional two-way fixed effects event studies can produce biased estimates of dynamic treatment effects when effects are heterogeneous across adoption cohorts ([Goodman-Bacon, 2021](#); [de Chaisemartin and D’Haultfoeuille, 2020](#); [Sun and Abraham, 2021](#)). We therefore use the group-time average treatment effect framework of [Callaway and Sant’Anna \(2021\)](#), estimated separately for low- and high-intensity adopters relative to later adopters in the same eventual intensity group while those firms have not yet adopted. This design is better suited to heterogeneous treatment timing than a single two-way fixed effects event study. Because adoption is selected, we treat pre-treatment event-study paths as central diagnostics and report never-treated comparisons as descriptive evidence rather than the preferred counterfactual.

### 3 Data

We combine two sources of firm-level information. Ramp line-item spend identifies whether a firm pays AI vendors, how much it spends, and when AI spending begins. Revelio Labs workforce records provide monthly measures of employment at those same firms. The linked panel is observed monthly from January 2021 through February 2026 to reduce reporting-lag bias in workforce data.

#### 3.1 Ramp Business Spend Data

Ramp is a financial operations platform that processes corporate card and bill pay transactions. We observe cleared spend at the line-item level, including date, amount, merchant, firm identifiers, and line-item detail when available. We aggregate these records to the firm-month level.

Using the Ramp AI Index ([Kharazian, 2026](#)), we identify AI spend from Ramp line items using a vendor and line-item classification designed to capture spending on foundational large language models, GPU cloud, model serving and inference, coding agents, API tokens, AI image and video generators, and AI search and research software.

For each firm  $i$ , let  $AI_{it}$  denote monthly spend at AI vendors. We define AI adoption as the first month of the earliest three-consecutive-month spell in which AI spend is at least \$100 in every month:

$$G_i = \min \{t : AI_{it} \geq 100, AI_{i,t+1} \geq 100, AI_{i,t+2} \geq 100\}.$$

Treatment is absorbing after  $G_i$ . Firms that never satisfy this rule are classified as never treated. This definition excludes one-off employee experiments and incidental AI payments while preserving firms that make sustained, organization-level AI purchases.

Among treated firms with an observable pre-treatment window, we measure adoption intensity as monthly AI spend per employee over the first three months after adoption

begins. The denominator is baseline headcount, measured in the month before adoption:

$$PEPM_i = \frac{\sum_{m=0}^2 AI_{i,G_i+m}}{3 \times HC_{i,G_i-1}}.$$

Eligible treated firms are sorted by  $PEPM_i$ : Low combines the bottom two PEPM terciles, and High is the top PEPM tercile. The estimator then compares adopters to firms in the same eventual intensity group that have not yet adopted by the relevant event time. We use never-treated firms for descriptive comparisons and robustness checks.

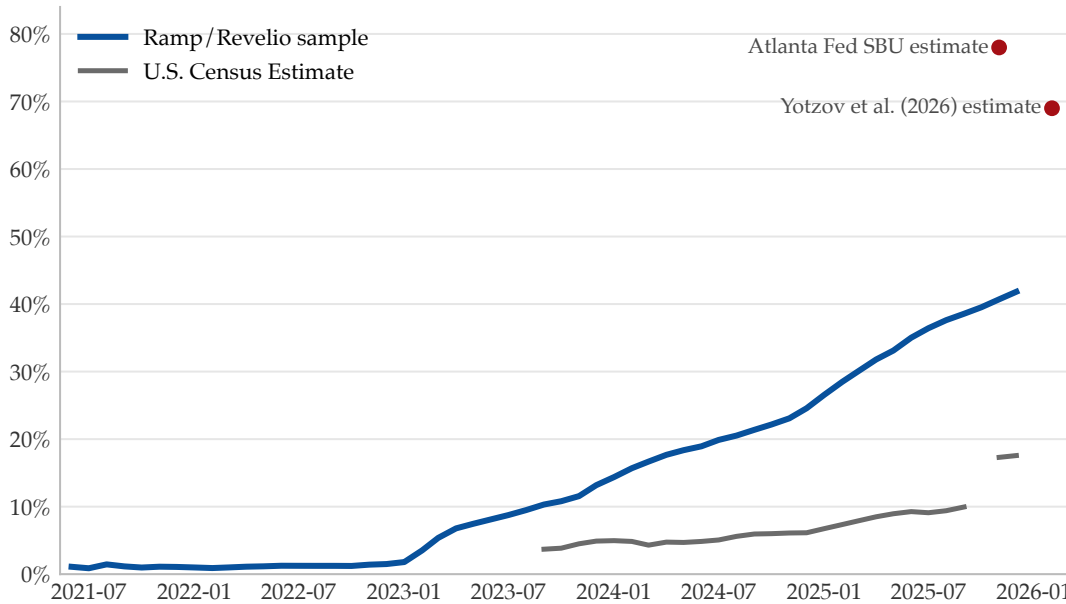
### 3.2 Benchmarking AI Adoption Rates

AI adoption is difficult to measure consistently because different data sources capture different margins of use. Figure 1 benchmarks the adoption rate in the linked Ramp–Revelio analytical panel against public survey estimates. The Ramp–Revelio series is not a nationally representative estimate of informal AI use. It measures the share of active in-sample firms that have crossed the paper’s sustained paid AI adoption threshold, which is closest to a revealed paid-use measure. While our sample likely skews toward tech-forward firms and knowledge work, we likely underestimate actual adoption relative to findings in recent executive surveys.

The comparison highlights why “AI adoption” does not have a single empirical rate. The U.S. Census Bureau’s biweekly Business Trends and Outlook Survey (BTOS) provides a high-frequency, nationally representative business survey (U.S. Census Bureau, 2026b). For most of our overlapping period, its AI question asked whether a business used AI in producing goods or services; in November 2025, Census broadened the wording to AI use in any business function (U.S. Census Bureau, 2026a). Under the revised question, Census reports that overall AI usage hovered between 17 and 20 percent from December 2025 to May 2026 (U.S. Census Bureau, 2026a). In the 2026 BTOS AI supplement, Bonney et al. (2026) report that 18 percent of firms, or 32 percent on an employment-weighted

basis, used AI in at least one business function during November 2025–January 2026.

**Figure 1: AI Adoption Benchmarks**



Notes. The Ramp–Revelio series reports paid-AI adoption in the linked analytical panel. Public benchmarks come from Census BTOS, executive surveys in [Yotzov et al. \(2026\)](#), and the Atlanta Fed Survey of Business Uncertainty ([Federal Reserve Bank of Atlanta, 2026](#)).

Other survey evidence points to much higher adoption. [Yotzov et al. \(2026\)](#) survey senior executives in the United States, United Kingdom, Germany, and Australia and report that about 69 percent of firms actively use AI. In the U.S. Survey of Business Uncertainty, fielded by the Federal Reserve Bank of Atlanta, the corresponding U.S. estimate is about 78 percent in November 2025 ([Yotzov et al., 2026](#); [Federal Reserve Bank of Atlanta, 2026](#)). These estimates are closer to the Ramp–Revelio paid-adoption series than the firm-weighted BTOS benchmark, but they still measure a different object. Executive surveys capture broad organizational use and put more weight on respondents with firm-wide information; Ramp–Revelio captures adoption that leaves a payment trace among firms linked to workforce records; BTOS is broad and nationally representative but sensitive to question wording, firm weighting, and response rate bias.

We therefore interpret Ramp AI spend as a revealed-adoption measure for a selected, business-spend-active firm population. Its value for this paper is not that it reproduces

every public adoption estimate. Its value is that it observes the timing and intensity of firm-level paid AI adoption directly, which is the variation required for our employment analysis.

### 3.3 Revelio Workforce Data

We measure employment outcomes using worker-level career histories from Revelio Labs. Revelio aggregates information from millions of publicly available online professional profiles and uses these records to construct longitudinal employment histories for individual workers.

Each profile contains information on employers, job titles, employment dates, locations, and other career characteristics. By aggregating these worker-level records, we construct monthly measures of firm employment and workforce composition.

We use Revelio’s seven-point seniority field to define entry-level headcount as seniority 1–2, non-entry headcount as seniority 3–7, and manager-plus headcount as seniority 4–7.

A key advantage of these data is their granularity. Because employment histories are observed at the individual level, the data can be used to measure not only total firm employment but also hiring, attrition, occupational composition, worker seniority, and other workforce characteristics that are difficult to observe in many traditional data sources.

Coverage in online professional profiles varies across workers, occupations, industries, and geographies. To address these differences, Revelio applies a weighting methodology that adjusts for differential platform usage and improves representativeness across the labor market. Prior validation exercises have shown that weighted Revelio measures closely track benchmark labor market statistics, including Revelio Public Labor Statistics (RPLS) ([Revelio Labs, 2026](#)). Similar worker-level employment histories have been used extensively in academic research to construct firm-level measures of turnover, hiring, workforce composition, and employee mobility ([Li et al., 2022](#); [Benson and Shaw,](#)

2025; Labaschin et al., 2025).

For this study, we aggregate worker-level employment records to the firm-month level and use these measures to track changes in employment following AI adoption. We examine both total employment and employment within specific occupational groups to understand how workforce composition evolves after adoption.

### 3.4 Matching and Analytical Sample

We link Ramp businesses to Revelio companies by exact domain matching. Candidate domains come from Ramp customer records, the account-resolution layer, and alias domains. When multiple Revelio entities share a domain, we choose one match using normalized name similarity and firm size, then impose two-sided deduplication so each Ramp business maps to at most one Revelio company and each Revelio company maps to at most one Ramp business.

A firm becomes qualified for analysis once it has at least six consecutive months of total Ramp spend of \$5,000 or more. Qualification is absorbing. A firm-month is marked in-sample when the firm is qualified, current-month total Ramp spend is at least \$5,000, and weight-adjusted headcount is at least five. The estimator uses only firm-month observations with this in-sample flag equal to one, so all treatment and comparison observations in the estimating equations correspond to active qualified Ramp/Revelio firm-months.

Firms are excluded from the main intensity analysis if AI adoption occurs before enough Ramp history is observed to form a clean pre-treatment period, or if the first three post-adoption months are not yet observed. Operationally, treated firms receive a low or high AI-intensity assignment only if  $G_i$  is at least three months after the firm's first qualified month and  $G_i + 2$  is observed in the panel.

Table 1 reports baseline summary statistics by treatment group for the final estimating sample.

**Table 1:** Sample Summary by AI Intensity

	AI Intensity		
	Never	Low	High
Number of firms	15,926	3,969	1,664
Headcount (mean)	103.9	193.4***	26.7***
Headcount, Engineering (mean)	20.6	52.8***	7.7***
Headcount, Entry-Level (mean)	65.5	100.0**	9.9***
Headcount Growth, YoY (median %)	+1.6	+5.4	+8.9
Mean Salary (\$)	93,847	103,916***	125,683***
Tech-Adjacent Share (%)	25.9	50.2***	63.8***
Monthly AI Spend per Employee (\$)	–	2.78	33.67
Months Post-Treatment (mean)	–	14.7	13.4

Notes: Treated firms are measured in the month before sustained AI adoption; never-treated firms are measured in December 2024. Low contains the bottom two terciles of monthly AI spend per employee, and High contains the top tercile, using baseline headcount from  $g_i - 1$  as the denominator. Headcount growth is the median 12-month percent change ending at baseline. Tech-adjacent firms are in Information, Finance and Insurance, or Professional, Scientific, and Technical Services. Stars on mean rows report Welch’s t-test against the *Never* column: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 4 What Drives AI Adoption?

AI adoption is highly selected across firms. Before describing our methodology, we first describe which firms adopt AI in the linked Ramp–Revelio panel. This exercise is descriptive rather than causal: the goal is to characterize the firms that select into adoption, and to make transparent the baseline differences that motivate our preferred specification and takeaways.

For this descriptive section, we use our main analytical sample. Treated firms are observed in the month immediately before AI adoption, and never-treated firms are observed in December 2024. Table 2 collapses the treated firms into a single adopter group.

Adopters are more technical, higher-paying, and more likely to be venture-backed. Categorizing firms by AI intensity reveals further differences. For example, small firms are less likely to adopt AI, but when they do, they adopt it more intensely. They are also more concentrated in tech-adjacent sectors and have higher mean salaries. These differences are useful economically, but they also make clear why a simple adopter-versus-

never-adopter comparison is not enough for the main causal design.

**Table 2:** Adopter Characteristics

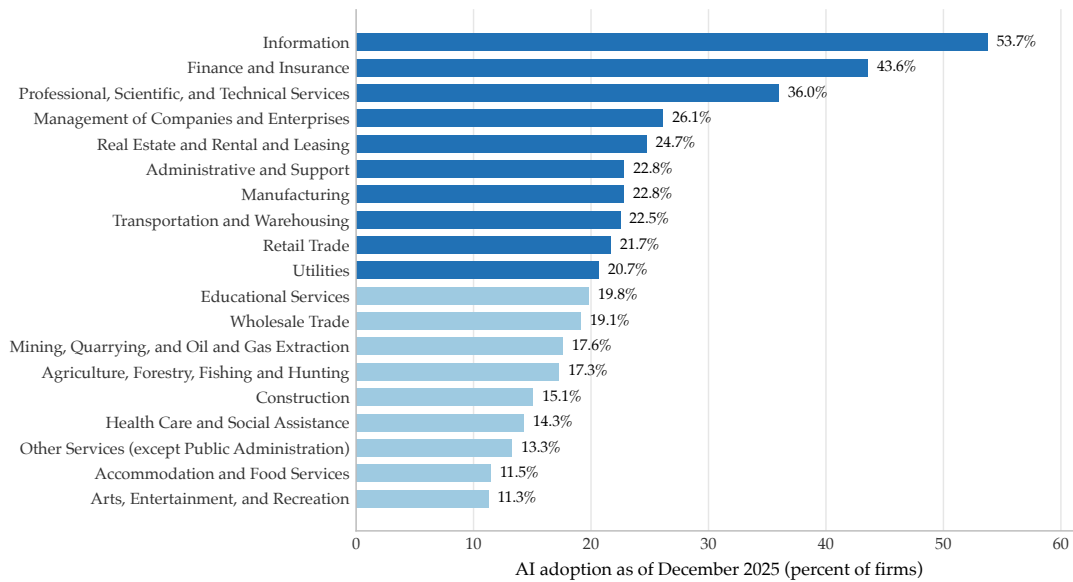
	Never treated	AI adopters	Difference
Number of firms	15,926	5,633	–
Headcount, mean	103.9	144.2	+40.3
Headcount, median	26.8	46.8	+20.0
Headcount growth, YoY (median %)	1.6	6.0	+4.4
Mean salary (\$)	93,847	110,346	+16,499
Engineering share (%)	18.6	26.8	+8.2
Entry-level share (%)	50.4	42.1	-8.3
Admin share (%)	16.8	11.9	-5.0
Sales share (%)	25.6	24.2	-1.4
Tech-adjacent sector share (%)	25.9	54.2	+28.3
VC-backed share (%)	9.5	34.0	+24.5
PE-backed share (%)	6.0	9.4	+3.4

Notes: AI adopter firms are measured in the month before sustained AI adoption; never-treated firms are measured in December 2024.

Sector composition is one of the clearest predictors of adoption. Figure 2 shows that AI adoption measured as of December 2025 is highest in Information, Finance and Insurance, and Professional, Scientific, and Technical Services. It is much lower in construction, health care, other services, arts and entertainment, and accommodation and food services. These differences align with a task-based interpretation of early AI demand: firms whose workforces are more concentrated in software, analytical, professional, and information-processing work are more likely to find immediate uses for generative AI tools.

Firm scale and engineering intensity also predict adoption. In Figure 3, the sustained-adoption rate rises from 12 percent among firms below 10 workers to about 41–43 percent among firms above 250 workers. We observe high levels of adoption at firms without an engineering presence too, but adoption is much more common when a firm has a technical workforce. Firms with no measured engineering headcount adopt at a rate of about 9 percent; adoption rises with engineering share and peaks near 36 percent among firms where engineers represent 30–50 percent of employment. At higher shares, AI adoption falls slightly. Our best guess is that AI adoption is especially common among mixed

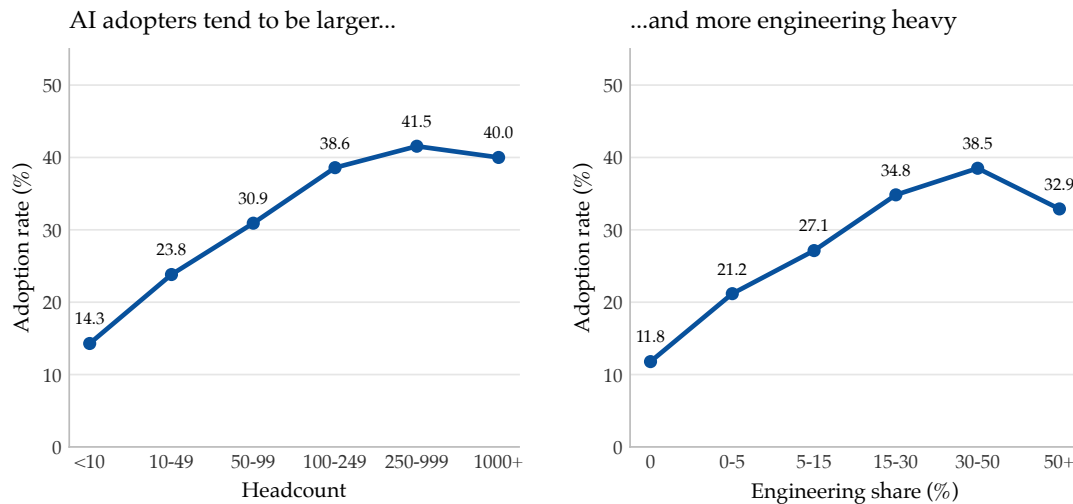
**Figure 2: AI Adoption by Sector**



Notes. The sample is the firm-level baseline sample used in Table 2; sectors with fewer than 100 firms are omitted. Adoption is measured as of December 2025.

workforces with enough technical capacity to evaluate and integrate AI tools, but enough non-engineering work to drive experimentation and diffusion across functions.

**Figure 3: AI Adoption by Size and Engineering Share**

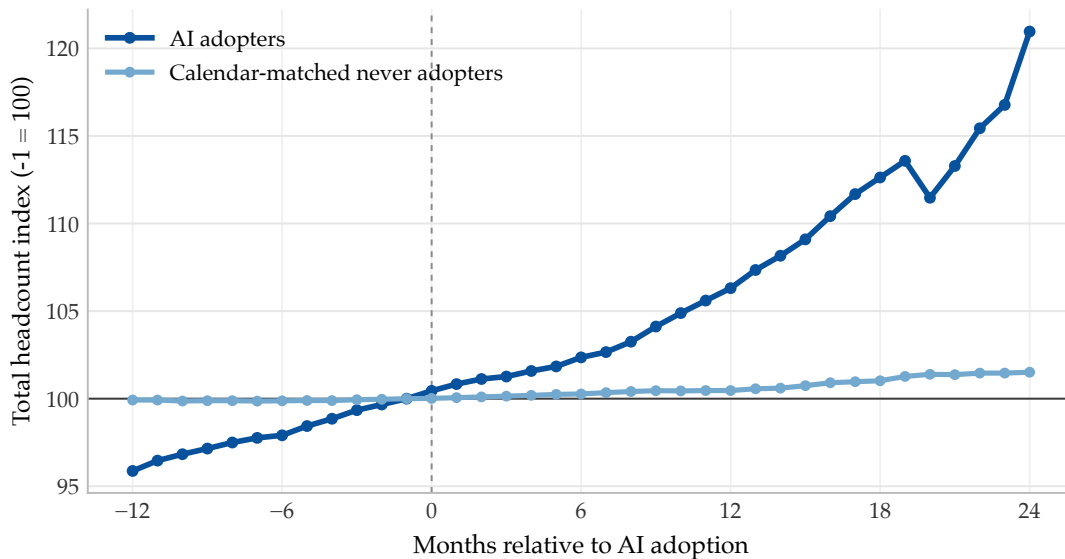


Notes. The sample is the firm-level baseline sample used in Table 2.

## 4.1 Pre-Adoption Growth and the Preferred Comparison

These selection patterns show up directly in employment trajectories. Figure 4 compares indexed total headcount for AI adopters and calendar-matched never-adopters over the full event window. The series are already separated before adoption: adopters are on a faster growth path, while never-adopters are comparatively flat. A design that uses never-adopters as the counterfactual would therefore confound AI adoption with pre-existing differences in firm growth.

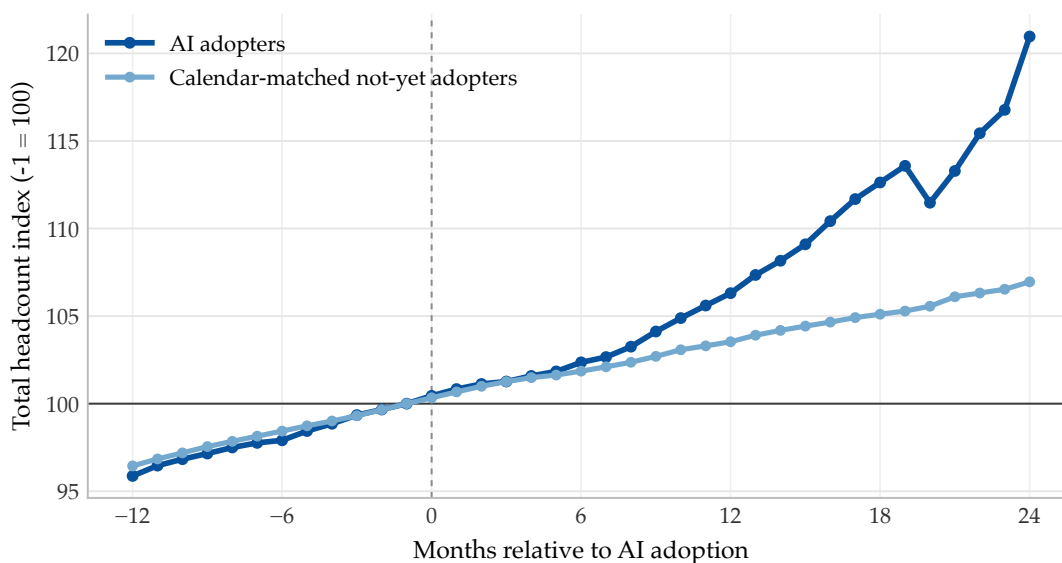
**Figure 4:** Indexed Headcount and Pre-Trends: AI Adopters vs. Never Adopters



Notes. The figure plots mean total headcount for sustained adopters and calendar-matched never-adopter firms. For each adoption cohort, both series are indexed to 100 in event month  $-1$  and then averaged across cohorts using treated cohort firm counts as weights.

The preferred specification therefore uses later adopters as controls while they are not yet treated. Figure 5 shows the corresponding indexed total-headcount paths for all adopters and for later adopters that have not yet adopted in the matched calendar month, normalized to 100 in the month before adoption. In the pre-treatment window, the two series track closely. This alignment supports the parallel-trends assumption needed for the Callaway–Sant’Anna design: the comparison is between earlier and later adopters, not between adopters and firms that never adopt.

**Figure 5:** Indexed Headcount and Pre-Trends: AI Adopters vs. Not-Yet Adopters



Notes. The figure plots mean total headcount for sustained adopters and later adopters that have not yet adopted in the matched calendar month. For each adoption cohort, both series are indexed to 100 in event month  $-1$  and then averaged across cohorts using treated cohort firm counts as weights.

The takeaway from this section is that AI adoption is highly selected, and that selection is visible before treatment when adopters are compared with never-adopters. Later adopters provide a more appropriate comparison because their pre-adoption headcount paths are much more closely aligned. The methodology and results below therefore place greatest weight on the not-yet-treated design, estimated within eventual AI-intensity group and with sector fixed effects. Never-treated firms remain useful for describing who adopts AI and for showing how adopters differ from permanent non-adopters, but they are not the preferred counterfactual.

## 5 Empirical Methodology

We estimate the relationship between firm-level AI adoption and employment using the staggered-adoption difference-in-differences framework of [Callaway and Sant'Anna \(2021\)](#). The unit of observation is a firm-month. For firm  $i$ , let  $G_i$  denote the first month of sus-

tained AI adoption, defined as the first month in the earliest three-month spell in which the firm records at least \$100 of AI vendor spend in each month. Treatment is absorbing after  $G_i$ . Estimation is conducted on active, qualified Ramp firm-months matched to Revelio workforce records.

## 5.1 Treatment Intensity

We classify treated firms into low and high groups of early AI adoption intensity. For treated firm  $i$ , intensity is measured as monthly AI vendor spend per employee over the first three months after sustained adoption begins:

$$PEPM_i = \frac{\sum_{m=0}^2 AI_{i,G_i+m}}{3 \times HC_{i,G_i-1}},$$

where  $AI_{it}$  is AI vendor spend and  $HC_{i,G_i-1}$  is baseline headcount in the month before adoption. Low combines the bottom two terciles of  $PEPM_i$ , and High is the top tercile among treated firms with at least three months between first Ramp qualification and AI adoption, a positive baseline headcount denominator, and all three post-adoption months observed.

We tested several metrics to determine AI intensity, including total spend, share of total spend, number of AI vendors used, and others. We selected the PEPM-based metric for three reasons. First, it correlated closely to adoption of more complex AI products and tools, like coding agents and APIs, as opposed to simpler chat enterprise subscriptions. Second, it allowed us to definite intensity close to treatment timing rather than using the full realized post-treatment spending path (we found the first three months correlated closely to realized spending two years out). Third, using baseline headcount in the denominator keeps the employee scale predetermined relative to adoption and makes intensity comparable between firms with very different headcount levels.

## 5.2 Estimator and Identification

In Section 4, we showed that never-adopter firms differ sharply from eventual adopters before adoption. We therefore compare adopters to later adopters while those later adopters have not yet adopted. To keep the comparison local to firms with similar eventual AI spending intensity, we estimate each intensity group separately: Low adopters are compared to later adopters whose eventual intensity group is Low while they are not yet treated, and High adopters are compared analogously to later adopters whose eventual intensity group is High.

For outcome  $Y_{it}$ , treatment cohort  $g$ , calendar month  $t$ , and intensity group  $k$ , the group-time average treatment effect is

$$ATT_k(g, t) = \mathbb{E} [Y_{it}(g) - Y_{it}(0) \mid G_i = g, i \in k],$$

where  $Y_{it}(g)$  is the potential outcome for firm  $i$  in month  $t$  if first treated in month  $g$ , and  $Y_{it}(0)$  is the potential outcome if the firm remains untreated through month  $t$ . In the preferred not-yet-treated design,  $Y_{it}(0)$  is identified using firms in the same eventual intensity group whose own adoption month is later than  $t$ .

The identifying assumption is conditional parallel trends within eventual intensity group: absent earlier AI adoption, treated firms in a given cohort and intensity group would have followed the same employment trajectory as later adopters in the same intensity group after accounting for broad sector composition. The event-study estimates aggregate group-time effects by relative month  $e = t - g$  over the window  $e \in [-12, 24]$ . The main scalar estimate in the results tables is the simple post-treatment aggregate ATT across post-adoption months in this window.

Our preferred specification includes NAICS sector fixed effects and no other controls. Sector adjustment addresses the most transparent source of selection into early AI adoption: adopters are disproportionately concentrated in information, finance, and

professional services. We do not include additional growth or firm-level controls in the preferred specification because earlier specifications showed little improvement in pre-treatment diagnostics after adding them.

### 5.3 Outcomes and Diagnostics

For total headcount, the dependent variable is log headcount. For role, education, and entry-level outcomes, we use  $\log(1 + Y_{it})$ , which preserves firm-months with zero employment in a category. Estimates are reported in log points; they can be read approximately as percent changes.

The paper outcomes are total headcount, total entry-level headcount, engineering, entry-level engineering, sales, marketing, administrative, finance, scientist, operations, customer service, bachelor's-degree, MBA, JD, and PhD headcount.

We also estimate workforce-composition outcomes as shares of total firm headcount: entry-level, engineering, entry-level engineering, sales, administrative, marketing, finance, scientist, operations, customer service, bachelor's-degree, MBA, JD, and PhD shares. These outcomes help distinguish broad employment growth from changes in the composition of the workforce.

We assess pre-treatment fit using universal-base event-study coefficients for  $e < 0$ , anchored at the month immediately before adoption. These pre-period estimates are our main diagnostic for conditional parallel trends. The  $e = -1$  period is normalized to zero, so pre-period flag counts use event times  $e = -12$  through  $e = -2$ . Pre-period flags are low or zero across most outcomes, providing confidence for our preferred specification.

### 5.4 Comparison-Group Checks

We also report descriptive and robustness evidence using never-adopter firms. These comparisons are useful for showing that AI adopters differ from permanent non-adopters,

but we do not treat them as the preferred comparison given their divergent pre-trends. The main analysis therefore places greatest weight on the not-yet-treated, within-intensity estimates, while using never-adopting comparisons to characterize selection and the broader adopter/non-adopter contrast.

## 6 Results

### 6.1 High-Intensity AI Adopters Grew Headcount by Ten Percent

Table 3 reports the preferred estimates over the first 24 months after sustained AI adoption. Each column compares adopters in a low- or high-AI intensity group to firms in the same eventual intensity group that have not yet adopted. All specifications include NAICS sector fixed effects. Figure 6 shows the total-headcount result visually.

High-intensity AI adopters expand employment following AI adoption. The coefficient on High vs. Not-yet for total headcount is 0.097, meaning high-AI intensity firms grew headcount 10.2 percent in the two years after adoption. The Low estimate is  $-0.006$ , or -0.6 percent, and is not statistically significant. Pre-period diagnostics are clean for Low, with 0 of 11 flagged pre-periods, but less clean for High, with 3 of 11 flagged pre-periods.

The headline result therefore shows that not all firms show the same gains from AI. Low-AI adopters did not increase total headcount, did not increase entry-level headcount, and did not show the broad functional expansion observed among high-PEPM adopters. In Table 1, Low adopters average \$2.78 per employee per month in first-three-month AI spend, while High adopters average \$33.67. The gains appear only among firms that make sustained, material investments, likely going beyond simple chat subscriptions to more advanced coding agents, API usage, and adoption of multiple models or vendors.

**Table 3: Headcount by AI Intensity**

Outcome	Average ATT (months 0–24)		Pre-period flags (lower is better)	
	Low vs. Not-yet	High vs. Not-yet	Low	High
<b>Headline Outcomes</b>				
Total headcount	−0.006 (0.016)	0.097*** (0.022)	0/11	3/11
Total entry-level headcount	−0.017 (0.018)	0.113*** (0.030)	0/11	2/11
Non-entry headcount	0.004 (0.016)	0.074*** (0.023)	0/11	1/11
Manager-plus headcount	0.010 (0.015)	0.065*** (0.020)	0/11	0/11
<b>Job Category Outcomes</b>				
Engineering headcount	−0.004 (0.023)	0.070** (0.031)	1/11	1/11
Entry-level engineering headcount	−0.012 (0.023)	0.061** (0.026)	0/11	0/11
Sales headcount	−0.005 (0.017)	0.098*** (0.027)	1/11	0/11
Marketing headcount	0.014 (0.018)	0.055* (0.029)	0/11	3/11
Admin headcount	−0.001 (0.019)	0.075*** (0.021)	0/11	1/11
Finance headcount	−0.012 (0.018)	0.045* (0.024)	0/11	0/11
Scientist headcount	0.032* (0.018)	0.054*** (0.017)	0/11	0/11
Operations headcount	0.000 (0.016)	0.022 (0.021)	0/11	0/11
Customer service headcount	−0.020 (0.019)	0.061** (0.025)	0/11	0/11
<b>Education Outcomes</b>				
Bachelors headcount	0.010 (0.017)	0.098*** (0.026)	0/11	4/11
MBA headcount	0.017 (0.017)	0.049** (0.021)	0/11	5/11
JD headcount	0.008 (0.010)	0.012 (0.011)	0/11	4/11
PhD headcount	0.040*** (0.014)	0.024 (0.015)	2/11	0/11

Notes: Each column compares one low- or high-intensity AI adoption group to firms in the same first-three-month PEPM intensity group that have not yet adopted by that event time. Low contains the bottom two PEPM terciles, while High contains the top PEPM tercile. The Callaway–Sant’Anna doubly robust estimator is run on active qualified Ramp/Revelio firm-months and includes NAICS sector fixed effects. Standard errors are in parentheses. Pre-period flags use universal-base CS diagnostics and count non-normalized pre-treatment event times  $e = -12$  through  $e = -2$  whose pointwise 95 percent confidence intervals exclude zero;  $e = -1$  is the normalized base period. Coefficients are log points for total headcount and  $\log(1 + y)$  points for role, education, seniority, and entry-level outcomes. Stars denote \* $p < 0.10$ , \*\* $p < 0.05$ , and \*\*\* $p < 0.01$ .

## 6.2 High-Intensity Employment Gains Compounded Over Time

The total-headcount event study shows that high-intensity employment gains may follow a learning curve in firms' use of AI. In Figure 6, the high-intensity coefficient is 0.003 in the adoption month and 0.020 three months after adoption. It then rises to 0.071 by month 6, 0.188 by month 12, and 0.277 by month 18. The high-intensity ATT reaches 0.277 log points by month 18, about 32 percent; the month-24 ATT is larger, 0.452 log points, about 57 percent, but comes from the late event-study window with wider confidence intervals.

The earliest signs of growth emerge roughly 6-12 months after adoption, perhaps after firms establish best practices, integrate AI tools into workflows, and are subsequently able to make new investments and hire staff. Later post-adoption estimates continue to rise, though the late-window confidence intervals are wider as fewer firms are observed that far after adoption.

## 6.3 Entry-Level Employment Also Rose Among High-Intensity Adopters

Entry-level employment also rises among high-intensity adopters. In Table 3, the high-intensity coefficient for entry-level headcount is 0.113, or a 12.0 percent increase after adoption, and statistically significant at the 1 percent level. Low-intensity adopters do not show a statistically significant change in entry-level headcount. Figure 7 shows the entry-level result visually. These estimates counter the view that intensive AI adoption is already reducing junior employment at adopting firms.

Table 4 reports the preferred design for workforce-composition shares. These estimates ask whether AI adoption changes the mix of workers inside firms, rather than firm scale. Entry-level share falls by 0.52 percentage points for Low adopters and rises by 1.15 percentage points for High adopters. Because non-entry is the complement of entry-level, its share moves mechanically in the opposite direction. Manager-plus share rises by 0.66 percentage points for Low adopters and falls by 1.52 percentage points for High adopters.

The entry-level result is our earliest evidence that high-intensity adopters end up hiring different kinds of employees. Entry-level headcount expands enough relative to the rest of the firm that entry-level workers make up a higher share of employment at the end of the 24-month analysis period.

**Table 4: Workforce Shares by AI Intensity**

Outcome	Average ATT (months 0–24)		Pre-period flags (lower is better)	
	Low vs. Not-yet	High vs. Not-yet	Low	High
<b>Headline Outcomes</b>				
Entry-level share	−0.520** (0.262)	1.147** (0.579)	0/11	0/11
Non-entry share	0.522* (0.287)	−1.142* (0.591)	0/11	0/11
Manager-plus share	0.655** (0.269)	−1.517** (0.617)	0/11	0/11
<b>Job Category Outcomes</b>				
Engineering share	−0.169 (0.331)	−0.552 (0.611)	1/11	0/11
Entry-level engineering share	−0.218 (0.214)	0.015 (0.386)	0/11	0/11
Sales share	−0.034 (0.249)	0.286 (0.532)	0/11	1/11
Marketing share	0.258 (0.289)	−0.358 (0.523)	0/11	0/11
Admin share	−0.140 (0.149)	0.401 (0.316)	1/11	0/11
Finance share	−0.219 (0.198)	0.052 (0.389)	1/11	0/11
Scientist share	0.325** (0.143)	0.280 (0.287)	0/11	0/11
Operations share	−0.021 (0.103)	−0.109 (0.234)	2/11	0/11
Customer service share	−0.274** (0.137)	0.501* (0.268)	1/11	3/11
<b>Education Outcomes</b>				
Bachelors share	0.762*** (0.285)	0.380 (0.564)	1/11	0/11
MBA share	0.207** (0.100)	−0.033 (0.256)	0/11	2/11
JD share	0.039 (0.041)	0.091 (0.118)	0/11	0/11
PhD share	0.197* (0.103)	0.311 (0.209)	2/11	0/11

Notes: Share outcomes divide role, seniority, entry-level, or education-group headcount by total firm headcount in the same month; total headcount share is omitted because it is mechanically one. Coefficients and standard errors are reported in percentage points. Each column compares one low- or high-intensity AI adoption group to firms in the same first-three-month PEPM intensity group that have not yet adopted by that event time, using the preferred sector-adjusted CS specification. Low contains the bottom two PEPM terciles, while High contains the top PEPM tercile. Pre-period flags use universal-base CS diagnostics and count event times  $e = -12$  through  $e = -2$ ;  $e = -1$  is the normalized base period. Stars denote  $*p < 0.10$ ,  $**p < 0.05$ , and  $***p < 0.01$ .

## 6.4 Employment Gains Were Broad Across Roles

Table 3 also shows that growth is broad across job functions. High-intensity adopters have 10.3 percent higher sales headcount, 7.8 percent higher administrative headcount, 7.3 percent higher engineering headcount, 6.3 percent higher entry-level engineering headcount, 6.3 percent higher customer-service headcount, and 5.6 percent higher scientist headcount. Finance headcount rises by 4.6 percent and marketing headcount rises by 5.7 percent, with both estimates marginally significant; marketing also has weaker pre-period diagnostics than most other high-intensity role estimates. Operations is the only job category in which high-intensity firms did not grow headcount.

Low-intensity adopters show mixed or minimal gains across roles. The few negative low-intensity estimates are small and not statistically significant. Finance headcount falls by 1.2 percent and customer-service headcount falls by 2.0 percent for low-intensity adopters. The positive low-intensity estimates are narrow: scientist headcount rises by 3.2 percent at the 10 percent level, and PhD headcount rises by 4.1 percent.

Education outcomes point in the same direction as the occupational results, but with more diagnostic caution. Bachelor's headcount rises by 10.3 percent and MBA headcount rises by 5.0 percent for high-intensity adopters, though both outcomes have several flagged pre-periods. JD headcount is essentially unchanged. PhD headcount is positive but not statistically significant in the high-intensity group.

Our best interpretation is that high-intensity adopters are expanding firm scale across multiple functions, with the clearest average gains in total headcount, entry-level headcount, sales, administration, engineering, customer service, and bachelor's-degree workers. There does not appear to be a strong correlation between a function's perceived AI exposure and its increase or decrease in headcount.

## 6.5 Early Employment Gains Are Concentrated in Information Firms

Table 5 reports the results of our preferred specification separately by broad sector.

The gains from AI adoption are unevenly distributed across sectors. We find that these gains are concentrated in Information and, at present, are statistically significant only in that sector group, which includes many software, internet, and media firms. High-intensity adopters in Information grow total headcount by 0.126, or about 13.4 percent, with no flagged pre-periods. The Low estimate in Information is small, negative, and not statistically significant. Professional and technical services also has a positive high-intensity estimate, but it is smaller and not statistically significant. The remaining sector groups are close to zero, imprecise, or both.

The result is consistent with our expectations given the uneven distribution of AI adoption. AI adoption is already uniquely concentrated in Information, and the most commercially mature AI gains appear clearest in coding-agent and software-engineering workflows. For software and technology firms, AI can make core output cheaper or faster to produce: writing code, debugging, building internal tools, producing technical documentation, and supporting product development. Lower production costs in these workflows can raise the return to expanding the whole firm, not just the engineering team.

The sector evidence should not be construed to suggest that AI is only meaningful for tech firms or that other industries will never show gains. It is still early in the AI adoption curve. AI use cases may diffuse more slowly outside software and media, especially where deployment requires workflow redesign and in non-technical sectors where AI vendors have yet to find product-market fit. We intend to update these estimates with future adoption cohorts and longer post-adoption windows beyond the initial 24 months to track developments across sectors.

**Table 5: Headcount by Sector**

Sector group	Average ATT (months 0–24)		Pre-period flags (lower is better)	
	Low vs. Not-yet	High vs. Not-yet	Low	High
Information	–0.028 (0.044)	0.126** (0.058)	0/11	0/11
Professional and technical services	–0.036 (0.028)	0.069 (0.049)	0/11	3/11
Finance and insurance	0.058 (0.078)	–0.030 (0.080)	0/11	0/11
Industrial, trade, logistics, and resources	0.008 (0.024)	–0.005 (0.044)	0/11	0/11
Health, education, and public services	–0.055 (0.042)	0.052 (0.144)	0/11	3/11
Consumer, administrative, real estate, and other services	0.015 (0.029)	–0.012 (0.041)	0/11	0/11

Notes: Each row estimates the preferred total-headcount Callaway–Sant’Anna design within the listed broad NAICS-sector group. Columns compare low- and high-intensity first-three-month PEPM adopters with not-yet-treated firms in the same PEPM intensity group and sector group. Low contains the bottom two PEPM terciles, while High contains the top PEPM tercile. Specifications include available NAICS sector fixed effects and no additional controls. Standard errors are in parentheses. Dashes indicate cells without enough treated or comparison-firm support after sector stratification. Pre-period flags use universal-base CS diagnostics and count event times  $e = -12$  through  $e = -2$  whose point-wise 95 percent confidence intervals exclude zero;  $e = -1$  is the normalized base period. Stars denote  $*p < 0.10$ ,  $**p < 0.05$ , and  $***p < 0.01$ .

## 6.6 Gains Are Even Higher Against Never-Adopters

Table 6 reports the same intensity design against never-treated firms. The total-headcount estimates are larger than in the preferred not-yet-treated specification and higher for High than for Low. The ATT is 0.063 for low-intensity adopters and 0.118 for high-intensity adopters, implying headcount gains of about 6.5 percent and 12.6 percent, respectively.

These estimates should be interpreted cautiously. Adopters were already growing at a faster rate relative to never-adopters: total-headcount pre-period flags are 11 of 11 for low intensity and 8 of 11 for high intensity. We therefore include this specification for completeness rather than as the preferred design.

**Table 6: Never-Treated by AI Intensity**

Outcome	Average ATT (months 0–24)		Pre-period flags (lower is better)	
	Low vs. Never	High vs. Never	Low	High
<b>Headline Outcomes</b>				
Total headcount	0.063*** (0.011)	0.118*** (0.015)	11/11	8/11
Total entry-level headcount	0.043*** (0.012)	0.109*** (0.017)	11/11	0/11
Non-entry headcount	0.072*** (0.009)	0.103*** (0.015)	11/11	9/11
Manager-plus headcount	0.069*** (0.009)	0.093*** (0.014)	11/11	10/11
<b>Job Category Outcomes</b>				
Engineering headcount	0.091*** (0.013)	0.130*** (0.018)	11/11	11/11
Entry-level engineering headcount	0.052*** (0.013)	0.090*** (0.016)	11/11	6/11
Sales headcount	0.081*** (0.012)	0.122*** (0.015)	11/11	5/11
Marketing headcount	0.052*** (0.011)	0.061*** (0.016)	11/11	10/11
Admin headcount	0.051*** (0.013)	0.088*** (0.015)	11/11	3/11
Finance headcount	0.043*** (0.010)	0.042*** (0.014)	10/11	0/11
Scientist headcount	0.055*** (0.010)	0.064*** (0.011)	11/11	5/11
Operations headcount	0.034*** (0.008)	0.029** (0.011)	9/11	0/11
Customer service headcount	0.061*** (0.012)	0.099*** (0.014)	10/11	1/11
<b>Education Outcomes</b>				
Bachelors headcount	0.071*** (0.010)	0.120*** (0.017)	11/11	9/11
MBA headcount	0.050*** (0.010)	0.045*** (0.012)	11/11	5/11
JD headcount	0.005 (0.006)	0.006 (0.007)	0/11	4/11
PhD headcount	0.030*** (0.008)	0.046*** (0.010)	10/11	5/11

Notes: Each column compares one low- or high-intensity first-three-month PEPM adopter group with the common never-treated comparison group. Low contains the bottom two PEPM terciles, while High contains the top PEPM tercile. The specification includes NAICS sector fixed effects and no additional controls. Standard errors are in parentheses. Pre-period flags use universal-base CS diagnostics and count event times  $e = -12$  through  $e = -2$  whose pointwise 95 percent confidence intervals exclude zero;  $e = -1$  is the normalized base period. Coefficients are log points for total headcount and  $\log(1 + y)$  points for role, education, seniority, and entry-level outcomes. Stars denote  $*p < 0.10$ ,  $**p < 0.05$ , and  $***p < 0.01$ .

## 7 Conclusion

This paper finds that AI adoption is associated with higher employment, especially among high-intensity adopters. The preferred results show no evidence that AI adoption is producing layoffs. Instead, firms that make sustained and intensive AI investments expand employment faster than otherwise comparable firms. The gains are visible in total headcount, entry-level work, and AI-exposed functions such as engineering, sales, customer service, and finance. The pattern across intensity groups strengthens our interpretation; the firms using AI more intensively are the firms with the largest employment gains.

These results also raise important questions for the developers of this technology and for policymakers.

First, the gains are concentrated. The clearest sector-level employment gains are in Information, the sector that includes many software, internet, media, and technology firms. These firms are already more likely to adopt AI, and they also appear to be the first firms translating AI adoption into measurable employment growth. Other sectors lag both in adoption and in observed hiring gains. That does not mean AI is irrelevant outside of tech. It means that the realized jobs gains are, at this stage, concentrated where the technology, organizational capacity, and commercial use cases are most developed.

Second, the results raise a product issue. If the largest gains are concentrated in Information, then the next frontier is not simply better models. To realize the broad economic promises of AI, developers need to build tools for firms whose work is not organized around code and whose users are less technical as-is.

Third, the results raise a policy question. AI adoption requires resources, managerial attention, technical capacity, and credible channels for learning. Firms with venture backing, founder networks, or other adoption channels may be better positioned to make those investments. Firms without those channels may fall behind. What role should government play in helping firms make durable, productivity-enhancing investments?

Fourth, the estimates point to a threshold and a learning curve. Enterprise chat sub-

scriptions do not appear to be enough to drive gains. Nor are a few months of experimental spending. This is consistent with a technology whose benefits require complementary investments, organizational change, and learning inside the firm. Many firms may buy subscriptions, run pilots, and then fail to make the sustained investments required to benefit. Developers and policymakers should therefore ask how to support long-term adoption rather than short-term experimentation.

Fifth, the mechanism remains unresolved. We now know that AI adopters, and especially high-intensity adopters, grow faster. We do not yet know which operational practices generate that growth. The relevant mechanisms may include product acceleration, sales productivity, engineering leverage, customer support automation, faster internal analysis, better targeting, or new business lines. But we don't know, and the firms that have figured it out have no incentive to share how they did it. Further research should identify how these fast-growing firms are driving change using AI, and how these results can be replicated across later-adopters.

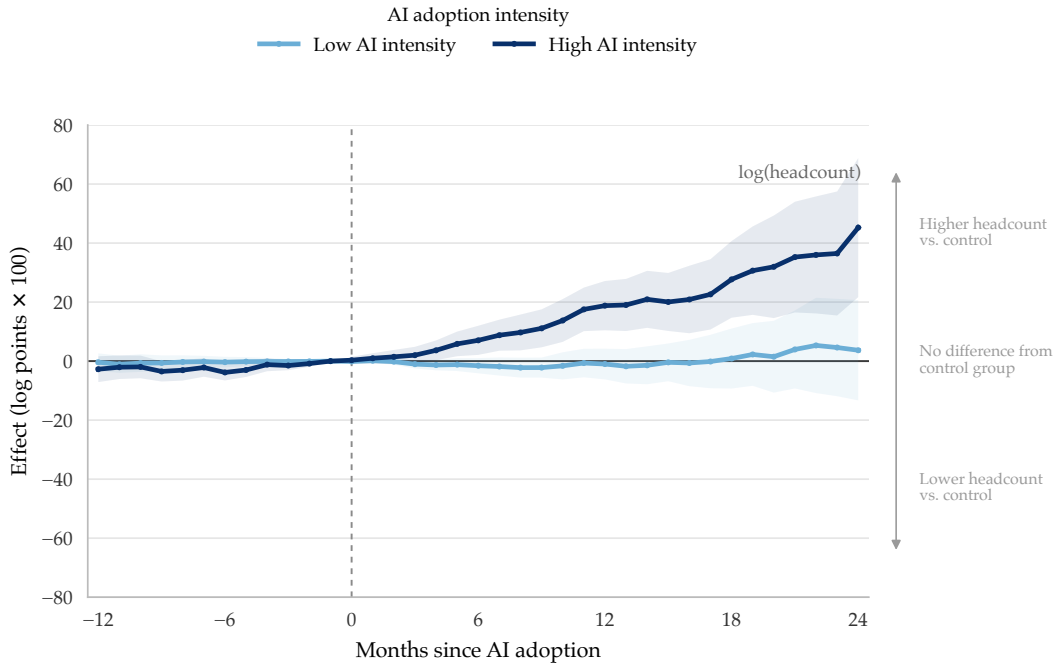
Sixth, AI adoption faces a trust problem. Governments are still deciding how to regulate the technology. Consumers worry about safety and resource use. AI developers alternate between narratives of extraordinary wealth creation and warnings of broad job loss. Some CEOs blame layoffs on AI, even though our firm-level evidence says the opposite. These trust issues are a headwind on adoption that will slow and reduce the ultimate productivity gains AI could deliver.

Finally, here are some takeaways for readers outside of the AI and policy scene. If you are a young person entering the labor market, and you are choosing between otherwise similar firms, choose the one that is using AI. If you are an engineer worried that AI will eliminate engineering jobs, the firms adopting AI are hiring engineers faster, not slower. And if you are reading headlines where CEOs blame layoffs on AI, be skeptical.

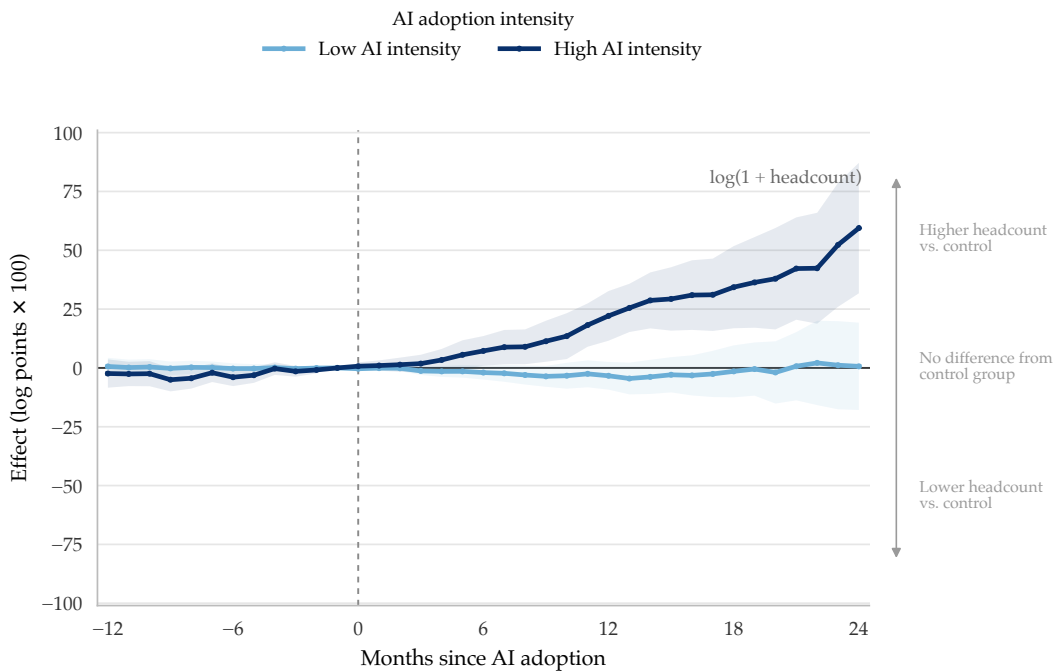
## 8 Event-Study Figures

The figures below plot the preferred-specification dynamic Callaway–Sant’Anna estimates for each paper outcome.


**Figure 6: Total Headcount**




**Figure 7: Entry-Level Headcount**




**Figure 8: Non-Entry Headcount**




**Figure 9: Manager-Plus Headcount**



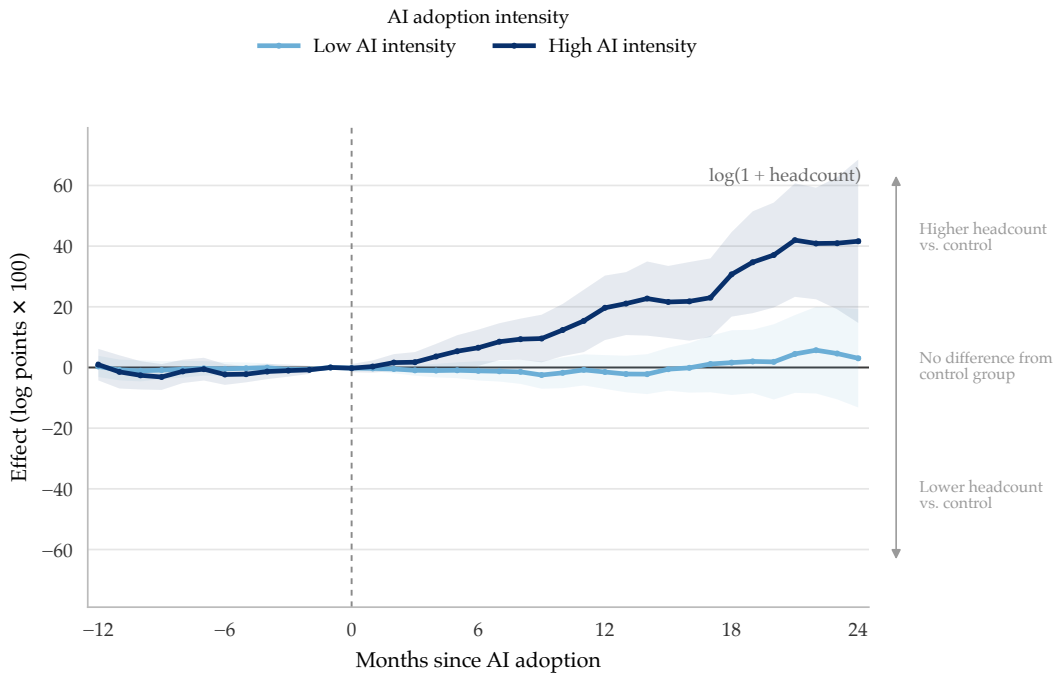
**Figure 10: Engineering Headcount**



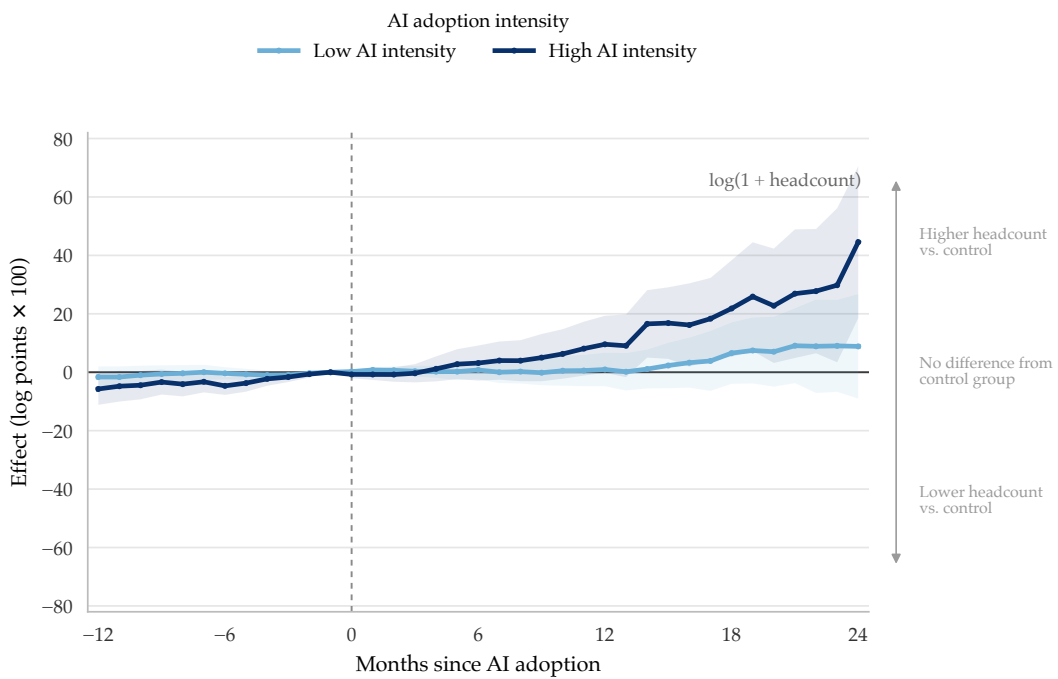
**Figure 11: Entry-Level Engineering**




**Figure 12: Sales Headcount**




**Figure 13: Marketing Headcount**




**Figure 14: Admin Headcount**




**Figure 15: Finance Headcount**




**Figure 16: Scientist Headcount**



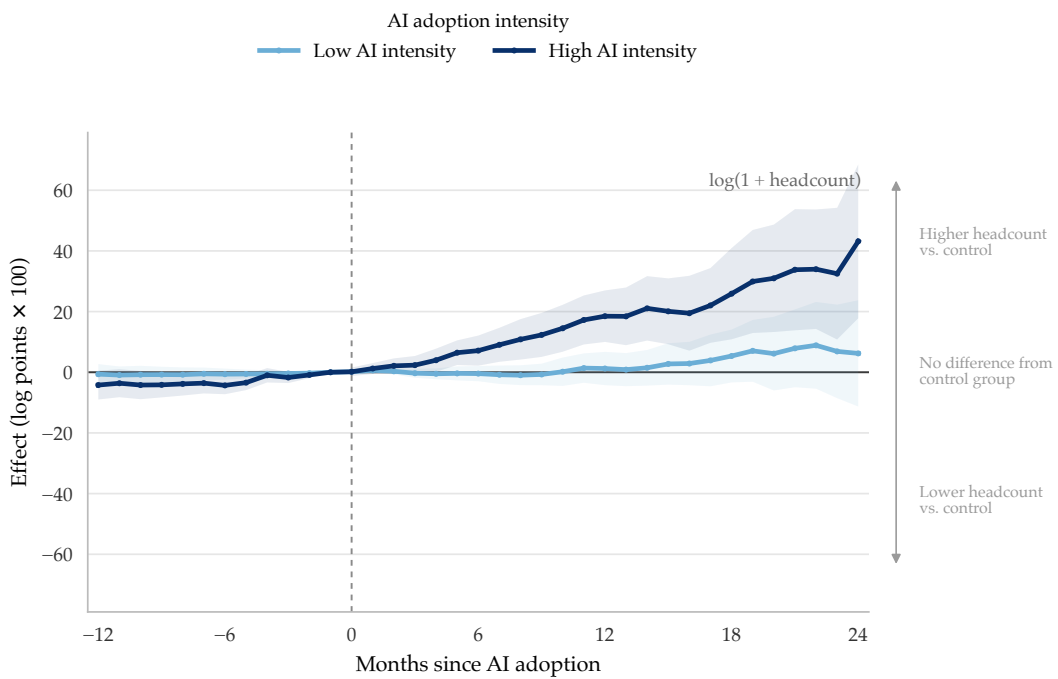
**Figure 17: Operations Headcount**




**Figure 18: Customer Service Headcount**



**Figure 19: Bachelor's Headcount**



**Figure 20: MBA Headcount**



**Figure 21: JD Headcount**


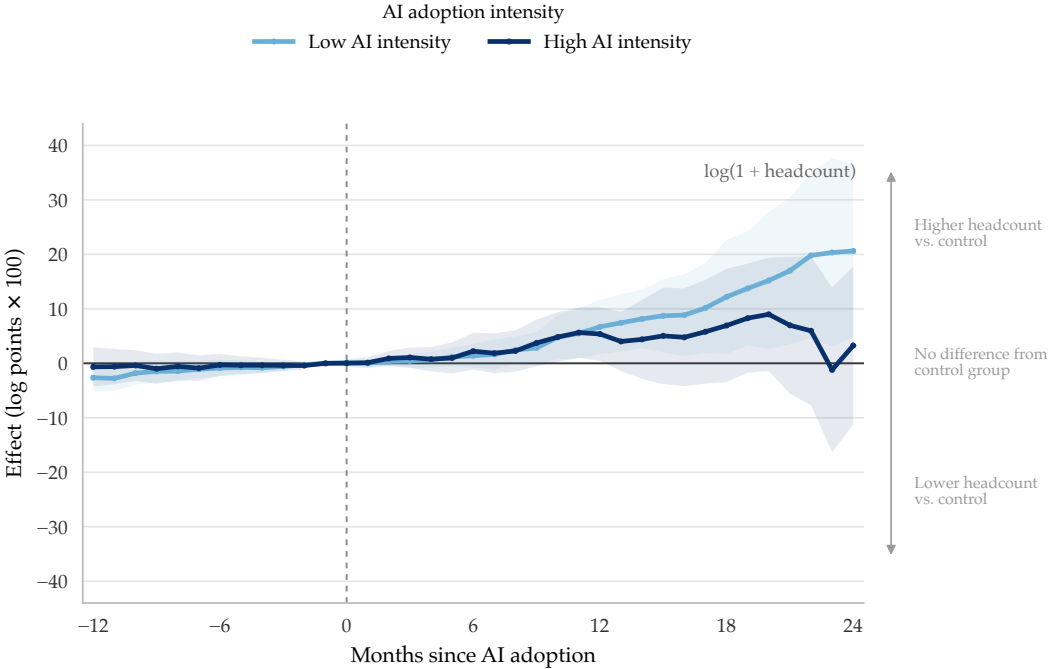


Figure 22: PhD Headcount



## References

- Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics* 40(S1), S293–S340.
- Acemoglu, D. and P. Restrepo (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 P. Restrepo (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives* 33(2), 3–30.
- Autor, D., C. Chin, A. Salomons, and B. Seegmiller (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics* 139(3), 1399–1465.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118(4), 1279–1333.
- Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151, 103745.
- Benson, A. and K. Shaw (2025). What do managers do? an economist’s perspective. *Annual Review of Economics* 17, 635–664.
- Bonney, K., C. L. Breaux, E. Dinlersoz, L. S. Foster, J. C. Haltiwanger, and A. A. Pande (2026). The microstructure of AI diffusion: Evidence from firms, business functions, and worker tasks. NBER Working Paper No. 35141; U.S. Census Bureau CES Working Paper 26-25.
- Brynjolfsson, E., B. Chandar, and R. Chen (2025). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab Working Paper.

- Brynjolfsson, E., D. Li, and L. R. Raymond (2025). Generative AI at work. *The Quarterly Journal of Economics* 140(2), 889–942.
- Callaway, B. and P. H. C. Sant’Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics* 225(2), 200–230.
- de Chaisemartin, C. and X. D’Haultfœuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110(9), 2964–2996.
- Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). GPTs are GPTs: An early look at the labor market impact potential of large language models. *Science* 384(6702), eadl2347.
- Federal Reserve Bank of Atlanta (2026). Survey of business uncertainty.
- Felten, E. W., M. Raj, and R. Seamans (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal* 42(12), 2195–2217.
- Gimbel, M., J. Kendall, and R. Kulsakdinun (2026). Labor market AI exposure: What do we know? The Budget Lab at Yale.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225(2), 254–277.
- Hampole, M., D. Papanikolaou, L. Schmidt, and B. Seegmiller (2025). Artificial intelligence and the labor market. NBER Working Paper No. 33509.
- Kharazian, A. (2026). Ramp AI index. <https://ramp.com/data/ai-index>.
- Labaschin, B., T. Eloundou, S. Manning, P. Mishkin, and D. Rock (2025). Extending ‘GPTs are GPTs’ to firms. In *AEA Papers and Proceedings*, Volume 115, pp. 51–55.

- Li, Q., B. Lourie, A. Nekrasov, and T. Shevlin (2022). Employee turnover and firm performance: Large-sample archival evidence. *Management Science* 68(8), 5667–5683.
- Massenkoff, M. and P. McCrory (2026). Labor market impacts of AI: A new measure and early evidence. Anthropic Research.
- Noy, S. and W. Zhang (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654), 187–192.
- Revelio Labs (2026). Revelio public labor statistics.
- Stevens, R. (2026). Payrolls to prompts: Firm-level evidence on the substitution of labor for AI. arXiv:2602.00139.
- Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2), 175–199.
- U.S. Census Bureau (2026a). AI use at U.S. businesses. America Counts, May 26, 2026.
- U.S. Census Bureau (2026b). Business trends and outlook survey.
- Webb, M. (2020). The impact of artificial intelligence on the labor market. Stanford University Working Paper.
- Yotzov, I., J. M. Barrero, N. Bloom, P. Bunn, S. J. Davis, K. M. Foster, A. Jalca, B. H. Meyer, P. Mizen, M. A. Navarrete, P. Smietanka, G. Thwaites, and B. Z. Wang (2026). Firm data on AI. NBER Working Paper No. 34836.