AI Adoption and Labor Market Responses: Cross-Country Evidence from Job Postings*

Yusuke Aoki † Ahmad Firdaus‡ Takuji Fueki § Koji Takahashi¶

March 2026

Abstract

We study the dynamic impacts of artificial intelligence (AI) adoption on productivity and labor market outcomes in Japan and the United States, using local projection method. We measure AI adoption by the share of AI-related vacancies based on a novel dataset of establishment-level job postings. We find that increases in AI adoption generate sizable productivity gains in both countries. However, labor market responses differ sharply across the two countries. In the United States, these productivity gains are accompanied by significant labor market adjustment costs. AI adoption leads to persistent declines in employment and total compensation growth over the short to medium run. In contrast, in Japan, productivity gains translate into persistent and significant increases in employment and wage growth. These findings imply that labor market adjustment is highly country-specific.

Keywords: Artificial intelligence, Labor market, Investment.
JEL Classification codes: J23, J24

*The views expressed in this paper are those of the authors and do not necessarily reflect the official views of the Bank of Japan. The authors thank Shingo Watanabe, Daisuke Ikeda, Yoshiyuki Nakazono, and the participants at the seminar at the Bank of Japan.
†Indeed Hiring Lab. Email: yaoki@indeed.com
‡Hitotsubashi University. E-mail: ed255114@g.hit-u.ac.jp
§Hitotsubashi University Email: takuji.fueki@r.hit-u.ac.jp
¶Bank of Japan. Email: kouji.takahashi-2@boj.or.jp

1 Introduction

What are the effects of artificial intelligence (AI) adoption on productivity and the labor market? AI is increasingly used across many sectors of the economy. These developments have generated widespread expectations that AI may raise productivity and reshape labor demand, with potentially important implications for wage dynamics. Understanding the consequences of AI adoption on the labor market is therefore of first-order importance for both economists and policymakers.

Theoretical studies suggest that increased AI adoption raises productivity and may reduce employment, while its impacts on wages remain ambiguous (e.g., Acemoglu, 2025, among others). On the other hand, empirical evidence on the macroeconomic effects of AI adoption is still limited, largely due to the lack of comprehensive measures of AI adoption (Babina et al. (2024)).

This study fills this gap by constructing a novel dataset on AI adoption and estimating its dynamic macroeconomic effects using panel local projections. Specifically, we provide new evidence on the effects of AI adoption on productivity and labor market outcomes, including employment and wages.

We use industry-level exposure to AI as a proxy for AI adoption. Industry-level AI exposure is measured using job-posting data from Indeed, a leading online employment platform. Job postings are classified based on standardized keyword-based rules that identify references to AI technologies, including machine learning, natural language processing, and computer vision. The resulting measure captures AI-related labor demand and provides a timely proxy for firms' adoption and deployment of AI.

Focusing on Japan and the US, we quantify the dynamic macroeconomic effects of AI adoption on productivity and the labor market using local projections. Our main findings are as follows.

- AI adoption acts as a powerful driver of efficiency . Increases in AI adoption result in significant productivity gains in both countries.
- In the US, these productivity gains are accompanied by significant labor market adjustment costs. AI adoption is associated with persistent negative effects on employment and total compensation growth in the short to medium term.
- In sharp contrast to the US, productivity gains from AI adoption are associated with persistent and significant increases in employment and wage growth over the short to medium horizon in Japan.

These differences may be explained by differences in labor market conditions associated with population aging, as well as differences in industrial structure. A key takeaway from these results is that the aggregate gains from AI primarily operate through efficiency improvements, while labor market adjustment is highly country-specific. Accordingly, the appropriate policy response to AI adoption may differ markedly across countries.

2 Literature Review

This paper relates to a growing literature studying the economic effects of artificial intelligence (AI), automation, and technological change on productivity, employment, and wages. While theoretical work has long emphasized the potential for technology to reshape labor demand, empirical evidence on AI in particular remains limited due to measurement challenges. Our study contributes to this literature by providing new evidence on the dynamic macroeconomic effects of AI exposure using industry-level job posting data.

2.1 Theoretical Perspectives on AI and Labor Markets

A foundational strand of the literature models technological change through a task-based framework, in which technology substitutes for labor in some tasks while complementing it in others. Seminal work by Autor et al. (2003) shows that computerization disproportionately affected routine tasks, leading to job polarization. Building on this framework, recent models explicitly incorporate AI as a general-purpose technology capable of automating cognitive and non-routine tasks.

Acemoglu and Restrepo (2018) develop a unified framework in which AI generates both displacement effects, as workers are replaced by machines, and reinstatement effects, as new tasks and occupations emerge. In this setting, productivity gains from AI do not necessarily translate into higher employment or wages in the short run. Extending this framework, Acemoglu and Restrepo (2020) demonstrate that automation can reduce labor demand at the local labor market level, even as aggregate productivity rises.

More recent theoretical work emphasizes that AI-driven productivity gains may materialize primarily through efficiency improvements rather than employment expansion. For example, Wang and Wong (2025) show that firms adopt AI to reduce marginal labor costs, which can depress employment and wage growth in the medium run despite higher output. These predictions align closely with the empirical patterns examined in this paper.

2.2 Empirical Evidence on Automation and Employment

Empirical studies on automation have largely focused on industrial robots due to the availability of systematic adoption data. Using variation in robot exposure across U.S. commuting zones, Acemoglu and Restrepo (2020) find that increased robot adoption leads to lower employment and wages, particularly for workers in routine occupations. Similar results are reported for European economies by Graetz and Michaels (2018), who find that robots increase productivity while having limited or negative effects on aggregate employment.

Compared to robots, AI is less tangible and more diffuse, making its economic effects harder to measure. Recent studies have begun to address this challenge using firm-level data and novel proxies. Babina et al. (2024) show that AI adoption is associated with firm growth and product innovation but also with workforce restructuring rather than broad-based job creation. At the task level, Brynjolfsson et al. (2025) document large productivity gains from generative AI in customer support tasks, with little immediate impact on employment.

Our paper complements this emerging empirical literature by focusing on industry-level dynamics and estimating the medium-run responses of employment, compensation, output, and productivity to AI exposure shocks.

2.3 Job Postings as a Measure of AI Adoption

A growing literature uses online job postings to measure labor demand, skill requirements, and technological change in real time. Job postings have been shown to closely track official vacancy measures and provide timely information on labor market conditions Barrero et al. (2021). Text-based analysis of postings has been used to study skill upgrading, task reallocation, and the diffusion of new technologies across firms and industries Deming and Kahn (2018).

In the context of AI, job postings are particularly informative because firms often signal AI adoption through explicit skill requirements, such as machine learning or natural language processing. Prior research demonstrates that changes in job posting content can reveal shifts in production technologies before they appear in traditional investment or productivity data Hershbein and Kahn (2018). By measuring AI exposure as the share of AI-related job postings at the industry level, our approach captures firms' revealed demand for AI-related capabilities and provides a timely proxy for AI deployment.

2.4 Macroeconomic Effects of Technological Change

This paper also contributes to the broader macroeconomic literature on technological change and productivity. Traditional growth accounting frameworks emphasize that productivity gains may not immediately translate into higher employment or wages due to adjustment costs and reallocation frictions. Recent empirical work highlights the importance of studying dynamic responses rather than static correlations.

To this end, we employ the local projection method developed by Jordà (2005), which allows for flexible estimation of impulse responses without imposing strong parametric assumptions. Our outcome variables are drawn from the quarterly industry-level dataset constructed by Hobijn et al. (2025), which provides internally consistent measures of labor input, output, and productivity. By combining these data with a novel measure of AI exposure, our analysis provides new evidence on how AI-driven productivity gains propagate through the macroeconomy over time.

Overall, the existing literature suggests that AI adoption can raise productivity while generating labor market adjustment costs. Our findings reinforce this view and extend it by quantifying the dynamic macroeconomic effects of AI exposure using novel data and a flexible empirical framework.

3 Data

This section describes the construction of our quarterly panel dataset for 13 U.S. industries over the period 2017Q2–2025Q2 and our measure of industry-level AI exposure. A key challenge is to quantify the extent of AI adoption across industries. We address this challenge by constructing a measure of industry-level AI adoption from establishment-level job postings. The measure captures the share of vacancies requiring AI-related skills and provides a timely proxy for firms' deployment of AI technologies. We discuss in detail how this measure is constructed.

3.1 Indeed Job Posting Data

Our primary data source is proprietary job posting information from Indeed, a leading online employment platform. The Indeed dataset aggregates job advertisements from multiple sources, including postings published directly on Indeed's platform and those collected from numerous external job boards, company career pages, and recruitment agencies. This comprehensive coverage captures the vast majority of online job listings in the United States. The data undergo a deduplication process to ensure that identical

positions advertised across multiple platforms are counted only once. It is important to note that online job postings do not correspond one-to-one with actual job openings, as some vacancies may not be advertised online, while a single posting may sometimes represent multiple positions. Nevertheless, Indeed Job Postings Index closely tracks the Job Openings and Labor Turnover Survey (JOLTS) published by the Bureau of Labor Statistics. This strong correspondence has been extensively documented by Indeed Hiring Lab in regular reports comparing the two series (Indeed Hiring Lab, 2025), and the Indeed data are publicly available through the Federal Reserve Bank of St. Louis (FRED) economic database. Prior research has validated the use of online job posting data for measuring labor market dynamics: Forsythe et al. (2020) demonstrate that Indeed postings track JOLTS closely during the COVID-19 period, while Adrjan et al. (2025) employ Indeed data across 20 OECD countries for labor market analysis.

Online job posting data have been widely used in recent labor economics research to study skill requirements (Deming and Kahn, 2018), technological change and labor demand (Hershbein and Kahn, 2018), AI adoption and employment (Acemoglu et al., 2022), and the labor market impacts of AI in developing economies (Copestake et al., 2024). Our approach to measuring AI adoption through keyword-based classification of job postings follows methodologies established in this literature.

Indeed employs advanced text analysis algorithms to extract structured information from raw job advertisements, including job titles, posting dates, geographic locations, and other key attributes. This automated extraction process enables systematic construction of industry-level and occupation-level measures of labor demand and technological adoption.

AI Posting Share. We measure industry-level AI adoption using job-posting data from Indeed. Each posting j is classified as AI-related ($d_j = 1$) if its title or description contains at least one term from a standardized keyword list covering machine learning, natural language processing, computer vision, and related technologies. Appendix A.1 describes the classification methodology in detail.

Let $\mathcal{J}_{it}$ denote the set of job postings in industry i during quarter t, and let $|\mathcal{J}_{it}|$ denote its cardinality. We define the industry-level AI posting share as:

$$s_{it} = \frac{\sum_{j \in \mathcal{J}_{it}} d_j}{|\mathcal{J}_{it}|} \times 100 \tag{1}$$

where $d_j \in \{0, 1\}$ is the AI classification indicator for posting j. The measure is expressed in percentage points so that estimated coefficients can be interpreted as the response to a one-percentage-point increase in AI adoption.

This measure captures the intensity of firms' revealed demand for AI-related skills at the industry level. Under the assumption that firms post AI-related vacancies when expanding AI-related production capacity, s_{it} provides a timely proxy for the pace of AI deployment. We discuss potential endogeneity concerns and our instrumental variable strategy in Section **??**.

This scaling expresses the adoption measure in percentage-point terms and ensures that estimated coefficients can be interpreted as the effect of a one-percentage-point increase in AI adoption. The analysis focuses exclusively on AI-related postings; references to robotics and other automation technologies are excluded from the baseline measure, though we explore broader definitions in robustness checks. Our measure captures revealed demand for AI-related skills. Under the assumption that firms would post vacancies when expanding AI-related production capacity, the share of AI-related postings provides a proxy for AI adoption intensity.

While online job postings provide a timely and high-frequency measure of labor demand, Several measurement considerations warrant discussion. First, job postings do not map one-to-one to realized hires or investment decisions. A single posting may represent multiple positions, and some vacancies may be filled internally without public advertisement. Second, online platforms may underrepresent certain industries or occupations that rely less on digital recruitment channels. If such underrepresentation is time-invariant within industries, our industry fixed effects absorb these differences. Third, our AI Share measure captures revealed demand for AI-related skills rather than direct expenditures on AI capital. To the extent that firms post AI-related vacancies when expanding AI-related production capacity or reorganizing tasks around AI technologies, the measure provides a proxy for AI adoption intensity. However, if AI-related hiring responds to anticipated productivity or demand shocks, the measure may contain endogenous components. We address this concern in Section 4.2 using an instrumental-variable strategy.

Total Index of Job Postings. To capture aggregate labor demand, we utilize Indeed's Total Index, which quantifies the seasonally adjusted volume of job advertisements relative to a pre-pandemic baseline. The index equals 100 on February 1, 2020, chosen as a reference point representing stable labor market conditions prior to the COVID-19 disruption. Construction of the Total Index involves aggregating de-duplicated job postings across all industries and occupations, followed by seasonal adjustment using standard X-13-ARIMA-SEATS procedures. This index provides our principal indicator of overall hiring activity and labor demand intensity.

The February 2020 baseline enables us to track labor market dynamics throughout the pandemic period and subsequent recovery phase. Index values exceeding 100 indicate

posting volumes above the pre-pandemic benchmark, while values below 100 reflect lower labor demand relative to that baseline.

Our sample spans 2017Q2 to 2025Q2 and covers the United States. Data are aggregated to the industry-quarter level to balance the need for sufficient within-cell observations with adequate cross-sectional and temporal variation for identifying dynamic effects.

3.2 Industry Classification

To align postings with labor market and productivity data, all job advertisements are mapped into 13 broad industry categories. This mapping is implemented using a standardized occupation-to-NAICS translation that links occupational information embedded in the postings to the corresponding industry classification. Appendix Appendix A.2 provides the complete occupation-to-industry mapping. Aggregation to these broader categories ensures consistency across data sources and produces a stable industry panel suitable for empirical analysis.

3.3 Outcome Variables

All outcome variables are drawn from the quarterly industry-level dataset developed by Hobijn et al. (2025). This dataset provides internally consistent quarterly measures of labor market and productivity outcomes for U.S. industries, constructed using harmonized data from the Bureau of Economic Analysis and the Bureau of Labor Statistics.

The outcomes used in the analysis include employment, total compensation, nominal output, real output, and labor productivity. Employment is measured as total hours worked at the industry level, while total compensation includes wages and supplemental labor income. Nominal output corresponds to industry-level value added measured at current prices, and real output is obtained by deflating nominal value added using industry-specific price indices. Labor productivity is defined as real output per hour worked. All outcome variables are seasonally adjusted. Growth rates are computed as log differences and expressed in annualized percentage terms. Using a single, unified data source for all outcomes ensures consistency across variables and avoids discrepancies arising from differences in measurement concepts or deflators across datasets.

The full dataset covers 2017–2025. However, substantial volatility and structural disruptions introduced by the COVID-19 pandemic generate pronounced deviations in both job-posting dynamics and outcome variable behavior.

3.4 Descriptive Statistics

Table 1 presents summary statistics for our sample. Panel A reports industry-level averages, ordered by AI adoption intensity. The Information sector exhibits the highest AI adoption (12.9%), followed by Utilities (5.2%) and Professional & Business Services (2.4%). In contrast, Retail Trade (0.2%) and Leisure & Hospitality (0.3%) show the lowest AI adoption. These patterns are consistent with the nature of AI technologies, which are most applicable to information-intensive and technical industries.

Panel B reports full sample statistics. The mean AI share across all industry-quarter observations is 2.3%, with substantial variation (standard deviation of 3.6%). Labor productivity growth averages 1.4% per quarter with considerable heterogeneity across industries and time periods.

Table 1: Summary Statistics

Panel A: Summary Statistics by Industry

Industry	N	AI Share (%)	Robot Share (%)	LP Growth (%)	Emp Growth (%)	Output Growth (%)
Information	33	12.87	3.19	6.58	0.44	7.44
Utilities	33	5.18	14.20	−0.02	0.86	1.22
Professional & Business	33	2.42	0.51	3.93	0.86	4.88
Financial Activities	33	2.29	0.41	2.09	0.89	2.79
Manufacturing	33	2.00	4.12	1.70	0.21	1.94
Wholesale Trade	33	1.64	0.41	−1.01	0.65	−0.21
Other Services	33	0.84	0.41	−1.08	0.22	−0.78
Mining & Logging	33	0.73	0.63	3.95	−0.48	4.27
Education & Health	33	0.58	0.95	1.20	1.86	3.06
Construction	33	0.49	1.35	−0.78	1.25	0.77
Transportation & Utilities	33	0.36	0.54	−1.73	3.31	1.52
Leisure & Hospitality	33	0.32	0.07	0.34	0.74	1.02
Retail Trade	33	0.17	0.05	3.33	−0.46	2.67

Panel B: Full Sample Statistics

Variable	N	Mean	SD	Min	Max
AI Share (%)	429	2.30	3.55	0.05	23.95
Labor Productivity Growth (%)	429	1.42	11.63	−51.42	53.87
Employment Growth (%)	429	0.80	13.69	−192.87	81.54
Total Compensation Growth (%)	429	4.55	14.94	−188.08	78.47
Real Output Growth (%)	429	2.35	19.26	−237.00	148.09

Notes: This table reports summary statistics for the quarterly industry-level panel from 2017Q2 to 2025Q2 (13 industries, 33 quarters). Panel A shows industry means ordered by AI Share. Panel B shows full sample statistics. AI Share denotes the percentage of job postings containing AI-related keywords. Growth rates are annualized.

Figure 1 displays the evolution of AI-related job postings over time. The aggregate AI share across all industries increased from approximately 1.5% in 2017 to over 4% by 2025. The growth accelerated notably after 2022, coinciding with the widespread adoption of generative AI technologies following the release of ChatGPT in late 2022.

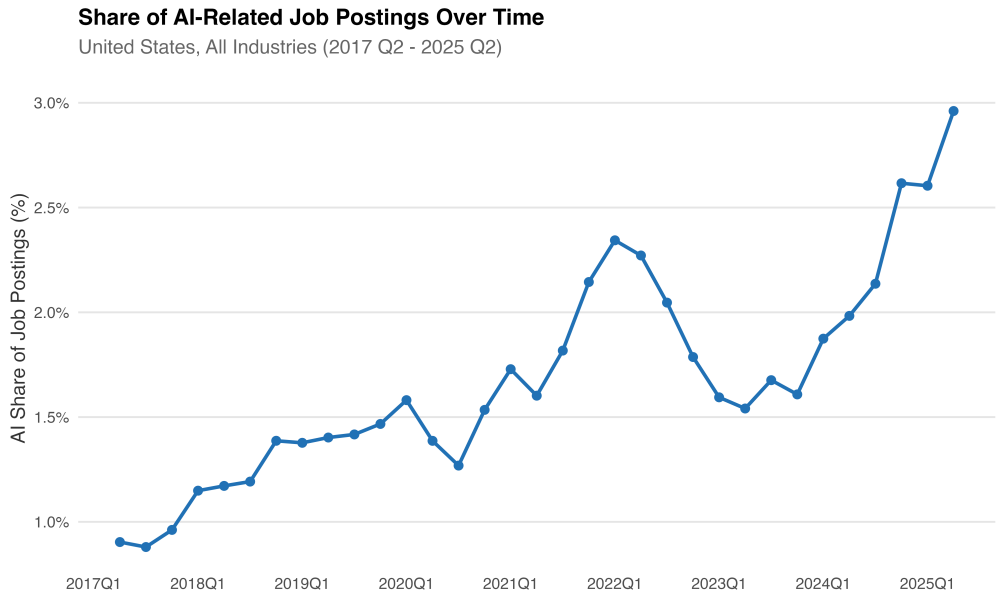


Figure 1: Share of AI-Related Job Postings Over Time

Notes: This figure shows the share of job postings containing AI-related keywords across all industries in the United States. The sample period is 2017Q2–2025Q2.

Figure 2 illustrates the substantial heterogeneity in AI adoption across industries. It displays time series of AI share for the six industries with highest average AI adoption. The Information sector stands out with AI share exceeding 20% by 2025, reflecting the technology-intensive nature of this industry, with its share more than doubling from the pre-pandemic period to the most recent years.

Figure 3 provides a cross-sectional comparison of AI adoption across all industries. The ranking of the AI adoption is based on 2023–2025 averages. The Information sector consistently leads, with its share more than doubling from the pre-pandemic period to the most recent years.

4 Empirical Methodology

The dynamic effects of industry-level AI adoption are estimated using the local projection method introduced by Jordà (2005). Local projections provide a flexible framework for estimating impulse responses in settings with heterogeneous dynamics and limited time-

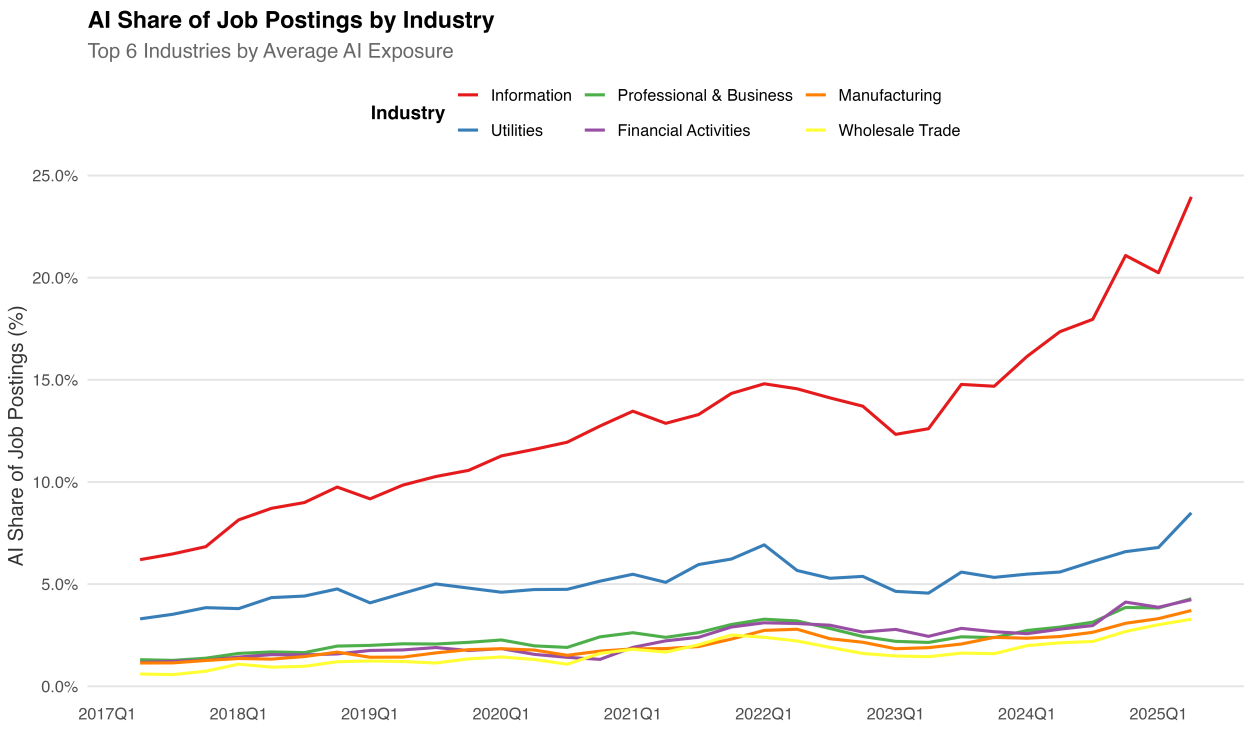


Figure 2: Top 6 Industries

Notes: This figure shows the share of job postings containing AI-related keywords by industry. It displays the top 6 industries by average AI adoption. The sample period is 2017Q2–2025Q2.

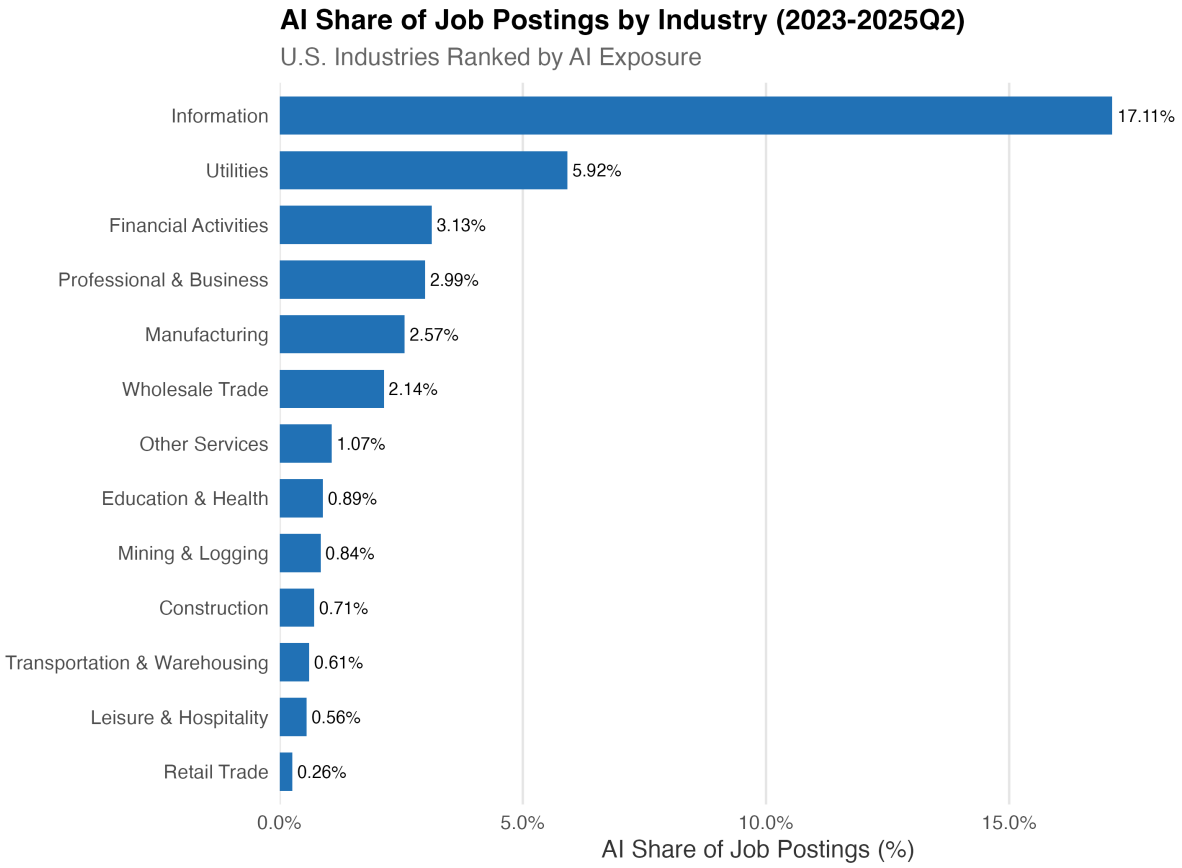


Figure 3: Industry Ranking (2023–2025)

Notes: This figure compares AI adoption across industries. It shows industry rankings based on 2023–2025 averages.

series length. Unlike vector autoregressions, local projections do not impose a specific parametric structure on the propagation of shocks and allow coefficients to vary freely across horizons, making this approach well suited for quarterly industry-level data with potential heterogeneity across sectors.

4.1 Baseline Specification

Our empirical strategy exploits within-industry variation in AI exposure over time. Intuitively, we compare industries that experience larger increases in AI-related hiring to those with smaller increases, while controlling for industry fixed effects and common macroeconomic shocks. Identification therefore comes from differential changes in AI exposure across industries within the same quarter.

To trace the dynamic effects of AI adoption, we employ the local projection (LP) method of Jordà (2005). This approach allows us to estimate impulse responses horizon by horizon without imposing strong parametric restrictions on the propagation of shocks. Specifically, we estimate the following for each horizon $h = 0, 1, \ldots, H$

$$y_{i,t+h} = \beta^h s_{i,t} + \phi_1^h y_{i,t-1} + \phi_2^h y_{i,t-2} + \theta_{i,A}^h D_{i,t+h}^A + \theta_{i,R}^h D_{i,t+h}^R + \alpha_i^h + \lambda_{t+h}^h + \varepsilon_{i,t+h}^h, \tag{2}$$

where i indexes industries and t indexes quarters. The dependent variable $y_{i,t}$ denotes the outcome of interest for industry i at time t, such as labor productivity growth. In the baseline specification, all dependent variables are defined as first differences of logarithms. The variable $s_{i,t}$ captures the AI adoption shock for industry i at time t, measured as the share of new AI-related job postings. To account for persistence in the outcome variable, we include two lags of the dependent variable, $y_{i,t-1}$ and $y_{i,t-2}$.

In addition, the COVID-19 pandemic generated extraordinary disruptions in both labor demand and output dynamics. To ensure that pandemic-related volatility does not confound our estimates, we include industry-specific dummy variables for (i) the acute contraction phase (2020Q2–2020Q3) and (ii) the recovery phase (2020Q4–2021Q4). These controls allow the pandemic shock to affect industries differentially, thereby isolating AI-specific variation from crisis-driven adjustments. The variables $D_{i,t+h}^A$ and $D_{i,t+h}^R$ are industry-specific dummy variables that take the value of one when $t + h$ falls within 2020Q2–2020Q3 (the acute phase) and 2020Q4–2021Q4 (the recovery phase), respectively. We also include industry fixed effects α_i^h and quarterly time fixed effects λ_t^h.

The horizon-specific coefficient $\beta^{(h)}$ represents the impulse response function (IRF) at horizon h, that is, the estimated effect of a one-percentage-point increase in AI adoption on the outcome variable h quarters after the shock. Identification relies on differential changes

in AI exposure across industries within the same quarter. Because time fixed effects absorb aggregate technological trends and macroeconomic shocks, the coefficient $\beta^{(h)}$ is identified from industries whose AI exposure increases more rapidly relative to others at a given point in time. A potential concern is that AI exposure may respond endogenously to anticipated productivity or demand shocks. For example, industries expecting higher future growth may expand AI-related hiring in advance. We address this concern using an instrumental-variable strategy described below. To characterize cumulative effects over time, we construct cumulative impulse response functions (CIRFs) as

$$\text{CIRF}(h) = \sum_{j=0}^{h} \beta^{(j)}, \qquad \text{se}(\text{CIRF}(h)) = \sqrt{\sum_{j=0}^{h} [\text{se}(\beta^{(j)})]^2} \tag{3}$$

where the standard error assumes independence of horizon-specific estimates. Confidence intervals are constructed using normal critical values at the 68% and 90% levels. Standard errors for the horizon-specific impulse responses are clustered at the industry level to allow for arbitrary serial correlation and heteroskedasticity within industries over time. Cumulative impulse response functions are constructed by summing the horizon-specific coefficients, and their standard errors are obtained by analytically aggregating the squared standard errors across horizons.

Our identification strategy relies on within-industry variation in AI adoption over time, while controlling for industry-specific time-invariant characteristics and common macroeconomic shocks through the inclusion of industry and quarterly time fixed effects. The quarterly time fixed effects absorb aggregate trends in AI adoption and macroeconomic conditions.

4.2 Instrumental Variable Approach

A potential concern is that AI adoption may be endogenous to expected future productivity or labor market conditions. Industries anticipating productivity gains may invest more heavily in AI, which would bias our estimates upward. Conversely, industries experiencing labor shortages may adopt AI as a substitute for scarce workers, potentially biasing estimates in the opposite direction.

To mitigate these concerns, as a robustness check, we complement the baseline OLS estimates with an instrumental variables (IV) specification designed to address potential endogeneity in AI adoption. While the baseline results rely on within-industry variation in AI exposure over time, AI adoption decisions may still respond to anticipated future productivity or labor market conditions.

Following Acemoglu and Restrepo (2020), we instrument the U.S. AI posting share with the corresponding measure constructed from EU and UK job postings on Indeed. For each horizon h, the first-stage equation is:

$$s_{i,t}^{US} = \gamma^{(h)} z_{i,t}^{EU} + \phi_{1,\text{FS}}^{(h)} y_{i,t-1} + \phi_{2,\text{FS}}^{(h)} y_{i,t-2} + \boldsymbol{\gamma}_{i,\text{FS}}^{(h)\prime} \mathbf{D}_{i,t+h} + \alpha_{i,\text{FS}}^{(h)} + \lambda_{t+h,\text{FS}}^{(h)} + \nu_{i,t}^{(h)} \tag{4}$$

where z_{it}^{EU} denotes the share of AI-related job postings in industry i in the EU and UK at time t, constructed analogously to equation (1) using Indeed posting data from nine European countries.[1] The first-stage equation includes the same controls, fixed effects, and industry-specific COVID interactions as the second-stage equation (2).

The identifying assumption is that AI adoption in the EU is correlated with U.S. AI adoption through common global technological developments, while being orthogonal to U.S.-specific industry demand shocks. Quarterly time fixed effects absorb common aggregate trends, so identification comes from differential changes in AI exposure across industries within the same quarter. We report horizon-specific first-stage F-statistics to assess instrument relevance.

The identifying assumption is that AI adoption in the EU is correlated with U.S. AI adoption through common global technological developments, but is orthogonal to U.S.-specific industry demand shocks. This assumption is plausible because AI technologies diffuse internationally along industry lines, while short-run demand conditions differ across regional labor markets.

The IV estimation is implemented within the same Aligned Backshift Local Projection (AB-LP) framework as the baseline specification. For each horizon h, the lagged U.S. AI adoption shock $s_{i,t-h}^{US}$ is instrumented using the corresponding lagged EU measure $z_{i,t-h}^{EU}$. The first-stage regression is estimated separately for each horizon and includes the same lagged outcome controls, fixed effects, and industry-specific COVID interactions as the second-stage equation. We report horizon-specific first-stage coefficients and F-statistics to assess instrument relevance.

5 Estimation Results

5.1 Baseline Estimation: Effect of AI adoption

This section presents the baseline dynamic effects of AI exposure estimated using local projections. We document three central findings. First, AI adoption generates sizable

[1] France, Germany, Spain, Ireland, United Kingdom, Italy, Netherlands, Sweden, and Belgium.

and persistent productivity gains. Second, in the United States, these productivity gains are accompanied by short- to medium-run declines in employment and wage growth. Third, the wage effects are highly heterogeneous across AI-related and non-AI-related jobs, consistent with task substitution.

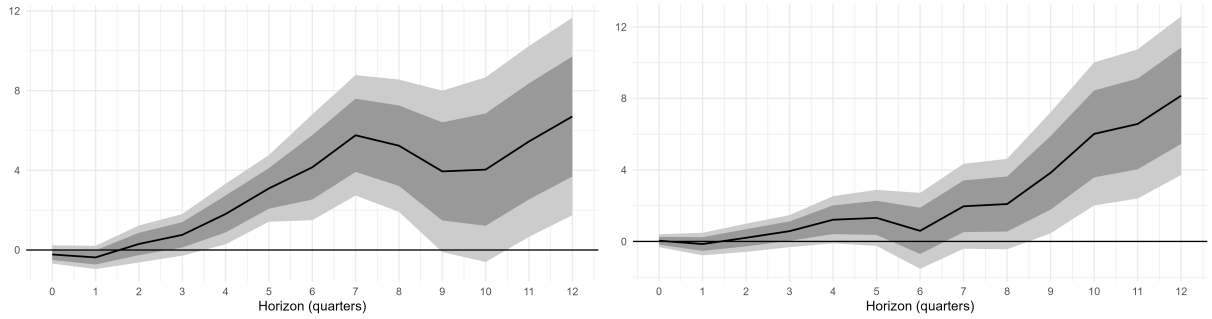
Figure 4 presents the cumulative impulse response functions to AI-related job postings. Panel (a) reports the response of labor productivity, showing a gradual increase beginning after five quarters and reaching a peak of approximately 6 percentage points after seven quarters. After eight quarters, the response remains significantly positive, albeit with some fluctuations. This pattern suggests that AI adoption contributes to higher labor productivity by improving business efficiency.

Panel (b) of Figure 4 displays the impulse response of real output. During the first eight quarters, real output increases slowly, although the response is not statistically significant. After nine quarters, however, real output rises significantly and continues to increase through the twelfth quarter.

Panel (c) illustrates the impulse response of employment growth to AI-related job postings. In the first five quarters, the response remains close to zero, indicating that AI adoption has little short-term effect on employment. After six quarters, employment growth begins to decline and becomes nearly statistically significant. Subsequently, after ten quarters, employment turns positive, although the response is not statistically significant.

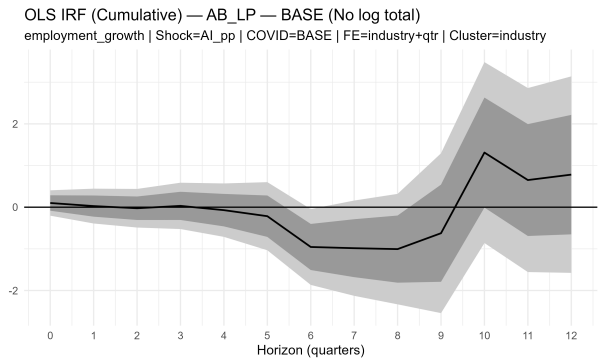
Finally, Panel (d) depicts the impulse response of total compensation to AI-related job postings. The response exhibits a pattern similar to that of employment growth. Specifically, total compensation declines significantly between the sixth and ninth quarters before turning positive after the tenth quarter. This pattern in total compensation can be attributed to changes in labor demand, as shown in Panel (c).

Taken together, these impulse responses suggest that during the first four quarters, AI adoption leads to a gradual increase in productivity through efficiency gains. This productivity improvement becomes more pronounced after five to six quarters and is associated with a temporary reduction in employment. As labor demand declines, employment falls, resulting in a temporary decrease in total compensation. Over time, rising productivity contributes to higher real output, which eventually stimulates employment growth and leads to an increase in total compensation.

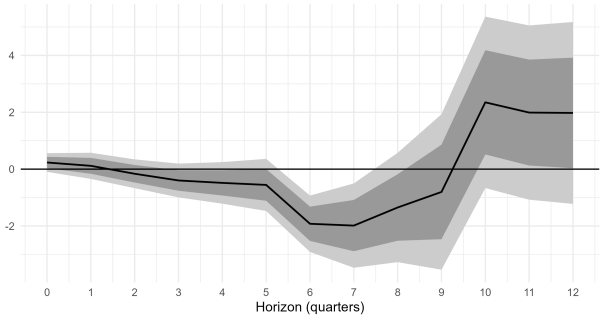


(a) Labor productivity

(b) Real output



(c) Employment growth



(d) Total compensation

Figure 4: IRFs to AI job postings

Notes: The figures report the estimated IRFs to changes in the share of new AI-related job postings, based on equation 2. The y-axis shows the percent change in the dependent variables in response to a one–percentage point increase in the share of AI-related job postings within total job postings in each sector.

5.2 Robustness Check: IV approach

In the previous analysis, we used new AI-related job postings as the explanatory variable of interest. However, one may argue that this variable is subject to an endogeneity problem. For instance, outcome variables such as productivity may be affected by confounding factors, including sectoral demand shocks. In such cases, the estimated impulse response functions may be inconsistent.

To address this endogeneity concern, we employ an instrumental variables approach following Acemoglu and Restrepo (2020). Specifically, we instrument AI-related job postings in the United States with AI-related job postings in the European Union. Because innovations in the AI industry diffuse rapidly across countries, AI-related job postings are likely to reflect common supply-side developments that are similar across the U.S. and the EU. In contrast, individual sectors in the EU are subject to different demand conditions. Therefore, using EU AI-related job postings as an instrument helps mitigate the influence of sector-specific demand shocks in the United States. We also note that aggregate demand conditions are controlled for by including time fixed effects, even in the baseline specification.[2]

The estimation results are presented in Figure 5. Panel (a) shows that labor productivity exhibits a pattern similar to that observed in Figure 4, although the confidence intervals are wider. The impulse response functions for real output, employment growth, and total compensation display similar patterns as well. Overall, these results provide evidence that the baseline findings are robust to concerns about endogeneity.

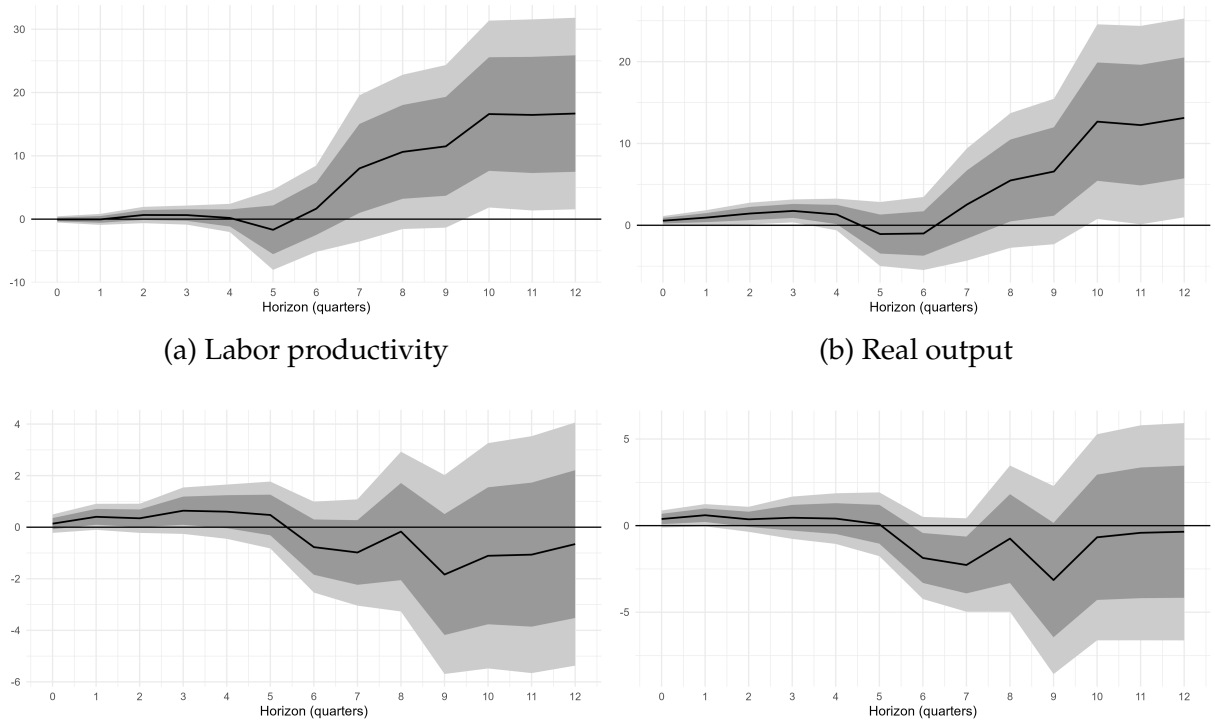
5.3 Wage for AI and non-AI Related Jobs

In the previous analysis, we examined the effects of AI-related job postings on the labor market. In this section, we exploit variation across job postings by distinguishing between AI-related and non-AI-related positions. To this end, we use job posting data from Indeed. Within each occupation, we compute the median wage separately for AI and non-AI job postings.

If our hypothesis—that AI adoption increases demand for AI-related jobs while reducing demand for non-AI jobs through efficiency gains within each sector—is correct, wages for AI-related jobs should not decline, whereas wages for non-AI-related jobs may decrease.

Figure 6 presents the cumulative impulse response functions of wages to AI adoption. Panel (a) shows that wages for non-AI-related jobs decline significantly after seven

[2]The first-stage F-statistic indicates that the estimation does not suffer from a weak-instrument problem.



(a) Labor productivity

(b) Real output

(c) Employment growth

(d) Total compensation

Figure 5: IRFs to AI job postings with IVs

Note:The figures report the estimated IRFs to the share of new AI-related job postings, based on equation 2 using AI-related job postings in EU as an instrumental variable. The y-axis shows the percent change in the dependent variables in response to a one–percentage point increase in the share of AI-related job postings within total job postings in each sector in the U.S.

quarters. In contrast, Panel (b) displays the impulse response of wages for AI-related jobs, showing that the response remains close to zero for up to twelve quarters. The divergence between the wage responses for AI-related and non-AI-related jobs suggests that AI adoption reduces labor demand for workers who are not directly engaged with AI technologies. The decline in labor demand for non–AI-related jobs at the sector level is consistent with the findings of Acemoglu et al. (2022), which are based on firm-level vacancy data from 2010–2018. While Acemoglu et al. (2022) conclude that the aggregate impact of AI is marginal at the aggregate level because its adoption was not widespread during that period, we find a significant impact on wages at the sector level. We extend their analysis by focusing on more recent periods, during which AI adoption has become more widespread.

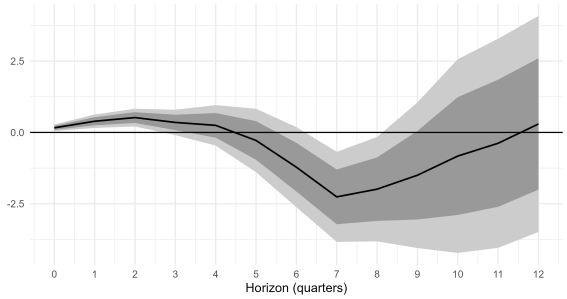
Finally, Panel (c) presents the impulse response of wages for all job postings, confirming that the impulse response of the median wage from Indeed is broadly comparable to that of total compensation shown in Panel (d) of Figure 4. In summary, the impulse response functions of the median wages for AI-related and non-AI-related jobs support the hypothesis that AI adoption reduces demand for non-AI workers by increasing efficiency.

5.4 Heterogeneous Impacts by Job Seniority

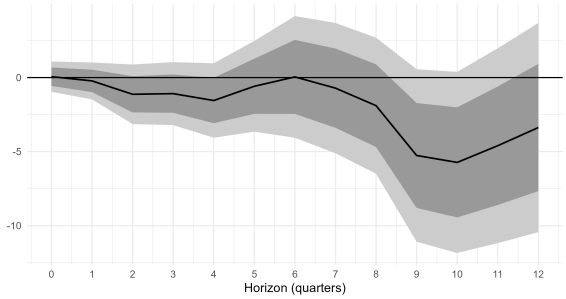
We find that AI adoption in the U.S. tends to increase labor productivity while reducing employment due to gains in efficiency. However, there is ongoing debate regarding which classes of workers are most affected by AI, even among those engaged in cognitive jobs. Some previous studies, such as Gambacorta et al. (2024), find that low-skilled workers benefit more from AI in computer coding. For policymakers, it is crucial to understand the heterogeneous effects on employment in order to design optimal policies.

To investigate the impact across job classes, we construct a measure of senior-level job postings based on information from job descriptions. Using this measure, we compute the senior-level job posting ratio by dividing the number of senior-level job postings by the total number of job postings. We use these two variables as dependent variables and estimate the LP model.

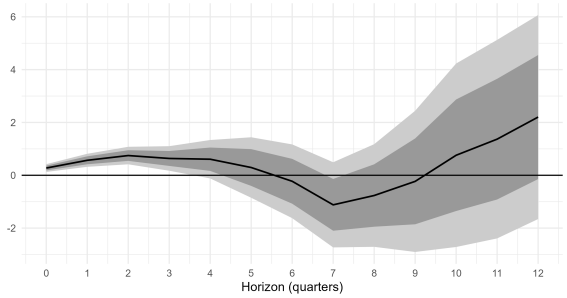
Panel (a) of Figure 7 reports the estimated IRFs for the share of senior-level job postings. The estimates indicate that AI adoption leads to an increase in the share of senior-level job postings, suggesting that labor demand shifts toward senior-level positions relative to junior positions following AI adoption. However, changes in the share do not necessarily reflect changes in levels, as the number of senior-level job postings may either increase or decrease.



(a) Median wage for non-AI job postings



(b) Median wage for AI job postings

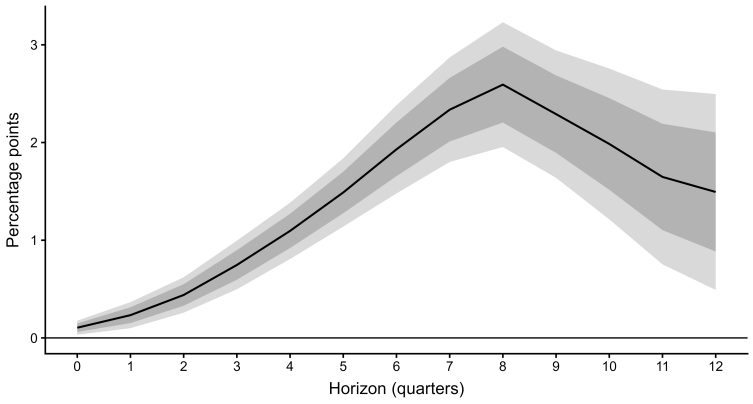


(c) Median wage for all job postings

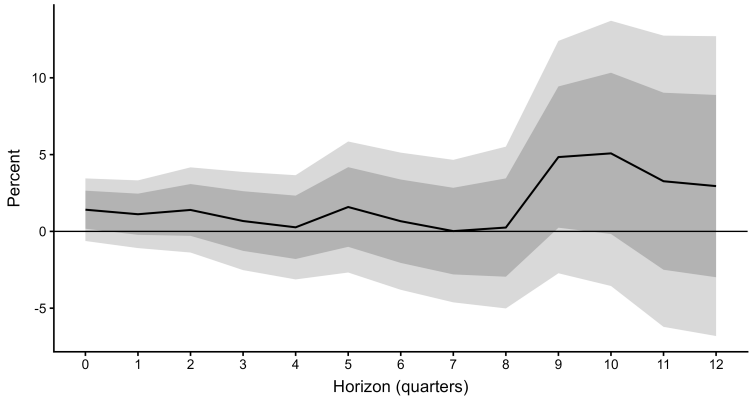
Figure 6: IRFs of wage for AI and non-AI jobs

Notes: The figures report the estimated IRFs to the share of new AI-related job postings, based on equation 2. The y-axis shows the percent change in the dependent variables in response to a one–percentage point increase in the share of AI-related job postings within total job postings in each sector.

To assess the impact on the level of senior-level labor demand, we re-estimate the LP model using the growth rate of senior-level job postings as the dependent variable. Panel (b) of Figure 7 shows that the growth rate of senior-level job postings remains statistically insignificant, or only marginally positive, following AI adoption. Taken together, these results imply that the increase in the share of senior-level job postings is primarily driven by a decline in non-senior-level job postings. In other words, AI adoption appears to exert a stronger substitution effect on jobs requiring less experience or specialized expertise.



(a) IRF of the share of senior-level job postings in the U.S.



(b) IRFs of growth rates of senior-level job postings in the U.S.

Figure 7: IRFs of senior-level job postings in the U.S.

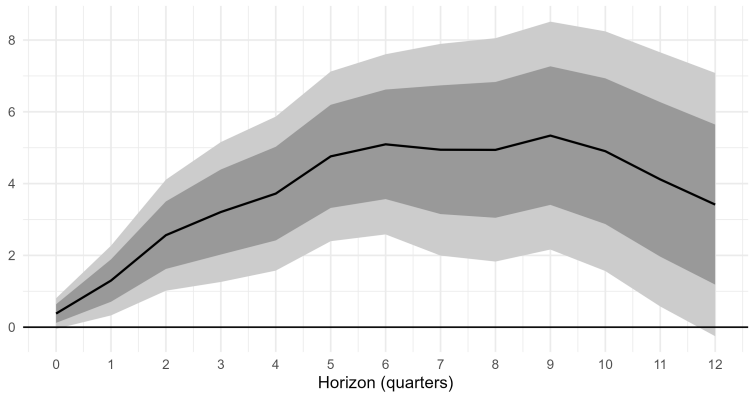
Notes: The figures report the estimated IRFs to the share of new AI-related job postings, based on equation 2. The y-axis shows the change in the dependent variables in response to a one–percentage point increase in the share of AI-related job postings within total job postings in each sector.

6 Effect of AI adoption in Japan

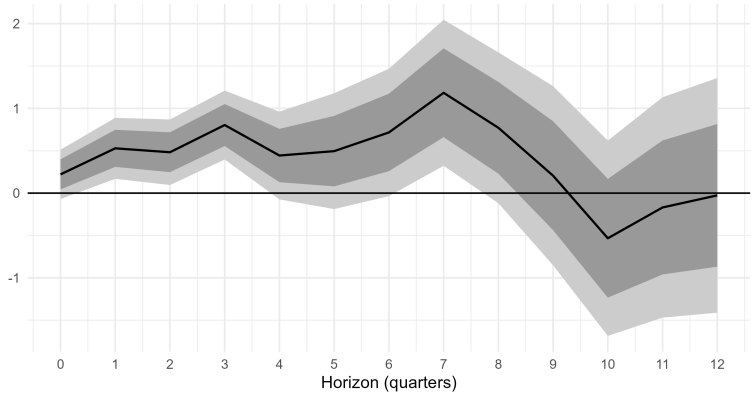
As discussed in the Introduction, existing literature provides ambiguous ex ante predictions regarding the sign of the impact of AI adoption on labor market outcomes such as wages. This ambiguity implies that the effects of AI adoption may vary depending on macroeconomic conditions, including the tightness of the labor market. For example, in countries where labor supply has been declining due to population aging, the available labor force may fall below its optimal level. In such settings, the introduction of AI can increase production efficiency and effectively augment labor input. As a result, AI adoption may expand economic activity, leading to higher labor demand and, on average, higher wages. In this section, by examining the impact of AI adoption in Japan, we demonstrate the possibility that AI adoption increases wages, with positive effects outweighing the substitution effect. (See Appendices B.2. to B.4. for details on the development of the AI job postings in Japan.) We note that these positive effects may arise through two main channels: (i) increased worker productivity due to complementary effects between AI and labor, and (ii) demand expansion through general equilibrium effects. Regarding outcome variables, we use labor productivity calculated by Japan Productivity Center.

Panel (a) of Figure 8 reports the IRFs of labor productivity, showing that a one–percentage point increase in the share of AI-related job postings in Japan leads to an increase of about 4 percentage points in labor productivity. Labor productivity rises gradually, reaches its peak after four to five quarters, and remains positive thereafter. The results for Japan are comparable to those for the United States, suggesting that productivity gains from AI adoption may be a common phenomenon across developed economies. Panel (b) of Figure 8 shows that the impact of AI adoption on employment in Japan is slightly positive. However, this positive effect materializes mainly within the first four quarters and dissipates within approximately three years.

To further investigate the impact of AI adoption on the labor market in Japan, we use wage data from official statistics and from Indeed for non–AI-related jobs as the dependent variable in the LP regression. Panel (a) of Figure 9 shows that wages slightly increase in the first few quarters and remain positive for three years, although the confidence interval widens as the horizon increases. This presents a substantial contrast with the results in the U.S., where employment slightly declines following the increase in AI adoption. Panel (b) of Figure 9 illustrates the estimated IRFs for wages in non–AI-related jobs, indicating that AI adoption significantly increases wages and that the effect is persistent. In addition, the magnitude is economically significant: a one-percentage-point increase in the AI job posting share leads to a five percent increase in wages for non-AI new job postings. In



(a) IRF of labor productivity in Japan



(b) IRF of employment to AI-related job postings in Japan

Figure 8: IRF to AI related job postings in Japan

Notes: The figures report the estimated IRFs to the share of new AI-related job postings, based on equation 2. The y-axis shows the change in the dependent variables in response to a one–percentage point increase in the share of AI-related job postings within total job postings in each sector.

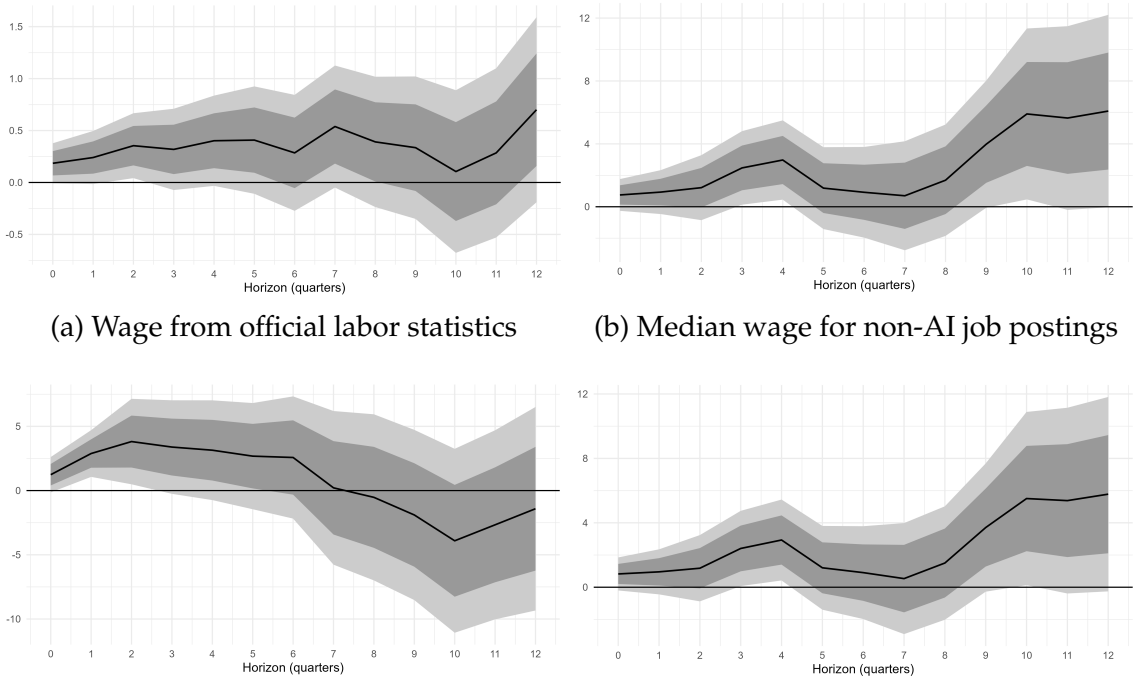
contrast, wages in AI-related jobs increase only slightly, and the effect becomes statistically insignificant at longer horizons. Panel (c) and (d) shows IRFs for wages of AI-related jobs and all jobs, indicating that an wage increases following AI adoption. These results stand in sharp contrast to those for the United States, where AI adoption does not significantly raise wages for non–AI-related jobs and instead is associated with a decline in those wages.

Finally, 10 shows that the share of senior-level job postings slightly increases in response to AI adoption in Japan. Panel (a) indicates a modest increase in the share of senior-level job postings, and Panel (b) shows a statistically significant increase in the number of senior-level job postings in Japan, albeit with a small impact of 0.3 percentage points at the peak. The impact is relatively small compared to that in the U.S.

Although a more detailed investigation of the underlying mechanisms is warranted, these findings suggest that, in Japan, the positive effects of AI—through higher worker productivity and overall economic expansion—outweigh the substitution effects associated with automation.

The stark differences between the results for Japan and the United States may be attributable to labor market conditions and differences in firms' adoption of information technology. For example, if advanced technologies have not yet been fully integrated into Japanese firms, AI adoption may act as a catalyst for broader digital transformation within organizations, generating substantial gains in production and productivity. In such a context, the positive effects of AI may dominate its negative impacts.

Another potential explanation is labor market tightness stemming from Japan's demographic trends, which may amplify the wage effects of productivity-enhancing technologies. The plausibility of these hypotheses can be examined using more granular firm- and worker-level data. We view this as an important topic for future research.



(a) Wage from official labor statistics

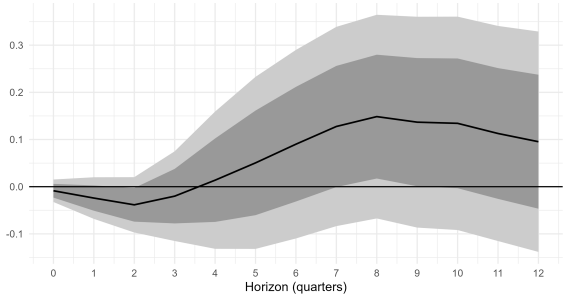
(b) Median wage for non-AI job postings

(c) Median wage for AI job postings

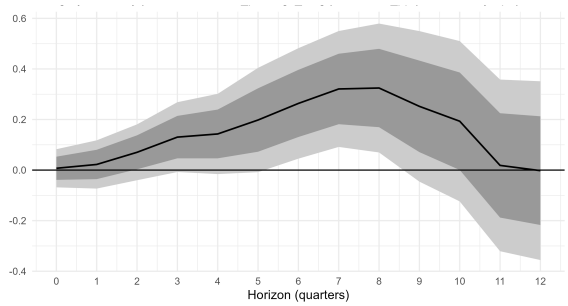
(d) Median wage for all job postings

Figure 9: IRFs of wage in Japan

Notes: The figures report the estimated IRFs to the share of new AI-related job postings, based on equation 2. The y-axis shows the percent change in the dependent variables in response to a one–percentage point increase in the share of AI-related job postings within total job postings in each sector in Japan.



(a) IRF of share of senior-level job postings in Japan



(b) IRF of the number of senior-level job postings in Japan

Figure 10: IRFs of senior-level job postings in Japan

Notes: The figures report the estimated IRFs to the share of new AI-related job postings, based on equation 2. The y-axis shows the change in the dependent variables in response to a one–percentage point increase in the share of AI-related job postings within total job postings in each sector.

7 Conclusion

This study provides novel empirical evidence on the dynamic macroeconomic consequences of AI adoption across U.S. and Japan industries. By utilizing a high-frequency dataset of AI-related job postings from Indeed as a proxy for adoption, we identify significant shifts in productivity and labor market dynamics.

Our findings indicate that AI adoption acts as a powerful driver of efficiency in both countries. However, more surprisingly, labor market responses differ sharply across the two countries. In the United States, these productivity gains are accompanied by significant labor market adjustment costs. AI adoption is associated with persistent negative effects on employment and total compensation growth in the short to medium term. This suggests that AI adoption reduces labor demand for workers whose tasks are not directly integrated with these new technologies.

On the other hand, in Japan, productivity gains from AI adoption are associated with persistent and significant increases in employment and wage growth over the short to medium horizon. This implies that the displacement effects predicted for workers whose tasks are directly integrated with AI do not materialize in this setting.

These results are robust to endogeneity concerns through an instrumental variables approach using EU job-posting data, confirming that the observed patterns are driven by supply-side technological advancements rather than sector-specific demand shocks. Ultimately, while AI holds the promise of significant aggregate productivity and output growth, its current trajectory poses challenges for labor demand and wage stability for a large segment of the workforce in the US. Policymakers and economists must therefore consider these distributional effects and adjustment costs when navigating the ongoing technological transition.

References

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

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D. and Restrepo, P. (2018). Artificial intelligence, automation, and work. *Journal of Economic Perspectives*, 32(2):197–236.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Adrjan, P., Ciminelli, G., Judes, A., Koelle, M., Schwellnus, C., and Sinclair, T. (2025).

Working from home after covid-19: What do job postings tell us? *Labour Economics*, 96:102751.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *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*. Forthcoming.

Barrero, J. M., Bloom, N., and Davis, S. J. (2021). Why working from home will stick. *Journal of Economic Perspectives*, 35(4):99–124.

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

Copestake, A., Marczinek, M., Pople, A., and Stapleton, K. (2024). Ai and services-led growth: Evidence from indian job adverts. *Working Paper*.

Deming, D. J. and Kahn, L. B. (2018). Skill requirements across firms and labor markets. *Journal of Labor Economics*, 36(S1):S337–S369.

Forsythe, E., Kahn, L. B., Lange, F., and Wiczer, D. (2020). Labor demand in the time of covid-19: Evidence from vacancy postings and ui claims. *Journal of Public Economics*, 189:104238.

Gambacorta, L., Qiu, H., Shan, S., and Rees, D. M. (2024). Generative ai and labour productivity: a field experiment on coding. *Bank for International Settlements WP*, 1208.

Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–768.

Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772.

Hobijn, B., Mestieri, M., Werquin, N., and Zhang, J. (2025). Quarterly industry-level labor productivity data for the us. *Economic Perspectives*, (1).

Indeed Hiring Lab (2025). Indeed's 2026 us jobs & hiring trends report. `hiringlab.org/2025/11/20/indeed-2026-us-jobs-hiring-trends-report/`.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Wang, C. and Wong, K. (2025). Artificial intelligence, labor demand, and productivity. *Journal of Economic Theory*. Forthcoming.

Appendix A Data Appendix

Appendix A.1 AI Keyword Classification

To identify AI-related job postings, we employ a keyword-based classification system developed by Indeed. Job advertisements are tagged as AI-related if their titles or descriptions contain terms associated with artificial intelligence technologies. The classification encompasses keywords organized into seven broad categories, following an approach similar to Babina et al. (2024).

Table 2 summarizes the keyword categories and provides representative examples. The categories cover core AI concepts as well as specific application domains. Machine Learning keywords include terms related to supervised and unsupervised learning algorithms, ensemble methods, and popular frameworks. Natural Language Processing (NLP) keywords capture text analysis, speech recognition, and language model technologies. Computer Vision keywords identify image and video processing applications. The remaining categories cover general AI terminology, neural network architectures, robotics applications, and autonomous systems.

The keyword list was developed through an iterative process combining expert review with data-driven validation. Initial keywords were selected based on established AI taxonomies and industry terminology. The list was subsequently refined by examining co-occurrence patterns with core AI terms (artificial intelligence, machine learning, deep learning, neural networks) in job posting text, following the methodology of Babina et al. (2024). Keywords with high co-occurrence rates were retained, while general technology terms with low AI specificity were excluded.

Job postings are classified as AI-related if they contain at least one keyword from the classification system. This binary classification is then aggregated to compute industry-level AI shares as described in Section 4.2.

Appendix A.2 Occupation-to-Industry Mapping

Job postings in the Indeed dataset are assigned to standardized occupation categories based on job titles and descriptions. To construct industry-level measures, we map these occupation categories to 13 broad industries following a classification consistent with Bureau of Labor Statistics industry definitions. Table 3 presents the complete mapping.

Table 2: AI Keyword Categories and Examples

Category	Representative Examples
Machine Learning	Deep Learning, Random Forest, XGBoost, Gradient Boosting, Support Vector Machines, K-Means Clustering, TensorFlow, PyTorch, Scikit-learn
Natural Language Processing	BERT, GPT, Transformer, Sentiment Analysis, Named Entity Recognition, Text Mining, Speech Recognition, Chatbot, Language Models
Artificial Intelligence (General)	Artificial Intelligence, AI, Machine Intelligence, Cognitive Computing, Expert Systems, AI Development, AI Strategy
Computer Vision	Image Recognition, Object Detection, Facial Recognition, Image Segmentation, OpenCV, YOLO, Convolutional Neural Networks
Neural Networks	Deep Neural Networks, Recurrent Neural Networks, LSTM, Autoencoders, Generative Adversarial Networks, Backpropagation
Robotics	Robotic Process Automation, Robot Operating System, Motion Planning, Cognitive Robotics
Autonomous Systems	Autonomous Vehicles, Self-Driving, ADAS, Sensor Fusion, Path Planning, LIDAR Processing

Notes: This table summarizes the AI keyword classification system used to identify AI-related job postings. Examples are illustrative and do not represent the complete keyword list.

Table 3: Occupation-to-Industry Mapping

Industry	Indeed Occupation Categories
Mining & Logging	Agriculture, Mining
Construction	Construction, Civil Engineering, Installation & Maintenance
Manufacturing	Chemical Engineering, Industrial Engineering, Production & Manufacturing
Utilities	Electrical Engineering, Mechanical Engineering, Utility
Wholesale Trade	Sales (wholesale)
Retail Trade	Retail
Transportation & Warehousing	Transportation, Driving, Warehouse, Aviation
Information	Media & Communications, IT Infrastructure, IT Systems & Solutions, Software Development, Creative
Financial Activities	Accounting, Banking & Finance, Insurance, Real Estate
Professional & Business Services	Administrative Assistance, Architecture, Consulting, Customer Service, Human Resources, Legal, Management, Marketing, Project Management, Security, Scientific Research, Data & Analytics
Education & Health Services	Nursing, Medical Information, Pharmacy, Dental, Physicians, Medical Technician, Therapy, Veterinary, Personal Care, Education, Childcare
Leisure & Hospitality	Food Service, Hospitality & Tourism, Leisure, Fashion, Sports, Arts & Entertainment
Other Services	Personal Services, Sanitation, Uncategorized

Notes: This table shows the mapping from Indeed occupation categories to 13 broad industries. The industry classification is consistent with BLS supersector definitions.

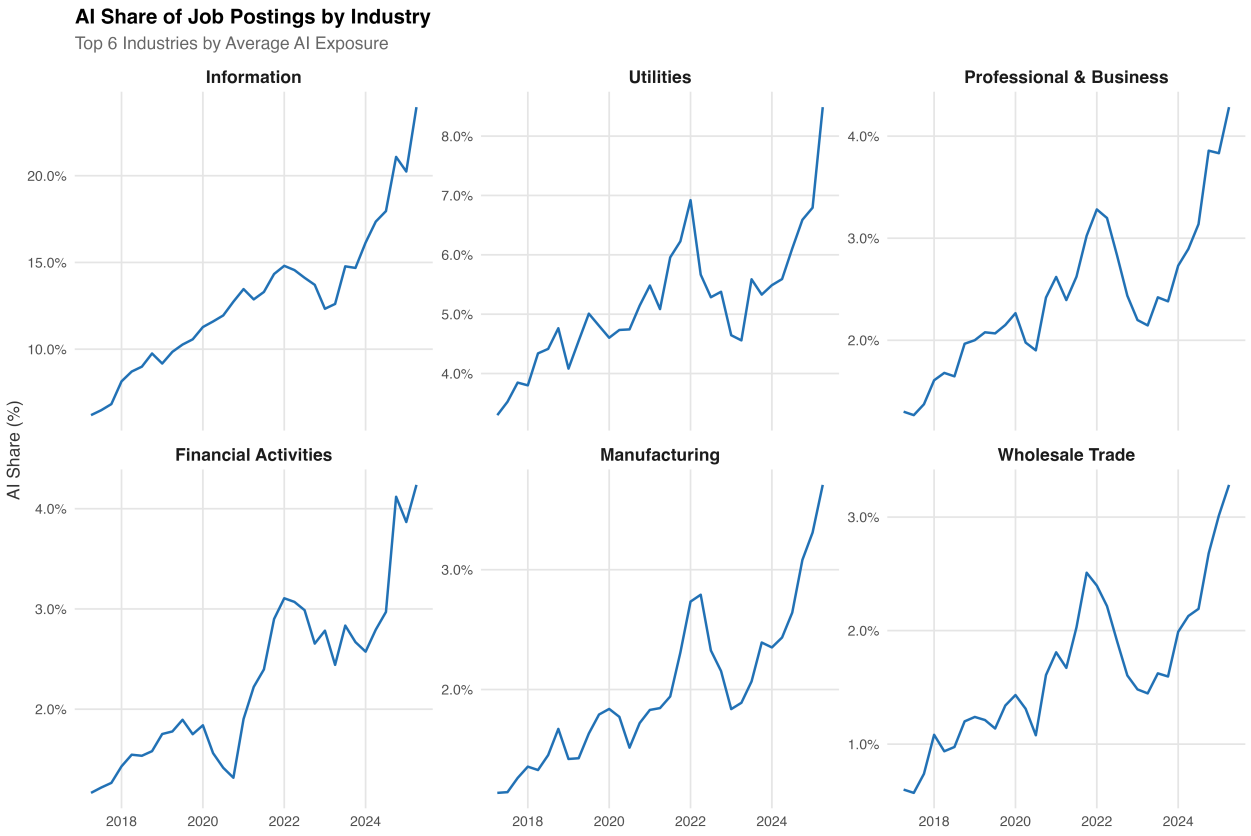


Figure 11: AI Share of Job Postings by Industry

Notes: This figure shows the share of job postings containing AI-related keywords by industry. It shows a faceted view with linear trend lines. The sample period is 2017Q2–2025Q2.

Appendix A.3 Industry-Level Trends

Appendix A.4 Occupation-Level Analysis

While the main analysis focuses on industry-level data at quarterly frequency, we also examine occupation-level patterns using monthly job posting data from Indeed. This higher-frequency data covers approximately 45 occupation categories from June 2017 to June 2025.

Figure 12 shows the monthly evolution of AI adoption. The monthly data confirms the acceleration in AI-related job postings following late 2022, with particularly sharp increases in early 2023 coinciding with the widespread adoption of generative AI technologies.

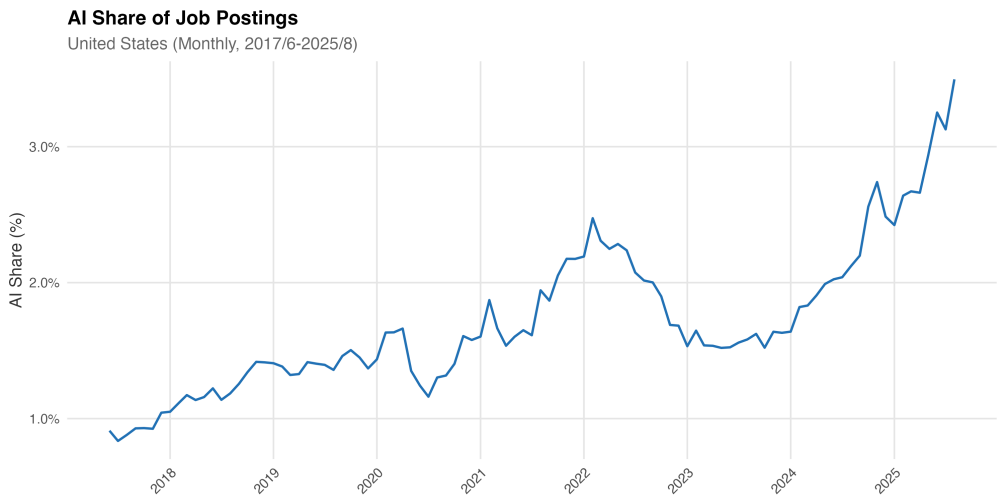


Figure 12: Monthly AI Share of Job Postings

Notes: This figure shows the monthly share of job postings containing AI-related keywords across all occupation categories in the United States. The sample period is June 2017 to June 2025.

Figure 13 displays AI adoption patterns across occupation categories. The figure shows the time series of AI share for the 10 occupation categories with highest average AI adoption. Software Development, Data & Analytics, and IT Systems & Solutions exhibit the highest AI shares, exceeding 25% by 2025. These patterns align with the technology-intensive nature of these occupations and are consistent with industry-level patterns showing high AI adoption in the Information sector.

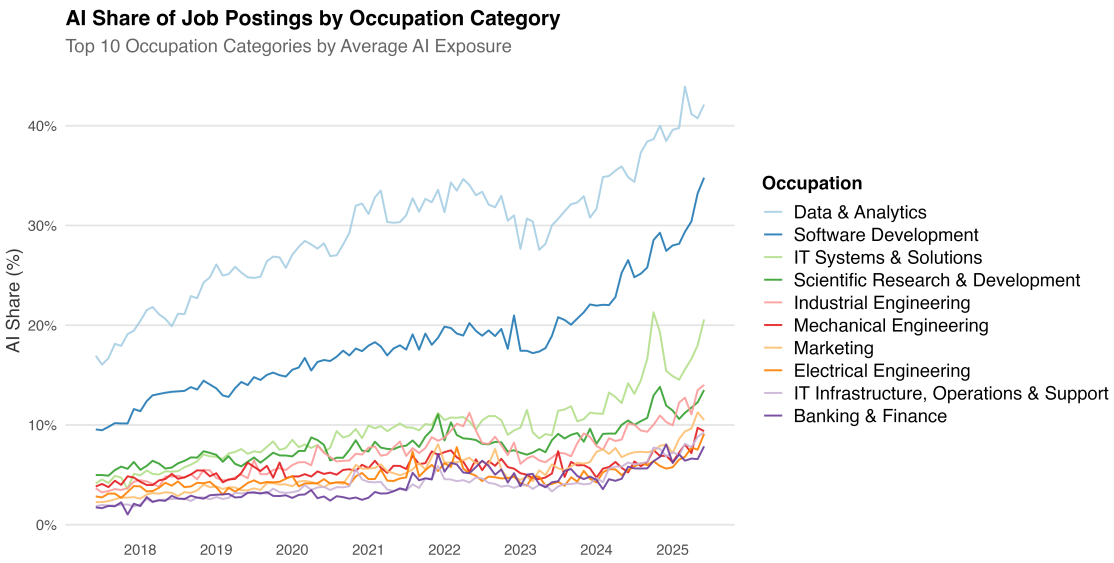


Figure 13: AI Share by Occupation Category (Top 10)

Notes: This figure shows the share of job postings containing AI-related keywords for the top 10 occupation categories. The sample period is June 2017 to June 2025 (monthly frequency).

Appendix B LLM-Based Industry Classification

Appendix B.1 Classification Methodology

Our baseline specification assigns each Indeed occupation category to a single NAICS industry classification using the deterministic mapping in Table 3. To assess sensitivity to this mapping, we construct an alternative probabilistic classification using a large language model (LLM) to assign job postings to industries based on their content.

The procedure consists of three steps. First, for each of the 61 Indeed occupation categories, we randomly sample 200 job descriptions from calendar year 2025, drawing from the same filtered universe of postings used in the baseline analysis. Second, for each sampled posting, we retrieve the job title and job description text from Indeed's job archive. Descriptions are truncated to 500 characters to retain sufficient information for classification while managing processing costs. Third, we submit batches of 20 job descriptions to the Anthropic API (Claude Sonnet 4.6) with the following prompt:

> *You are classifying job postings into NAICS industry classifications. For each job posting (identified by [ID: ...]), respond with ONLY a JSON array. Each element must be: {"id": "¡jobid¿", "industry": "¡industry_label¿"}.*
>
> *The industry must be exactly one of: "Construction", "Education & Health Services (EH)", "Financial Activities (FA)", "Information (IF)", "Leisure & Hospitality (LH)", "Manufacturing", "Mining & Logging (ML)", "Other Services (OS)", "Professional & Business Services*

(PB)", "Retail Trade (RT)", "Transportation & Warehousing (TW)", "Utilities", "Wholesale Trade (WT)".

Choose the SINGLE best-matching industry based on the job title and description. Respond with ONLY the JSON array, no other text.

Each job description is independently classified into one of 13 industries. Of the 12,200 sampled job descriptions, 89% were successfully classified. Unclassified postings are excluded when computing weights; the effective sample size ranges from 160 to 200 per occupation category.

For each occupation category j and industry s, the classification weight is computed as:

$$w_{js} = \frac{n_{js}}{\sum_{s'} n_{js'}} \tag{5}$$

where n_{js} is the number of job descriptions from occupation j classified into industry s. These weights represent the empirical probability that a posting in occupation j belongs to industry s. The alternative industry-level AI posting share is then constructed as:

$$\text{AI Share}_{st}^{LLM} = \frac{\sum_j w_{js} \cdot \text{AI Count}_{jt}}{\sum_j w_{js} \cdot \text{Total}_{jt}} \tag{6}$$

where the summation is over all 61 occupation categories. The classification weights are fixed across the sample period; only the underlying posting counts vary over time.

Appendix B.2 Classification Weights

Figure 14 displays the full weight matrix across 61 occupation categories and 13 industries. Medical and healthcare occupations (nursing, physicians, dental, therapy) are classified almost exclusively into Education & Health Services, with weights exceeding 95%. Similarly, driver and transportation occupations map predominantly to Transportation & Warehousing. These high-concentration categories closely match the baseline deterministic mapping.

The most informative cases are occupations where the LLM classification diverges from the one-to-one mapping. Software Development, assigned entirely to Information in the baseline, receives weights of 64% to Professional & Business Services, 14% to Information, and 10% to Manufacturing under the LLM classification. This reflects the fact that software engineers are employed across many industries, not exclusively in the technology sector. Similarly, Accounting splits between Financial Activities (39%) and Professional & Business Services (31%), and Pharmacy divides between Education & Health Services (62%)

and Retail Trade (36%), corresponding to hospital pharmacists versus retail pharmacies.

Of the 61 occupation categories, 12 have a top-1 weight exceeding 90%, 10 fall between 70–90%, and the remaining 39 exhibit meaningful multi-industry assignment.

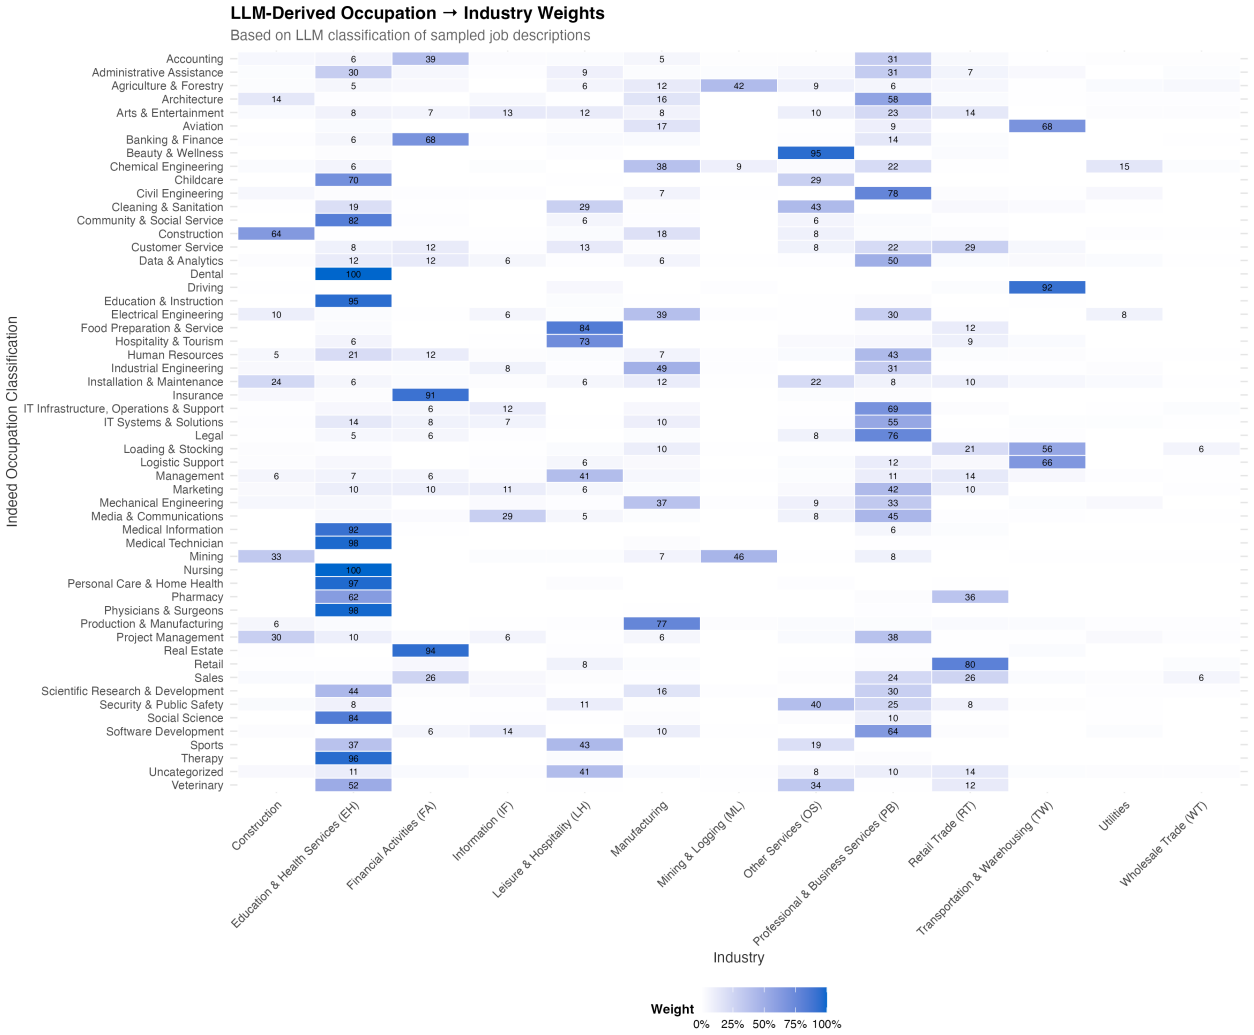


Figure 14: LLM-Derived Occupation-to-Industry Weights

Notes: Each cell shows the share of sampled job descriptions from a given occupation category classified into a given industry by the LLM. Cells with weights below 5% are left blank. Based on 200 randomly sampled job descriptions per occupation category, classified by Claude Sonnet 4.6 (Anthropic).

Appendix B.3 Comparison with Baseline Classification

Figure 15 compares the AI posting share under the two classifications. The top-left panel shows that the posting-weighted aggregate AI share is nearly identical, as the LLM weights

redistribute postings across industries without changing the total number of AI postings in the economy.

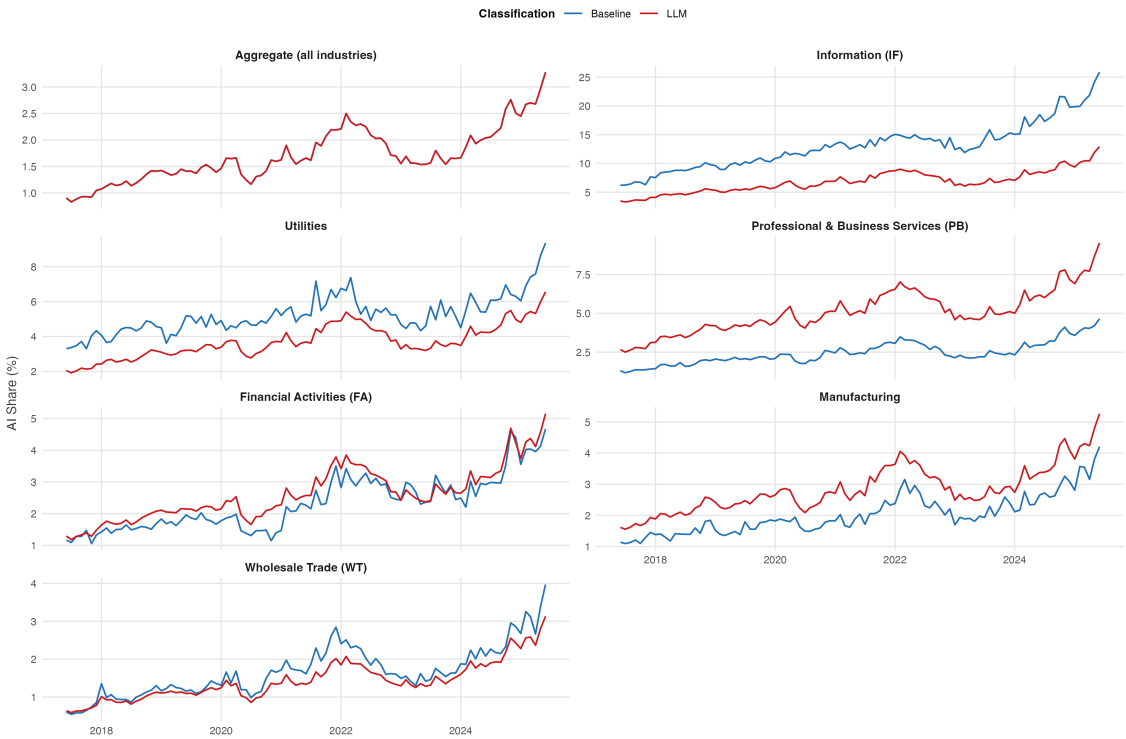


Figure 15: Aggregate AI Share and Top 6 Industry AI Share Over Time: Baseline vs. LLM Classification

Notes: This figure compares the AI posting share under the baseline (deterministic) and LLM (probabilistic) occupation-to-industry mappings. The top-left panel shows the posting-weighted aggregate across all industries. Remaining panels show the six industries with highest average AI adoption. Vertical axis scales differ across panels.

Figure 15 also compares industry-level AI share time series. The most notable difference is in Information, where the baseline AI share exceeds 20% by 2025 but falls to approximately 10% under the LLM classification. This reduction reflects the redistribution of technology-intensive occupations—software development, data analytics, and IT systems—to Professional & Business Services and Manufacturing. Conversely, Professional & Business Services exhibit higher AI shares under the LLM classification, as they absorb portions of these cross-cutting technology occupations.

Despite these differences in industry-level AI shares, the IV estimates using the EU instrument are qualitatively robust to the classification choice, as reported in Section **??**.

Appendix C Descriptive Statistics for Japan

This appendix presents descriptive evidence on the evolution of AI-related job postings in Japan using quarterly industry-level data from Indeed. The appendix documents summary statistics, aggregate trends, and industry-level heterogeneity in AI exposure. The purpose is to provide background information on the diffusion of AI across industries in Japan, without making causal claims.

Appendix C.1 Summary Statistics

Table 4 reports summary statistics for the main variables used in the Japan analysis. All variables are measured at quarterly frequency and aggregated at the industry level. The table reports means, standard deviations, and selected percentiles over the full sample period.

Table 4: Summary Statistics

Panel A: Summary Statistics by Industry

Industry	N	AI Share (%)	LP Growth (%)	Emp Growth (%)	Wage Growth (%)
Information & Communications	33	6.20	−0.29	4.55	1.24
Manufacturing	33	1.77	−0.68	−0.04	1.43
Professional & Technical	33	1.36	−0.48	2.31	1.61
Finance & Insurance	33	0.99	3.99	−0.48	2.01
Transport & Postal	33	0.63	0.09	0.32	1.11
Other Services	33	0.55	−0.04	2.08	1.41
Personal & Amusement	33	0.52	2.33	−0.14	1.94
Construction	33	0.48	−1.00	−0.02	2.01
Wholesale & Retail	33	0.47	−0.98	−0.24	1.73
Education	33	0.40	−3.65	1.57	0.82
Real Estate & Rental	33	0.37	−1.73	1.80	2.43
Health & Welfare	33	0.24	0.43	2.29	0.61
Accommodation & Food	33	0.13	0.48	1.28	1.37

Panel B: Full Sample Statistics

Variable	N	Mean	SD	Min	Max
AI Share (%)	429	1.09	1.67	0.00	10.35
Robot Share (%)	429	0.30	0.54	0.00	3.18
LP Growth (%)	429	−0.12	19.46	−169.69	109.78
Employment Growth (%)	429	1.17	10.68	−42.08	56.70
Wage Growth (%)	429	1.52	6.92	−24.41	24.28

Notes: Quarterly industry-level panel, 2017 Q2 to 2025 Q2 (13 industries, 33 quarters). AI Share (%) is defined as $100 \times$ (AI_not_robots_count + AI_and_robots_count)/total. Growth rates are annualized quarter-over-quarter changes: LP and Wage Growth are $400 \times \Delta \log(x)$; Employment Growth is 4× QoQ % change. All series are X-13 seasonally adjusted.

Appendix C.2 Aggregate Trends in AI-Related Job Postings in Japan

Figure 16 plots the quarterly share of AI-related job postings across all industries in Japan from 2017Q2 to 2025Q2. The AI share is defined as the fraction of job postings containing at least one AI-related keyword, as described in Appendix Appendix A.1.

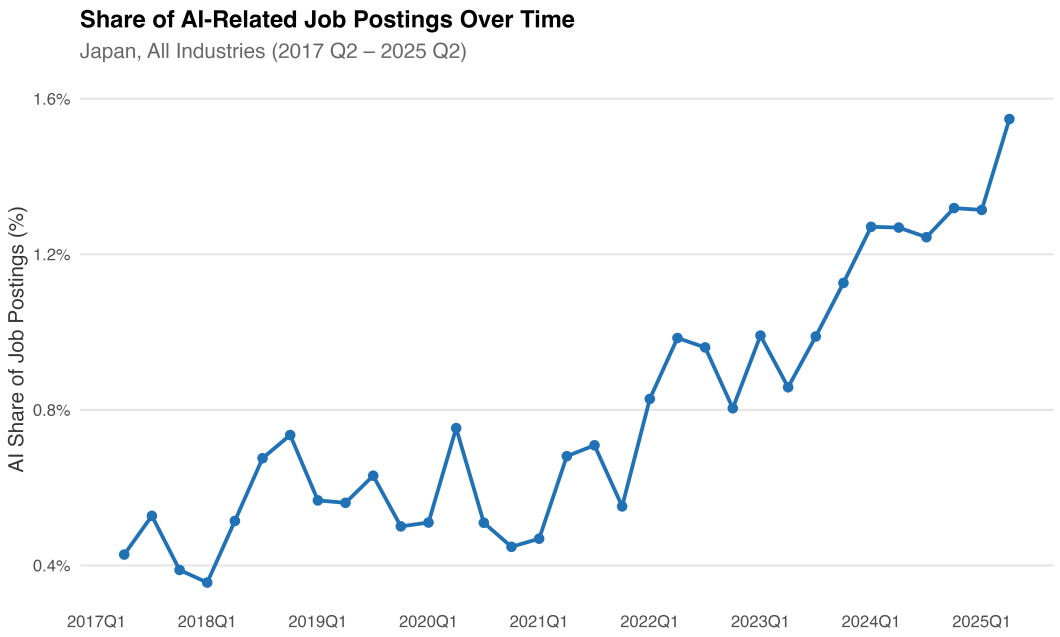


Figure 16: Aggregate Share of AI-Related Job Postings in Japan

Notes: This figure shows the quarterly share of job postings containing AI-related keywords across all industries in Japan. The sample period is 2017Q2–2025Q2.

The figure shows a gradual increase in AI-related job postings prior to 2020, followed by a temporary slowdown during the COVID-19 period. Beginning in late 2022, the AI share rises sharply, with a pronounced acceleration through 2024 and 2025.

Appendix C.3 Industry-Level Heterogeneity in Japan

Figure 17 presents AI adoption patterns for the six industries in Japan with the highest average AI exposure over the sample period.

Appendix C.4 Industry Trends and Long-Run Patterns in Japan

Figure 18 plots industry-specific AI shares together with fitted linear trends.

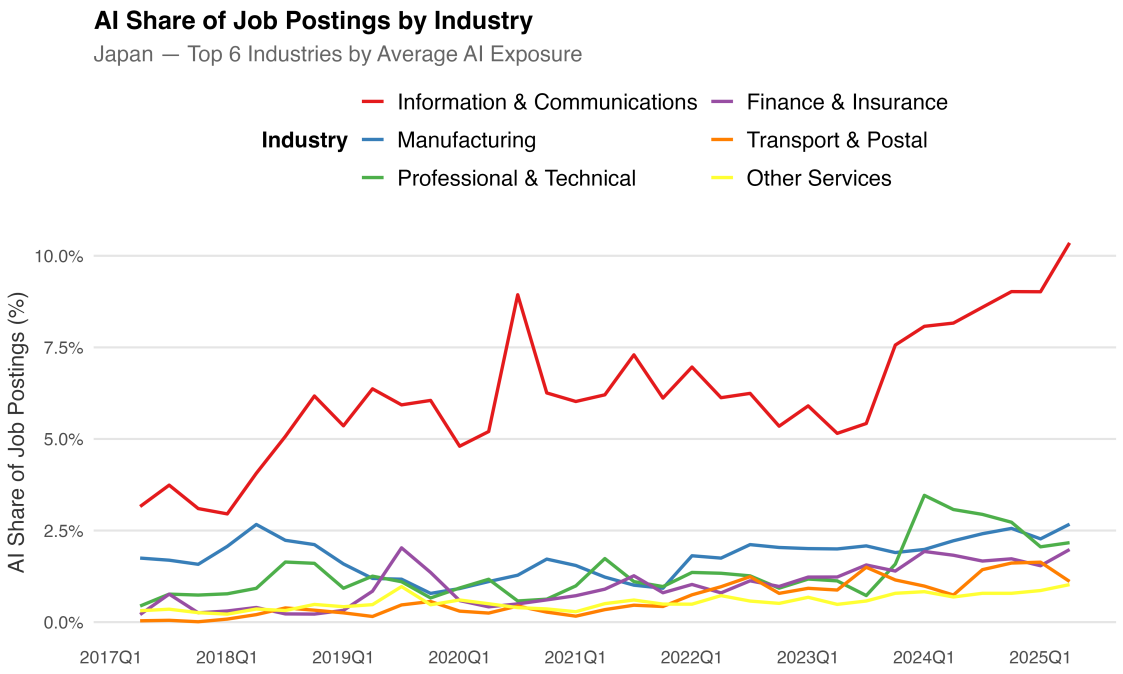


Figure 17: AI Share of Job Postings by Industry in Japan

Notes: This figure shows the quarterly AI share of job postings for the six Japanese industries with the highest average AI exposure.

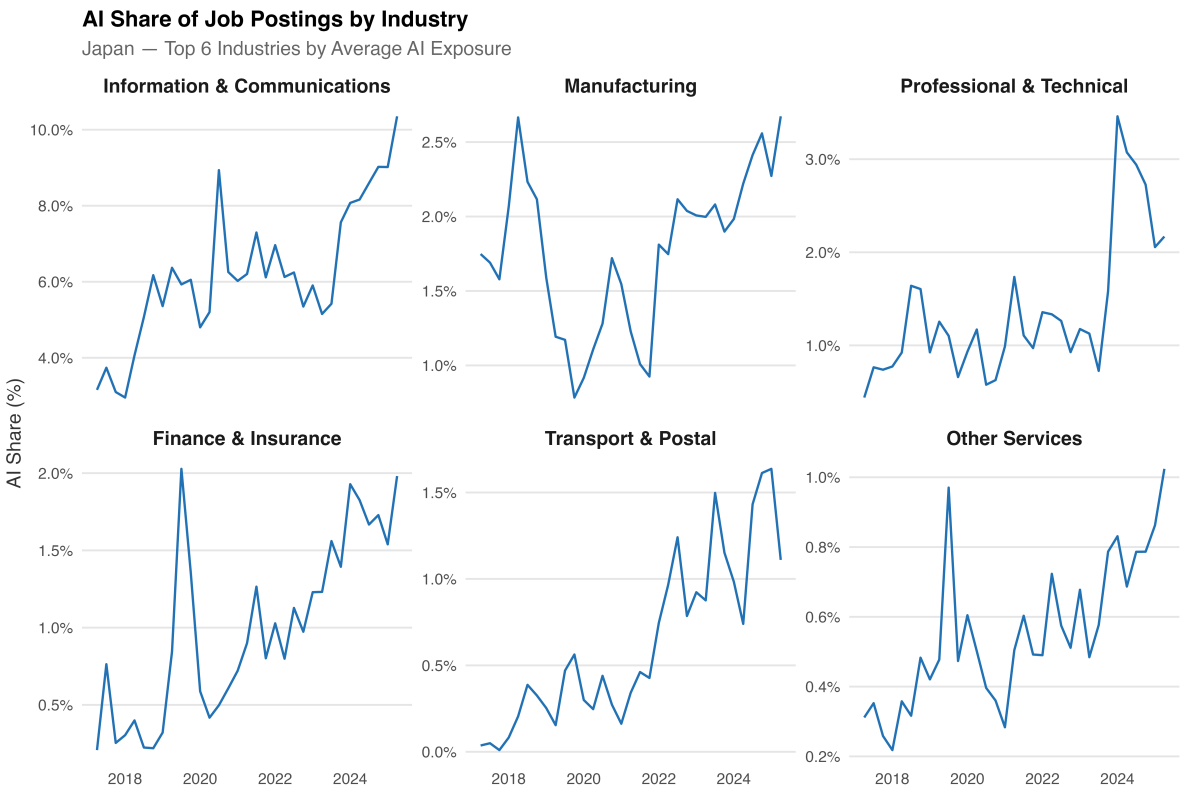


Figure 18: Industry-Level AI Trends in Japan

Notes: This figure displays quarterly AI shares by industry in Japan.