Discussion Paper Series

IZA DP No. 18740

June 2026

AI and Labor Market Outcomes: Evidence from China

Tony Fang
Memorial University of Newfoundland
and IZA@LISER

Carl Lin
Bucknell University

Qing Liu
Hefei University of Technology

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18740 June 2026

AI and Labor Market Outcomes: Evidence from China*

Abstract

As artificial intelligence (AI) spreads worldwide, its impact on labor markets is still unfolding and remains uncertain, yet potentially far-reaching in developing economies undergoing rapid structural transformation. This paper provides the first large-scale evidence from China, where AI investment and adoption have expanded rapidly, linking local AI labor demand to individual wage outcomes. We construct city–year measures of AI labor demand from 1.6 million online job postings between 2016 and 2024, capturing the intensity, breadth, and diversity of AI-related hiring, and merge them with nationally representative microdata from the China Family Panel Studies (2016–2022). Fixed-effects estimates show that local AI labor demand has positive impacts on individual wages: a one-unit increase in AI demand (1,000 postings, firms, or job titles) raises wages by about 0.2–0.3 percent. The benefits, however, are uneven. Women experience stronger gains—about 0.5–0.7 percent per unit increase—while men show no measurable effect. Wage effects are largest in Western provinces and in China's major AI-cluster cities—Jing–Jin–Ji, Yangtze River Delta, Pearl River Delta, and the Southwest–Central corridor—where complementary production factors and digital infrastructure are most developed. Occupational analyses further show that women's gains are concentrated in service-oriented, less skill-intensive jobs where AI complements interpersonal and coordination tasks rather than substituting them. Overall, AI diffusion generates meaningful but unequal labor market spillovers, with wage gains concentrated among women, dynamic regions, and human–AI complementary occupations, underscoring both the opportunities of technological transformation and the challenges of achieving inclusive growth.

JEL classification

I23, J24

Keywords

artificial intelligence (AI), labor market, wages, productivity, China

Corresponding author

Qing Liu

liuqingdm@hfut.edu.cn

* *Acknowledgements*: Fang acknowledges funding from the Social Sciences and Humanities Research Council (SSHRC) and the National Social Science Fund of China. Lin thanks the funding from the Bucknell Institute for Public Policy, the Center for Social Science Research, and the China Institute at Bucknell University. Liu thanks the funding provided by the National Natural Science Foundation of China (Grant No. 72373035).

Competing Interests: None of the authors has a conflict of interest to disclose.

1 Introduction

The rapid development and diffusion of artificial intelligence (AI) represent one of the most profound technological transformations of our time. AI is reshaping production processes, reorganizing tasks, and transforming labor markets. Policymakers, firms, and workers are grappling with a central question: who stands to gain—and who may lose—as AI adoption accelerates? This question is crucial for China, which combines the world's largest labor force with rapid AI investment and a labor market undergoing significant structural changes.

A central debate in the economics of AI concerns whether new technologies primarily substitute for workers—displacing labor and exerting downward pressure on wages—or instead complement human tasks, raising productivity and labor demand. Popular narratives often emphasize job loss and wage suppression, particularly for routine and middle-skill occupations. However, task-based models of technological change suggest that AI may also increase wages by creating new tasks, enhancing worker productivity, and raising demand for complementary skills. In this sense, technological change may generate both "winners" and "losers" in the labor market: some workers may experience displacement, while others retain their jobs and benefit from higher productivity and wages. As a result, the net effect of AI on wages is theoretically ambiguous and ultimately an empirical question. Conceptually, these opposing forces reflect substitution and complementarity channels whose relative strength may vary across workers, regions, and occupations.[1] In addition, demand-side adjustments may further shape wage outcomes, as productivity gains can expand product markets and increase labor demand depending on the elasticity of demand (Bessen, 2018).

[1] AI adoption can affect wages through at least two channels. Substitution effects may reduce demand for certain tasks, putting downward pressure on wages. Complementarity effects may raise productivity, expand output, and increase demand for labor in both AI-related and non-AI occupations, potentially benefiting some workers even as others face displacement. The relative strength of these channels is likely to vary across workers, regions, and occupations—an empirical pattern explored in this paper.

Beyond average wage effects, a key motivation of this paper is to examine whether the labor-market impacts of AI adoption differ systematically across workers, particularly by gender. Task-based theories of technological change emphasize that new technologies reshape the composition of tasks within jobs rather than eliminating occupations wholesale. Because men and women are unevenly distributed across occupations and task bundles—especially along dimensions such as interpersonal interaction, coordination, and routine processing—AI diffusion may generate heterogeneous wage responses even within the same local labor markets. In particular, AI technologies that automate routine components while complementing communication- and service-oriented tasks may raise the productivity and earnings of workers in roles where such tasks are prevalent. This perspective suggests that gender-differentiated wage effects are not an ex post anomaly but a plausible outcome of how AI interacts with existing patterns of occupational sorting and task specialization, which we are able to study using detailed measures of local AI labor demand constructed from online job postings.

This paper examines how local AI labor demand affects individual wage outcomes in China. Specifically, we ask whether cities undergoing stronger AI-related hiring also experience higher wage outcomes for workers, and whether these effects differ by gender, across regions, and between AI-cluster and non-cluster cities. We further explore the mechanisms underlying these patterns, focusing on occupational and skill complementarity as key channels through which AI adoption generates heterogeneous wage gains.

Our study makes four contributions. First, we provide the first large-scale evidence on the wage effects of AI adoption in China by constructing new city–year measures of AI labor demand from more than 1.6 million online job postings. Unlike conventional proxies based on patents or firm-level adoption, these data capture revealed labor demand for AI-related tasks and allow us to measure not only the intensity of AI hiring but also its breadth across firms and diversity across job titles. Second, leveraging this granular demand-side information, we move beyond average effects to document systematic heterogeneity in wage responses by gender, region, and urban technological ecosystems. Third, we offer a task-based interpretation of the observed gender differences, showing that women's stronger wage response arises primarily in occupations where AI complements interpersonal, coordination, and service-oriented tasks rather than substituting for them. Fourth, by linking detailed measures of local AI labor demand to nationally representative individual-level microdata, we present a unified empirical framework that traces how the diffusion of AI generates labor-market spillovers extending beyond narrowly defined AI workers.

To answer these questions, we merge a large dataset of AI-related job postings from major online recruitment platforms with individual-level data from the China Family Panel Studies (CFPS), covering the period from 2016 to 2022. We construct three complementary measures of local AI labor demand, including AI job postings, posting firms, and job titles, and then estimate fixed-effects models. A key challenge is potential endogeneity: cities that adopt AI more rapidly may differ systematically from others, or unobserved local factors may correlate with both AI adoption and wage outcomes. To mitigate these concerns, our specifications include an extensive set of controls: individual demographics, human capital, and employment characteristics; city and year fixed effects to absorb time-invariant local attributes and nationwide shocks; region–year interactions to capture broader macroeconomic trends; and province-specific time trends to account for gradual policy or structural changes.

Our findings show that stronger local AI labor demand is positively associated with individual wages. A one-unit increase in AI labor demand (equivalent to 1,000 postings, firms, or job titles) corresponds to statistically significant wage gains of approximately 0.2–0.3 percent. The benefits, however, are uneven. Women's wages increase by 0.5–3 percent per unit of AI labor demand, while men exhibit no significant effects. Regional disparities are also evident: the Western provinces experience the largest and most robust gains, effects in the East are modest, and impacts in the Central and Northeastern regions are statistically insignificant. The positive wage effects are concentrated in China's major AI-cluster cities—Jing–Jin–Ji (Beijing, Tianjin, and Hebei) region, Yangtze River Delta, Pearl River Delta, and the Southwest–Central corridor (Chengdu, Chongqing, and Wuhan)—where complementary production factors such as skilled labor, digital infrastructure, and innovative firms are well established. Outside these clusters, AI labor demand shows no statistically detectable effect on wages. Occupational evidence further suggests that the gender differences arise primarily from task complementarity—women benefit more in clerical, service, and production occupations where AI enhances, rather than replaces, human-centered work.

The remainder of the paper is organized as follows. Section 2 provides background on China's AI industry, outlining its policy environment, major investment trends, and spatial distribution of AI activity. Section 3 reviews the related literature. Section 4 presents data and summary statistics, detailing the construction of AI demand measures, the matched microdata, and key descriptive patterns. Section 5 describes the empirical strategy and reports the main results. Section 6 examines heterogeneous effects by gender, region, and AI-cluster cities. Section 7 investigates the mechanisms underlying the observed gender differences. Section 8 concludes.

2 China's AI Industry: Policy, Progress, and Regional Dynamics

China's AI industry has expanded rapidly over the past decade, driven by strong policy support, technological advancements, and increasing commercial applications. A clear strategic trajectory was established in 2017, when AI was included for the first time in the Government Work Report alongside the launch of the *New Generation Artificial Intelligence Development Plan* issued by the State Council of China. This blueprint set the goal of making China the world's leading AI power by 2030, with major milestones targeted for 2025. In 2024, the "AI+" initiative was elevated to a national priority at the Central Economic Work Conference, signaling the state's intent to integrate AI into productivity growth, industrial modernization, and the digital transformation of traditional sectors. These milestones underscore the central role of state policy in steering China's AI development.

Technological progress has reinforced these ambitions. Advances in data availability, algorithm design, and computing power have enabled machines to learn autonomously with minimal human programming. While early adoption focused on core algorithmic innovation, AI applications now span finance, healthcare, retail, and manufacturing, with growing penetration into frontier areas such as large language models, generative AI, autonomous driving, AI operations, and data product development. This diversification reflects both the maturation of the AI industry and its deepening integration into the broader economy. On the infrastructure side, China's national computing power reached an estimated 197 Exaflop[2] in 2023, with a target of 300 Exaflop by 2025. In early 2025, the government launched a 60 billion RMB *National AI Industry Investment Fund* to accelerate innovation and commercialization.

[2] An exaflop is a unit of computing performance that measures a supercomputer's ability to perform at least one quintillion (10^{18}) floating-point operations per second.

China's AI ecosystem benefits from a strong foundation of industrial capacity, technological expertise, and institutional support. The country has positioned itself as a global AI hub by leveraging the capabilities of major technology firms, including Alibaba, Tencent, Baidu, and Huawei, as well as vast datasets and a comprehensive industrial supply chain. State-owned enterprises have been directed to incorporate AI into the country's *15th Five-Year Plan* (2026–2030), while national and local governments have promoted investment through patient capital, startup incubation, and talent development programs.

Labor market evidence further highlights the growing premium for AI-related skills. AI positions consistently command salaries well above the national average. In 2023, the average monthly wage in the urban non-private sector reached 10,059 RMB and 5,700 RMB in the private sector, compared with 6,193 RMB and 3,813 RMB in 2017. Advertised AI positions remain far above these benchmarks, and the educational bar is high: roughly 80 percent of postings require at least a bachelor's degree, and an increasing share now demands a master's degree. These trends align with global evidence of substantial wage premiums for AI-skilled workers, underscoring the rising returns to advanced technical capabilities in China's labor market.

The geographic distribution of AI activity has also undergone significant evolution. Initially concentrated in Beijing, Shanghai, and Shenzhen, AI hiring has spread to a broader range of cities. Prominent clusters now include the Jing–Jin–Ji region (Beijing, Tianjin, and Hebei), Yangtze River Delta (Shanghai, Suzhou, Nanjing, and Hangzhou), Pearl River Delta (Shenzhen and Guangzhou), and the Southwest–Central corridor (Chengdu, Chongqing, and Wuhan). Nevertheless, a "diffusion gap" persists, leading firms and regions to be deeply engaged in AI adoption. At the same time, many small and medium-sized enterprises (SMEs) remain constrained by skill shortages and uneven infrastructure readiness.

Overall, China's AI industry has entered a new phase characterized by rapid expansion, technological diversification, and broader regional diffusion, supported by robust state policy and industrial integration. Persistent challenges remain, including talent shortages, uneven adoption across firms, and the tension between regulation and innovation. Yet, the combination of strong policy momentum, expanding industrial capacity, and rising labor market premiums for AI skills underscores the transformative potential of AI for China's economy. These developments provide the empirical foundation for the following analysis, which examines how local AI labor demand affects individual labor market outcomes.

3 Literature Review

The economic implications of AI have become a significant focus in recent research, spanning theory, measurement, and empirical evidence. While early debates emphasize automation and job displacement, newer studies—drawing on firm microdata, patent text, and online job postings—paint a more nuanced picture in which AI both substitutes and complements human labor. This literature collectively shows that the effects of AI depend crucially on how technologies interact with the task structure of work, the organization of production, and institutional contexts, as well as on demand-side adjustments that determine whether productivity gains translate into expanded labor demand (Bessen, 2018). This growing body of work is synthesized in recent surveys such as Lu and Zhou (2021), which emphasize the diverse channels through which AI affects productivity, labor demand, and inequality. Our study, which links local AI labor demand to individual wage outcomes in China, builds on this body of work while addressing a critical empirical gap.

The theoretical foundations of the AI–labor nexus rest on task-based models of technological change. In Acemoglu and Restrepo's (2018, 2020, 2025) framework, automation operates at the level of tasks rather than jobs, producing a displacement effect as machines substitute for human labor and a reinstatement effect as new, complementary human tasks emerge. The balance between these effects determines whether technological change raises or reduces labor demand and inequality. Recent evidence from China further highlights the heterogeneous effects of automation in developing-country contexts. For example, Zhang and Feng (2026) show that automation exposure reduces employment and wages overall—particularly for male and low-skilled workers—while simultaneously increasing employment opportunities for women, underscoring the complex distributional consequences of technological change.

Acemoglu (2025) estimates that current AI adoption patterns could increase U.S. total factor productivity by only 0.5–0.9 percent over a decade, suggesting that macroeconomic gains depend less on automation and more on creating complementary human tasks. Similarly, Agrawal, Gans, and Goldfarb (2019) conceptualize AI as a "prediction technology" that reduces the cost of inference and thereby transforms decision-making, organizational structures, and skill requirements. Brynjolfsson, Rock, and Syverson (2023) frame AI as a general-purpose technology whose productivity potential materializes only when accompanied by investments in data, human capital, and organizational adaptation. Together, these frameworks help explain why the microeconomic effects of AI adoption may become visible before its aggregate productivity impacts are apparent.

Building on these conceptual foundations, a growing body of empirical research seeks to measure the impact of AI technologies on occupations and tasks. Webb (2020) pioneered a text-based method that quantifies the overlap between the language of patent documents and the task descriptions in O*NET to measure occupational exposure. Unlike industrial robots or traditional software—which primarily affected low- and middle-skill, routine work—AI is found to target high-skill, cognitive, and analytical tasks such as diagnosis, prediction, and optimization. Based on historical substitution patterns, Webb (2020) projects that AI could modestly reduce overall 90–10 wage inequality while slightly increasing inequality at the top (99–90).

Subsequent studies have refined these exposure metrics. Sytsma and Sousa (2024) apply large-language-model techniques to classify over 18,000 tasks and find that by 2019, about 15 percent of U.S. workers were in highly AI-exposed occupations, concentrated in management, healthcare, and information services. Septiandria, Constantinides, and Quercia (2024) extend this approach by applying deep-learning semantic matching between more than 24,000 AI patents and 17,000 occupational tasks, demonstrating that AI has begun to impact complex, non-routine domains once considered automation-resistant. Complementary evidence from the Pew Research Center (2023) indicates that nearly one in five U.S. workers is in a highly AI-exposed occupation—twice the share in low-exposure jobs—with women slightly more exposed because of their overrepresentation in education, healthcare, and administrative fields. These studies show that AI differs fundamentally from earlier waves of automation by penetrating professional and analytical occupations rather than routine manual work.

At the firm level, analyses highlight heterogeneous productivity and employment effects of AI adoption. Using U.S. patent and business microdata, Alderucci et al. (2020) find that firms developing AI-related inventions experience roughly 25 percent faster employment growth and 40 percent faster revenue growth, accompanied by higher within-firm wage dispersion. Damioli, Van Roy, and Vertesy (2021) report similar productivity gains among European small and medium-sized enterprises, albeit with limited wage pass-through. Using newly linked U.S.

Census data, Dinlersoz et al. (2024) demonstrate that AI adoption increases output and labor productivity, particularly in information and professional services, though with limited evidence of widespread job displacement. More recent worker-level evidence points in a similar direction. Brynjolfsson et al. (2025), studying the rollout of a generative AI assistant in customer support, find that AI assistance increases worker productivity by about 15 percent on average, with especially large gains for less experienced and lower-skilled workers, highlighting the heterogeneous effects of AI within occupations.

Evidence from developing economies shows a similar pattern of duality. Copestake et al. (2024), using millions of online job postings from India, document a rapid increase in AI-skill demand since 2015 but find that adoption is accompanied by reductions in non-AI hiring and a widening wage polarization. At the aggregate level, Acemoglu, Dinlersoz, and Seamans (2024) use the U.S. Census Business Trends and Outlook Survey and show that between September 2023 and February 2024, the share of firms using AI increased from 3.7 to 5.4 percent, with larger and younger firms leading adoption. Most of these firms reported retraining employees and reorganizing workflows rather than cutting jobs, suggesting that the early phase of AI diffusion is characterized more by organizational adaptation than direct labor substitution.

Beyond firm productivity, researchers have increasingly examined how AI affects inequality and labor shares. Josifidis and Supic (2024) argue that AI may deepen the productivity-pay gap by strengthening capital's bargaining power relative to labor and expanding managerial control, thereby reviving classical concerns about the distribution of technological gains. Autor et al. (2024) emphasize that AI risks polarizing employment by eroding middle-skill jobs that combine routine and interpersonal tasks, while simultaneously creating new opportunities at both the high and low ends of the skill distribution. Comparative evidence from Webb (2020) and the Pew

Research Center (2023) further shows that the direction and magnitude of AI's labor-market effects vary with occupational structure, union coverage, and the strength of social institutions.

In developing contexts, the balance between complement and substitution effects depends on these same structural and institutional factors. Economies with large labor surpluses and uneven digital capacity—such as China—may experience more substantial displacement effects in some regions. At the same time, high-skill urban centers could benefit from complementarity and new task creation. Such heterogeneity underscores the importance of studying AI's impact at a localized level.

Taken together, three broad conclusions emerge from the existing literature. First, AI represents a qualitatively distinct wave of technological change centered on cognitive, analytical, and predictive tasks rather than physical automation. Second, its labor-market effects are uneven, boosting productivity and firm growth in AI-intensive sectors while widening within-firm and regional disparities. Third, aggregate productivity gains remain modest due to the slow diffusion of innovations and the high costs associated with building complementary capabilities. Despite the rapid expansion of this literature, however, most evidence still comes from advanced economies, leaving open questions about how AI interacts with labor markets undergoing structural transformation.

China's experience provides a particularly important yet underexplored case. Despite the country's position as a global leader in AI investment and home to the world's largest labor force, empirical evidence on how the diffusion of AI affects its labor markets remains limited.[3] Most existing research focuses on macroeconomic or industrial dimensions—such as patenting

[3] A few sectoral studies (Li and Li 2023; Sun et al. 2024) examine firm productivity and automation effects, yet they rarely connect AI adoption to wages or employment structures.

activity, R&D intensity, and innovation clustering in cities like Beijing, Shenzhen, and Hangzhou—while offering limited micro-level evidence on worker outcomes. Our study helps bridge this gap by linking city-level measures of AI labor demand, constructed from more than 1.6 million online job postings between 2016 and 2024, with individual-level wage data from the CFPS. In doing so, our paper provides new empirical evidence on how the diffusion of AI affects wage income across gender, region, and sector within China's rapidly evolving labor market.

4 Data and Summary Statistics

Online job postings provide a direct window into firms' revealed labor demand and the task content of jobs as new technologies diffuse. Unlike patent-based or firm-level adoption measures, job postings capture how AI adoption translates into concrete hiring needs across occupations, firms, and locations, making them particularly well-suited for studying task-level spillovers and heterogeneous wage responses in local labor markets.

4.1 Data Sources

We combine two complementary datasets to examine the causal relationship between AI labor demand and wage outcomes in China: a large-scale city-level dataset of AI-related job postings and individual-level data from the CFPS. Linking these datasets at the city–year level provides a novel opportunity to study how local AI demand is associated with individual labor market outcomes.

The first dataset consists of AI-related job postings collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS 直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58 同城), and Ganji.com (赶

集网), among others.[4] We identify AI-related jobs using 73 keywords across core AI technologies, natural language and speech, computer vision, big data/computing, and applications and smart living services (see Appendix Table A1).[5] After removing duplicates and incomplete entries, the final dataset comprises approximately 1.6 million unique job postings spanning 2016 to 2024.

To reduce potential duplication arising from reposted vacancies over time or across platforms, we implement a cleaning procedure before aggregation. Postings with identical combinations of firm name, job title, location, and posting period are treated as duplicates and removed from the dataset. In addition, the analysis does not rely solely on the raw number of postings. We construct three complementary measures of local AI labor demand—the number of AI job postings, the number of distinct firms posting AI jobs, and the number of distinct AI job titles—which capture the intensity, breadth, and diversity of AI-related hiring. Because repeated postings from a single employer would increase the postings measure but not the number of firms or distinct job titles, the consistency of results across these three measures helps mitigate

[4] Online recruitment platforms are widely used for formal-sector hiring in urban China but may not capture all recruitment activity, particularly among smaller informal firms or employers relying on offline hiring channels. However, evidence from our dataset suggests that online postings span a broad range of industries and firm sizes. For example, in a representative subsample of approximately 700,000 postings from the Zhaopin platform in 2021, production and manufacturing occupations account for about 13.8 percent of listings—exceeding the share of IT/Internet-related jobs (10.5 percent)—while sales, customer service, and service-sector positions together account for nearly 40 percent. In addition, nearly 70 percent of postings originate from small and medium-sized enterprises. Because our full dataset also includes postings from 58.com and Ganji.com—platforms widely used for blue-collar and local service employment—the combined sample likely provides even broader coverage. To the extent that some AI adoption occurs outside online recruitment channels, our measure may understate the true extent of AI diffusion, implying that the estimated coefficients should be interpreted as conservative estimates of its broader labor-market effects.

[5] While keyword-based classification provides a practical approach for identifying AI-related hiring in large-scale job-posting data, some classification error is possible. For example, broad technological terms such as “big data” may occasionally capture positions only indirectly related to AI, while firms may adopt AI technologies without explicitly labeling positions as AI-related. In addition, online job postings may not fully capture hiring activity by smaller, less formal firms or those relying on offline recruitment channels. To the extent that such misclassification or incomplete coverage occurs, the resulting measurement error is likely to be approximately classical and would tend to attenuate the estimated coefficients toward zero. The estimated effects should therefore be interpreted as conservative, and potentially lower-bound, estimates of the relationship between local AI labor demand and wage outcomes, particularly in regions where online recruitment is less prevalent—and may also contribute to weaker estimated effects in non-cluster cities where AI adoption is less fully captured by online postings.

concerns that the estimates primarily reflect reposting behavior rather than genuine expansion of AI-related labor demand.

These postings provide detailed information on firm name, job title, location (city and district), compensation (minimum and maximum monthly salary), education and experience requirements, number of openings, job type, and posting/closing dates. From these variables, we construct city–year measures of AI activity, including the total number of AI job postings, the number of distinct firms posting AI jobs, and the number of distinct AI job titles. These indicators capture both the scale and diversity of local AI labor demand and serve as the key explanatory variables in our analysis.

The second dataset is the CFPS, a nationally representative biennial longitudinal survey conducted by the Institute of Social Science Survey at Peking University. The baseline wave in 2010 covered 14,960 households and 42,590 individuals across 25 provinces, municipalities, and autonomous regions, representing about 95 percent of China's population (Xie and Hu, 2014). Follow-up surveys have been conducted every two years, collecting rich information on household structure, education, employment, income, health, and migration. Notably, the CFPS began distinguishing between formal and informal employment in 2014.

For this study, we focus on the adult sample in 2016, 2018, 2020, and 2022 waves to align with the job postings data. Key outcome variables include annual after-tax wage income (expressed in constant 2022 RMB), employment status, years of schooling, hours worked, hukou status, marital status, health status, and internet use. Respondents are geocoded to their city or county of residence, which allows us to match them with city–year measures of AI job postings.

We then merge the two datasets to link CFPS individuals with local AI activity at the city–year level. This matched dataset enables an analysis of the relationship between AI diffusion and individual wage outcomes. We restrict the sample to wage earners aged 18–60 who are employed in non-agricultural occupations. Observations with non-positive annual after-tax wages or non-positive monthly wages are excluded, and weekly working hours are trimmed at the 1st and 99th percentiles to address implausible reports. The final sample includes 6,905 working-age individual observations across four survey waves (2016–2022).

Because the analysis focuses on wage outcomes, the sample is restricted to non-agricultural wage earners with positive reported wages. Consequently, the estimates capture intensive-margin wage responses among employed workers, rather than employment, labor force participation, or entry and exit decisions. Self-employed and informal workers are therefore excluded from the wage analysis because comparable wage measures are not consistently available for these groups in the CFPS data. Moreover, because female labor force participation may be selective, the gender-differentiated results should be interpreted as applying to employed women in wage-paying jobs, rather than to all women in the labor force.

4.2 AI Job Postings and Firms

To provide background on the evolution of AI-related hiring in China, Figures 1 and 2 plot the total number of postings and the number of distinct firms posting AI jobs between 2016 and 2024. Together, these figures capture both the intensity and breadth of AI labor demand.

[Figure 1 about here]

[Figure 2 about here]

Two patterns stand out. First, postings and firms expanded rapidly in the late 2010s. Total postings rose sharply after 2016, while the number of firms engaged in AI hiring increased from fewer than 30,000 in 2016 to more than 70,000 by 2018. This parallel growth suggests that AI adoption was not confined to a small group of leading technology companies but diffused broadly across employers in sectors such as manufacturing, finance, and services.

Second, the early 2020s reveal considerable volatility. Both the number of postings and the number of posting firms fell sharply around 2020, coinciding with the COVID-19 pandemic and broader macroeconomic uncertainty. Although activity partially rebounded in subsequent years, neither measure returned to its 2018–2019 peak by 2024. These fluctuations underscore the cyclical sensitivity of AI labor demand: while the long-term trajectory indicates continued diffusion, short-term shocks can disrupt adoption and hiring.

Several factors may explain the decline in AI-related postings observed in the early 2020s. First, it partly reflects the broader contraction in hiring activity following the COVID-19 pandemic. Recruitment-platform reports indicate that overall hiring demand in China fell sharply during the early stage of the pandemic and recovered only gradually thereafter. For example, a report by Zhaopin (2022), one of China's largest online recruitment platforms, shows that overall recruitment volume declined by roughly 20 percent year-on-year in early 2020, and that by early 2021 total postings on the platform remained about 8 percent below the level observed during the same period in 2019. In addition to these cyclical effects, the pattern may also reflect changes in how firms adopt AI technologies. Recent policy and industry discussions increasingly emphasize the integration of AI into existing sectors ("AI+"), in which AI capabilities are embedded within workflows in manufacturing, services, healthcare, and education rather than implemented through the creation of new AI-specific job titles. As AI diffusion moves toward broader application, firms may rely more on task reorganization and skill upgrading within existing occupations rather than continued expansion of specialized AI job postings. Consequently, measures based solely on AI-specific job postings or titles may understate the extent of AI diffusion during later stages of adoption.[6]

Overall, Figure 1 and Figure 2 show that China's AI labor market experienced a phase of rapid expansion followed by a period of retrenchment. The scale of postings and the breadth of firms both signal strong underlying demand for AI skills, yet momentum has been uneven. These dynamics underscore the importance of analyzing not only whether AI-related hiring occurs, but how its intensity varies across time and place. In the following empirical analysis, we exploit these city–year postings and firm activity variations to investigate how local AI demand translates into wage outcomes for individuals.

4.3 Geographic Clusters of AI Job Postings

Figure 3 highlights the spatial concentration of AI-related hiring by showing the leading AI clusters across Chinese cities. The figure shows that AI demand is highly uneven, with a handful of metropolitan areas accounting for a disproportionate share of postings.

[Figure 3 about here]

Unsurprisingly, Beijing, Shanghai, and Shenzhen dominate as the primary centers of AI hiring, reflecting their roles as hubs for technology firms, research institutions, and venture capital. Other large coastal cities such as Guangzhou, Hangzhou, and Nanjing are also prominently featured, indicating that AI diffusion has been strongest in regions with advanced digital infrastructure and well-established innovation ecosystems.

At the same time, AI activity is not confined to the top tier. For example, cities such as Chengdu, Xi'an, and Wuhan also appear among the leading clusters, underscoring the role of

[6] Recruitment-platform reports indicate that overall hiring demand in China declined sharply during the early stage of the COVID-19 pandemic (Zhaopin, 2022), At the same time, recent labor-market analyses suggest that AI capabilities are increasingly embedded within a wide range of occupations rather than appearing only in specialized AI job titles (Microsoft and LinkedIn, 2024).

inland centers in China's AI development strategy. These cities often benefit from strong university systems and targeted government support, even if their scale of activity remains smaller than that of coastal leaders.

Taken together, Figure 3 underscores the geographic concentration of AI labor demand in China. While postings are present nationwide, activity is heavily clustered in a limited number of urban centers, primarily along the coast. This pattern highlights the importance of considering regional heterogeneity in the following empirical analysis: the labor market impacts of AI adoption are likely to differ sharply between leading clusters and other cities.

4.4 Summary Statistics Analysis

Table 1 presents descriptive statistics for the main variables used in the analysis, reported for the full sample and separately by gender. The sample consists of 6,905 working-age individual observations drawn from the 2016, 2018, 2020, and 2022 waves of the CFPS, matched to city–year measures of AI activity constructed from the job postings dataset.

[Table 1 about here]

The primary outcome of interest is annual after-tax wage income, expressed in constant 2022 RMB. The mean is 48,398 RMB, with substantial variation across observations (standard deviation approximately 36,500). Men earn significantly more than women, on average, with 55,137 RMB versus 40,347 RMB, resulting in a wage gap of nearly 14,800 RMB, or 36.6 percent in terms of level. In log terms, the difference is 0.38 log points, equivalent to a male wage premium of roughly 46 percent. Both perspectives confirm the persistence of a significant gender wage gap in China's labor market, consistent with prior evidence, and provide a benchmark for assessing whether local AI exposure amplifies or narrows gender gaps.

The key explanatory variables measure local AI intensity, which varies widely across the sample. On average, respondents live in cities with 4,138 AI job postings, 919 distinct firms posting AI positions, and 2,642 unique AI job titles per year. Women, on average, reside in labor markets with higher AI activity than men, particularly in posting firms (984 versus 863) and job titles (2,785 versus 2,523). The firm-level difference is statistically significant, indicating that women in the sample are more concentrated in cities with relatively high AI adoption activity.

Other independent variables capture important demographic and employment characteristics associated with labor market outcomes. Men are, on average, slightly older (42.1 versus 40.2 years) and more likely to hold an agricultural hukou (69% versus 63%), which is linked to restricted access to urban labor markets and social benefits. They also report working longer hours (52.6 hours versus 49.0 hours per week) and a slightly better self-reported health status. These patterns suggest that men's labor market advantage is partly linked to greater labor supply and health capital, in addition to structural differences in hukou status.

Women, by contrast, are more likely to be married (84% versus 82%), employed in the same city where they live (92% versus 84%), and regular internet users (84% versus 82%). These characteristics reflect stronger family responsibilities, less geographic mobility, and greater digital connectedness among women. Internet use is potentially crucial in AI adoption, as it facilitates the flow of information and opportunities for skill acquisition.

Educational attainment is broadly similar across gender groups, with an average of just over 11 years of schooling. This suggests that observed wage differences are unlikely to be explained solely by formal education, but rather by differential access to labor markets, sectoral sorting, and how men and women engage with emerging technologies.

Taken together, these statistics highlight three central features of the data. First, a substantial gender wage gap is evident. Second, local AI activity varies considerably, with women more often located in areas of higher AI penetration. Third, systematic demographic and employment differences across genders will likely influence labor market outcomes. These patterns provide essential context for the regression analysis, which examines how local AI exposure interacts with individual characteristics to affect wages in China.

5 Empirical Strategy and Results

5.1 Empirical Model

We combine individual-level data from the CFPS with city–year measures of local AI labor demand derived from large-scale online recruitment platforms. The CFPS provides detailed information on demographic characteristics, education, employment, wages, working hours, and household background.

Annual after-tax wage income is adjusted to 2022 RMB using the national consumer price index. The dependent variable is the natural logarithm of annual after-tax wage income. We construct three complementary city–year measures of local AI exposure. The first measure, AI job postings, represents the number of job postings containing AI-related keywords, rescaled to thousands. The second measure, AI posting firms, captures the number of unique firms that posted at least one AI-related job, also rescaled to thousands. The third measure, AI job titles, reflects the number of distinct AI-related job titles appearing in postings, rescaled to thousands. Together, these measures capture the intensity, breadth, and diversity of local AI labor demand.

At the individual level, we control for age and its square, years of schooling, gender, marital status, hukou type (agricultural vs. non-agricultural), whether the individual works locally, weekly working hours (log), health status, and internet usage. At the city level, we include GDP per capita, fiscal expenditure as a percentage of GDP, the value-added of the secondary and tertiary sectors, the industrial structure index, the consumption share of GDP, the ratio of higher-education students, the post and telecommunications share of GDP, and the industrial output share of GDP.

Our baseline specification is a two-way fixed effects model:

$$\ln(\text{wage}_{ict}) = \beta AI_{ct}^{m} + X'_{ict}\lambda + Z'_{ct}\delta + C_c + T_t + \varepsilon_{ict}, \tag{1}$$

where *i* indexes individuals, *c* cities, and *t* years. AI_{ct}^{m} refers to one of the three demand measures m∈{postings, firms, titles}. Thus, the key independent variables: AI_{ct}^{posts}: number of AI job postings (×1,000), AI_{ct}^{firms}: number of AI posting firms (×1,000), and $AI_{ct}^{poststitles}$: number of AI job titles (×1,000), where X'_{ict} is the vector of individual-level controls, and Z'_{ct} the vector of city-level controls. City fixed effects (C_c) absorb time-invariant city characteristics, while year fixed effects (T_t) capture common national shocks. Standard errors are clustered at the city level.[7]

In the baseline specification, the city–year AI labor demand variables are measured contemporaneously with the CFPS wage outcomes. The AI demand measures are constructed from online job postings aggregated at the city–year level and matched to the corresponding survey year in the CFPS data. Because wage outcomes in the CFPS are observed through 2022, the empirical analysis uses AI labor demand measures over the same period (2016–2022). This

[7] Conceptually, one could consider alternative empirical setups, such as framing the 2017 New Generation Artificial Intelligence Development Plan as a policy shock in a difference-in-differences design, or exploiting the panel structure of the CFPS via individual fixed effects or first differences. In our setting, however, AI diffusion does not occur as a discrete or staggered treatment across cities; instead, our key regressor is a continuous city–year measure of AI labor demand constructed from job postings. Moreover, CFPS wage outcomes are observed biennially (2016, 2018, 2020, 2022), which precludes a high-frequency event-study around 2017 and leaves no credible treated–control split for a nationwide policy. Panel-based specifications are feasible but would rely on limited within-person exposure variation across four waves—often driven by movers—and may reduce precision. For these reasons, we adopt a city–year exposure framework with rich fixed effects as our primary empirical strategy.

specification captures the local labor demand environment associated with AI-related hiring in the year the wage is observed. Job postings represent active recruitment demand and therefore reflect firms' current hiring needs rather than purely anticipatory investment.

At the same time, wage adjustments to technological adoption may occur gradually as firms integrate new technologies into production. Accordingly, the contemporaneous estimates should be interpreted as reflecting relatively short-run wage responses to changes in local AI labor demand, while longer-run effects may emerge over time as technological adoption diffuses more broadly across firms.

Moreover, because the empirical specification includes city and year fixed effects, identification relies on within-city changes in AI-related labor demand over time, which helps isolate local labor demand shifts associated with AI hiring from broader national trends or persistent cross-city differences. Because AI-related hiring tends to persist within cities over time, contemporaneous and short-horizon lagged measures of AI demand are likely to be highly correlated, suggesting that the baseline specification captures the dominant variation in local AI-related labor demand. Consistent with this interpretation, specifications using lagged AI labor demand measures yield qualitatively similar results (reported in Appendix Tables A3 and A4), indicating that the main findings are not sensitive to alternative timing assumptions.

To account for heterogeneous macroeconomic conditions, we include region- and year-fixed effects interactions and province-specific linear trends. These adjustments absorb differential regional shocks and gradual province-level structural changes. In robustness checks, we further allow for region-specific time trends.

The coefficients of interest, β, measure the semi-elasticity of wages with respect to local AI labor demand. Separate models are estimated for each of the three AI exposure measures to provide a comprehensive view of how different aspects of AI adoption affect wage outcomes.

Because AI labor demand is measured at the city–year level and linked to individual workers through their location in the CFPS data, the estimated coefficients should not be interpreted as capturing the direct effect of individual exposure to AI technologies. Instead, our specification identifies local labor market responses to changes in AI-related labor demand within a city. When AI-related hiring expands in a local labor market, firms may adjust production processes, task allocation, and hiring practices, potentially affecting wages more broadly through changes in labor demand, occupational task complementarities, and labor market competition. The estimated coefficients, therefore, capture average wage spillovers among workers residing in cities experiencing stronger growth in AI demand, rather than treatment effects at the individual or firm level. While city-level aggregation may mask heterogeneity in firm exposure and occupational sorting within cities, and gender differences may partly reflect differential sorting across cities with varying AI demand, this framework captures the local general-equilibrium effects of AI diffusion in urban labor markets, which are the primary focus of our analysis.

Although the empirical specification includes city and year fixed effects, region–year interactions, and province-specific time trends—absorbing many sources of persistent local heterogeneity and broad regional shocks—some endogeneity concerns may remain. In particular, local AI hiring may respond to unobserved city-level productivity shocks, industrial upgrading, or policy initiatives that simultaneously influence wage growth. For example, cities experiencing faster technological upgrading or receiving targeted policy support for digital industries may both attract AI-related hiring and experience rising wages for reasons not directly caused by AI adoption itself. Several features of the setting nevertheless mitigate these concerns. The AI labor-demand measure is constructed from firms' job postings seeking AI-related capabilities, which typically reflect medium- to long-term technology adoption and organizational restructuring rather than short-run wage fluctuations. Moreover, AI adoption in China during the period we study has been strongly shaped by national digitalization initiatives and firm-level technology upgrading, which are unlikely to respond immediately to local wage movements.

To further reduce potential simultaneity, we conduct an additional robustness check using lagged AI labor demand as the key explanatory variable. Specifically, we replace the contemporaneous AI labor-demand measure with its one-year lag. Because the CFPS is conducted biennially and AI data are not available for 2015, this specification excludes the 2016 survey wave and therefore relies on a slightly smaller sample than the baseline regressions. The estimated coefficients remain positive and broadly similar in magnitude to the baseline results, although they are estimated less precisely due to the reduced sample size. The gender-differentiated pattern also remains unchanged, with positive wage effects concentrated among women. The results are reported in Appendix Tables A3 and A4. Nevertheless, because technology adoption and local economic conditions may evolve jointly, the estimated coefficients should be interpreted as associations between local AI labor demand and wage outcomes—consistent with AI-related labor-demand spillovers—rather than strictly causal effects.

Conceptually, one could frame the 2017 State Council plan as a policy shock and estimate a difference-in-differences design. In our setting, however, AI diffusion does not occur as a discrete, staggered treatment at the city level; instead, our key variable is a continuous city–year measure of AI labor demand derived from postings. Moreover, CFPS wage outcomes are observed biennially (2016, 2018, 2020, 2022), which precludes a high-frequency event-study around 2017 and leaves no credible treated–control split for a nationwide policy. Another alternative is to exploit the CFPS panel with individual fixed effects or first differences. While feasible, identification would rely primarily on limited within-person exposure changes across four waves (often driven by movers), and may therefore reduce precision and amplify measurement error in local AI-demand proxies. For these reasons, we adopt a city–year exposure framework with rich fixed effects (city, year, region–year) and province trends, and we treat individual-FE/FD specifications as robustness checks.

5.2 Baseline Estimates of AI Labor Demand on Wage Income

Table 2 presents the main estimates of how local AI labor demand affects wages. The dependent variable is the log of annual after-tax income (CPI-adjusted to 2022 RMB), and the table reports separate regressions using three city–year demand measures: the number of AI job postings, the number of distinct firms posting AI jobs, and the number of different AI job titles. Columns 1–3 report separate regressions using three measures of AI activity at the city–year level: postings, firms, and job titles. The coefficient on AI postings is 0.0019, implying that 1,000 additional postings are associated with a 0.19 percent increase in wages. For firms, the coefficient is 0.0124, corresponding to a 1.24 percent wage increase per 1,000 additional firms. For job titles, the coefficient is 0.0034, or a 0.34 percent increase per 1,000 titles. Scaling by observed variation, a one–standard deviation increase in each measure implies wage gains of roughly 2–3 percent. Therefore, the estimates in Table 2 indicate a consistent and economically meaningful positive association between local AI demand and wages.

[Table 2 about here]

Across all three specifications, the coefficients on AI demand are positive and statistically significant. The point estimates are 0.0019 for postings, 0.0124 for firms, and 0.0034 for job titles (each measured in units of 1,000). Interpreted in levels, an additional 1,000 AI job postings is associated with a 0.19 percent increase in wages, while 1,000 additional firms corresponds to a 1.24 percent increase, and 1,000 new job titles corresponds to a 0.34 percent increase. Alternatively, a one–standard–deviation increase in postings (about 12,400) raises wages by roughly 2.4 percent. For firms, a one–standard deviation increase (2,340) implies a 2.9 percent gain, while for job titles (7,382), the corresponding increase is 2.5 percent. The magnitudes are similar across the three proxies, underscoring that the effects are economically meaningful.

The consistency of results across measures of intensity (postings), breadth (firms), and diversity (titles) suggests that the findings are not driven by measurement choice but reflect a general pattern: cities with more substantial AI-related hiring experience broader wage gains. Among the three proxies, the firm-based measure produces the largest effect. This aligns with the intuition that when more employers adopt AI, competition for labor intensifies and wage pressures diffuse more widely into local markets.

The magnitudes of the estimates, rather than statistical significance alone, warrant attention. A 2–3 percent wage premium associated with one standard deviation of AI hiring is sizable in a labor market where wage growth is modest, and mobility is constrained by hukou and sectoral segmentation. Moreover, these patterns remain after controlling for individual demographics, human capital, employment conditions, city-level economic structure, and an extensive set of fixed effects and trends, suggesting that the results are not driven by omitted location-specific shocks.

Taken together, the results provide robust evidence that AI-related hiring is positively associated with wage growth in Chinese cities. This relationship holds across multiple lenses, including the volume of postings, breadth across firms, and diversity of job titles. The results suggest that the diffusion of AI generates labor-market spillovers that extend well beyond the narrowly defined AI workforce.

6 Heterogeneous Wage Effects Across Gender and Region

6.1 Gendered Wage Responses to AI Labor Demand

Table 3 examines whether the wage effects of AI labor demand differ between men and women. Columns 1–3 interact each AI demand measure with a male dummy in the pooled sample, while Columns 4–9 re-estimate the models separately for men and women. In the pooled regressions, the interaction terms are negative and statistically significant, indicating that the positive association between AI hiring and wages is concentrated among women. The split-sample results confirm this pattern: coefficients for women are positive and significant across all three measures. At the same time, the corresponding estimates for men are small and not statistically different from zero. The estimates in Table 3 therefore suggest that women capture the bulk of the wage gains associated with rising AI labor demand.

[Table 3 about here]

The interaction results reveal a consistent pattern: the positive association between AI demand and wages is concentrated among women. In the pooled regressions, the interaction terms are negative and statistically significant across all three measures, indicating that the wage effects of AI hiring are weaker for men than for women. For example, the baseline coefficient on AI postings is 0.0049, but the male interaction term is –0.0053, leaving no significant net effect for men and a positive, sizable impact for women.

The split-sample regressions reinforce this conclusion. For women, coefficients are positive and significant across postings (0.0046), firms (0.0312), and job titles (0.0080). By contrast, the corresponding estimates are small and statistically indistinguishable from zero for men. Put differently, women capture most of the wage gains associated with the rise in AI labor demand.

How large are these effects in levels? A one-unit increase in AI demand (1,000 postings, firms, or job titles) is associated with wage gains of 0.46 percent for postings, 3.12 percent for firms, and 0.80 percent for job titles among women, while the corresponding effects for men are negligible. These statistically and economically meaningful differences indicate that women enjoy disproportionate wage benefits from AI-related hiring growth.

These findings align with broader evidence that technological change can affect gender wage gaps. One interpretation is that women may be more likely to enter or transition into emerging AI-related roles, or that the diffusion of AI hiring raises demand for complementary tasks in which women are relatively more disproportionately represented. Another possibility is that AI adoption improves transparency or reduces reliance on informal networks, which have historically advantaged men in Chinese urban labor markets. While the mechanisms cannot be identified directly here, the pattern is clear: AI-driven demand growth translated into stronger wage gains for women than men during 2016–2022. These results indicate differential wage responses to AI labor demand across genders rather than direct changes in the overall gender wage gap. We further examine potential mechanisms in Section 7.

6.2 Regional Heterogeneity in the Effects of AI Labor Demand

Figure 4 disaggregates the analysis across China's four major economic regions—Eastern, Central, Western, and Northeastern provinces. We ask whether the wage effects of AI labor demand are evenly distributed or concentrated in particular parts of the country. The figure shows that the effects vary sharply across space. In the Western region, all three AI demand measures are positive and statistically significant, with a one–unit increase (1,000 postings, firms, or job titles) associated with wage gains of 0.3–0.6 percent. In the Eastern region, the coefficients are smaller and statistically significant only for postings, while in the Central and Northeastern regions, the estimates are close to zero and insignificant. The estimates in Figure 4 therefore suggest that wage spillovers from AI labor demand are strongest in the Western provinces, modest in the East, and absent in the Central and Northeastern regions.

[Figure 4 about here]

The results indicate substantial regional heterogeneity. In the Western region, the coefficients on postings, firms, and job titles are all positive and significant. A one–unit increase (1,000 postings, firms, or job titles) corresponds to wage gains of 0.3–0.6 percent, and a one–standard deviation increase implies wage growth of 3–6 percent. These are the most significant effects observed in the country, indicating that inland provinces have generated the strongest wage spillovers from the adoption of AI. Several factors may contribute to this pattern.

In the Eastern region—the most developed part of the country with dense clusters of firms and advanced digital infrastructure—the evidence is modest. Only postings show a significant effect, and the implied wage gains are smaller than those in the West. This suggests that although AI demand is highest in the East, spillovers into wages have been less pronounced.

By contrast, the Central and Northeastern provinces show limited evidence of wage gains from AI hiring. Estimates are close to zero and statistically insignificant across all demand measures. This absence of an effect is consistent with slower economic restructuring and weaker absorptive capacity in these regions.

The stronger wage spillovers observed in Western provinces may reflect several complementary mechanisms. First, baseline wages and productivity levels tend to be lower in less-developed regions, so the marginal productivity gains from adopting new technologies may translate into relatively larger wage increases. Second, the industrial structure of many Western cities remains concentrated in manufacturing, logistics, and service sectors where early applications of AI—such as monitoring systems, process automation, and coordination tools—can complement existing task structures. Third, national development strategies, including the Western Development Strategy and regional technology cluster initiatives, have promoted investment in digital infrastructure and technology adoption in inland areas. These factors are broadly consistent with theories of catch-up growth and absorptive capacity, under which regions that are further from the technological frontier may experience larger marginal gains from adopting emerging technologies.

Taken together, the results underscore that the benefits of AI diffusion are uneven across the country. Wage spillovers are most substantial in the Western provinces, modest in the East, and absent in the Central and Northeastern regions. These patterns underscore the significance of local economic structure and absorptive capacity in influencing the impact of emerging technologies on labor market outcomes.

6.3 Wage Effects in China's AI Cluster Cities

Figure 5 examines whether the effects of AI labor demand on individual wages differ between China's four major AI-cluster cities and other cities. The four clusters are defined as: (1) Jing–Jin–Ji: Beijing, Tianjin, and Hebei in the North; (2) Yangtze River Delta: Shanghai, Nanjing, Hangzhou, and Suzhou in the East; (3) Pearl River Delta: Shenzhen and Guangzhou in the South; and (4) the Southwest–Central corridor: Chengdu, Chongqing, and Wuhan. Together, these regions represent the country's principal AI innovation hubs, characterized by dense industrial agglomerations, advanced digital infrastructure, and strong human capital endowments.

[Figure 5 about here]

The coefficients in Figure 5 show a clear contrast between AI-cluster and non-cluster cities. In AI-cluster cities, all three measures of AI labor demand are positive and highly significant. A one-unit increase (1,000 AI-related postings, firms, or job titles) is associated with wage gains of 0.42 percent for postings, 3.9 percent for firms, and 0.8 percent for job titles, each significant at the 1 percent level. These magnitudes are economically meaningful and consistent with the notion that AI adoption raises local productivity and is reflected in wages when complementary production factors (e.g., skilled labor, capital, and digital capabilities) are abundant.

In contrast, the coefficients for non-cluster cities are statistically indistinguishable from zero across all three measures. Although the point estimates (0.021, 0.079, and 0.034) appear larger in magnitude, their standard errors are wide, suggesting a non-robust wage response to AI demand. From an economic perspective, this indicates that outside the main AI hubs, local labor markets have not yet developed the necessary absorptive capacity to translate AI-related hiring into higher earnings. The lack of significant effects in these regions may reflect weaker digital infrastructure, smaller pools of skilled labor, and limited linkages between AI-adopting firms and the broader urban economy.

Taken together, the estimates in Figure 5 highlight a sharp spatial divide in how AI diffusion affects wages. AI-cluster cities are the only places where AI labor demand reliably translates into wage growth, whereas non-cluster regions show no detectable impact. The pattern underscores the role of agglomeration and technological ecosystems: in China's leading AI hubs, dense networks of firms and skilled workers amplify the wage benefits of AI adoption, while in the periphery, the gains from AI hiring remain latent or unrealized.

6.4 Discussion and Interpretation of Findings

Across three measures of AI demand, two genders, four regions, and four industrial clusters, our results present a consistent picture: local AI labor demand is positively associated with individual wage income in urban China. Across all three measures—AI postings, posting firms, and job titles—the coefficients are positive and statistically significant, indicating that the findings are not sensitive to how AI intensity is measured. The magnitudes are economically meaningful: a one-unit increase in local AI hiring (1,000 postings, firms, or job titles) is associated with wage gains of roughly 0.2–0.3 percent, suggesting that the spread of AI technologies is translating into tangible labor-market rewards.

Beyond these aggregate effects, the analysis reveals important distributional and spatial heterogeneity. First, the wage effects are strongly gendered. The gains from rising AI demand accrue disproportionately to women, who experience statistically significant wage increases of about 0.5–0.7 percent for a one-unit rise in AI hiring. At the same time, men show little to no benefit. This pattern suggests that AI diffusion may help narrow gender wage gaps in certain segments of the labor market, possibly by expanding employment opportunities in emerging occupations or by increasing the value of interpersonal, cognitive, and organizational skills, which are more commonly held by women.

Second, the effects vary sharply across regions. Western provinces exhibit the largest and most robust wage gains, resulting from active industrial upgrading and the faster absorption of new technologies. Eastern provinces exhibit modest effects, primarily concentrated in postings, whereas the Central and Northeastern regions display no measurable impact. These differences underscore that local absorptive capacity—reflected in human-capital endowments, digital infrastructure, and industrial structure—conditions how effectively AI adoption translates into higher wages.

Third, the evidence from AI-cluster cities highlights the importance of technological ecosystems. In China's major innovation hubs—Jing–Jin–Ji (Beijing, Tianjin, and Hebei), Yangtze River Delta, Pearl River Delta, and the Southwest–Central corridor (Chengdu, Chongqing, and Wuhan)—AI labor demand has consistently positive and statistically significant wage effects. Outside these clusters, the estimated coefficients are statistically indistinguishable from zero. This spatial divide suggests that AI-related wage gains are concentrated where complementary production factors and network externalities are strongest. The results, therefore, align with the broader literature on agglomeration and technological adoption (Glaeser et al., 1992; Moretti, 2010), which emphasizes how dense innovation environments amplify the local returns to new technologies.

Taken together, these findings underscore that the diffusion of AI is not merely a firm-level phenomenon but a geographically and demographically uneven process. The wage benefits of AI adoption are concentrated among women, in the Western provinces, and in more established AI clusters.

7 Why Do Female Workers Benefit More? Occupational Structure, Task Composition, and the Nature of AI Diffusion

A central finding of this paper is that wage gains associated with rising local AI labor demand accrue primarily to women, while the corresponding effects for men are generally small and statistically insignificant. Given that men are more heavily represented in technical and high-skill occupations, this pattern may appear counterintuitive. If AI is a high-technology phenomenon, why do male-dominated technical roles not exhibit stronger wage spillovers?

This apparent paradox reflects the nature of the AI exposure captured in our data. Our city–year AI measures are constructed from online job postings and therefore reflect the intensity and breadth of AI-related hiring across local labor markets rather than frontier AI research or advanced algorithmic innovation alone. During the 2016–2022 period, AI diffusion in China primarily involved integrating AI tools into existing organizational workflows—automating routine components, improving data processing, and augmenting decision-support systems—rather than exclusively expanding high-end AI engineering roles.

Under a task-based framework of technological change, wage effects depend not simply on formal education or skill level, but on how AI interacts with the task composition of occupations. If AI complements coordination, communication, monitoring, and customer-facing activities while automating routine components, productivity gains will emerge in occupations where such task bundles are prevalent. Because men and women are unevenly distributed across occupations and task structures, the same technological shock can generate gender-differentiated wage responses within the same local labor market.

To examine this mechanism directly, Figure 6 presents occupation-specific gender differences in the marginal effect of AI labor demand on log wages across seven occupational groups aggregated from the CFPS: (1) Clerical and Administrative, (2) Sales and Hospitality, (3) Healthcare, (4) Education, (5) Finance and Professional Services, (6) Production and Manufacturing, and (7) Public, Security, and Other Services. The figure plots the interaction term (Female × AI), which captures the difference in the marginal wage response to AI labor demand between women and men (β_3). Appendix Table A2 reports the full decomposition: the marginal effect for men (β_1), the implied effect for women ($\beta_1 + \beta_3$), and the gender difference (β_3).

[Figure 6 about here]

The estimates show a clear and economically meaningful pattern. Because the dependent variable is log wage and AI variables are measured in units of 1,000 at the city–year level, coefficients can be interpreted as approximate percentage changes in wages associated with an increase of 1,000 AI job postings (or posting firms or job titles).

Using AI job postings (Panel 1 of Figure 6), Sales and Hospitality exhibits a gender difference of 0.0118. Appendix Table A2 shows that the marginal effect for men in this occupation is −0.0053 (statistically insignificant), while the implied effect for women is 0.0065. Thus, a 1,000-unit increase in AI job postings is associated with approximately a 0.65 percent wage increase for women, compared with no statistically significant change for men. The 1.18 percentage-point differential reflects the stronger wage response for women relative to men.

A similar pattern appears in Production and Manufacturing. The interaction coefficient is 0.0129, while the marginal effect for men is −0.0032 and statistically insignificant. The implied female effect is 0.0097, corresponding to roughly a 1 percent wage increase per 1,000 additional AI postings. Clerical and Administrative occupations also show a positive differential (0.0036), with women's wages rising modestly (about 0.22 percent) and men's wages again showing no measurable response.

These results are robust across alternative measures of AI labor demand. When AI exposure is measured using the number of AI posting firms (Panel 2), magnitudes are larger but the ordering across occupations is unchanged. For example, in Sales and Hospitality, the gender difference is 0.0645; the marginal effect for men is −0.0137 (insignificant), while the implied female effect is 0.0508, corresponding to a 5.1 percent wage increase per 1,000 additional AI-posting firms. Production and Manufacturing shows a comparable pattern, with a gender difference of 0.0690 and an implied female effect of 0.0450. Results using AI job titles (Panel 3) confirm similar patterns, reinforcing that the gender asymmetry is not driven by the specific AI measure employed.

In contrast, analytically intensive and credential-heavy occupations display a different pattern. In Education, the marginal effect for men is positive and statistically significant (0.0213), implying a 2.1 percent wage increase per 1,000 additional AI postings. The interaction coefficient is −0.0124, indicating that women's wage response is 1.24 percentage points smaller than men's. The implied female effect remains positive (0.0089), corresponding to approximately a 0.9 percent wage increase, but the gains are substantially weaker than those for men. Similar negative differentials appear when AI exposure is measured using posting firms or job titles.

In Finance and Professional Services, the marginal effect for men is small and statistically insignificant, while the gender difference is positive (0.0085 for job postings). The implied female effect is approximately 0.0107, indicating modest but positive wage gains for women. Healthcare shows positive but imprecisely estimated effects for both genders, likely reflecting the smaller sample size in that occupation.

Taken together, the results do not suggest a universal female advantage from AI exposure. Instead, wage responses vary systematically across occupations. In most service-oriented and operational occupations—Clerical and Administrative, Sales and Hospitality, Production and Manufacturing, and Public/Security—men's wages show little measurable response to AI labor demand, whereas women experience statistically significant gains. In Education, both genders benefit, but men benefit more strongly. This systematic variation indicates that the gender asymmetry arises from differences in occupational task composition rather than from uniform changes in returns to education.

Occupations with the largest female wage gains combine routine operational components with substantial coordination, communication, and interpersonal tasks. In these settings, AI adoption may automate repetitive elements while enhancing the productivity of monitoring, customer interaction, and multi-task coordination activities. Conversely, in analytically intensive occupations such as Education, AI may complement higher-level cognitive and technical tasks, generating relatively larger gains for workers more concentrated in those task bundles.

In addition to task complementarity, occupational sorting and transitions may also contribute to gender-differentiated wage responses to rising AI labor demand. If workers move toward occupations that benefit more from AI, gender differences in mobility patterns could reinforce the observed wage effects. Recent cross-country studies suggest that women are often more concentrated in occupations with relatively high exposure to AI-related tasks, particularly clerical, administrative, and service roles (Demombynes et al., 2025, Pizzinelli et al., 2023). Some evidence also indicates that higher-wage workers—especially women—may transition from occupations at greater risk of automation to roles that are more complementary to AI technologies (Cazzaniga et al., 2025). Although our data do not allow us to track detailed occupational transitions over time, the occupation-specific estimates in Figure 6 indicate that gender differences in wage responses arise largely within broad occupational groups rather than through a uniform shift across occupations.[8] This pattern suggests that the interaction between

[8] Several factors suggest that large AI-driven sorting is unlikely to be the primary explanation for the patterns we document. The period we study (2016–2022) corresponds to an early phase of AI diffusion in China, during which adoption largely involved integrating AI tools into existing workflows rather than rapidly expanding new AI-specific occupations. Transitions into AI-intensive roles typically require specialized training and skill accumulation, while occupational gender segregation tends to

AI diffusion and occupational task composition is an important channel shaping wage outcomes, although gradual occupational reallocation may also play a complementary role.

These findings help resolve the earlier puzzle regarding the absence of average wage gains for men. During the 2016–2022 period, the primary margin of adjustment appears not to have been confined to frontier AI engineering roles, but rather the broader diffusion of AI-enabled tools into organizational processes. Wage spillovers were therefore strongest in occupations where AI augmented coordination and service-related tasks—roles in which women are disproportionately represented.[9] The gender-differentiated effects reflect structural differences in occupational sorting and task exposure rather than intrinsic differences in technological capability.

Overall, Figure 6 and Appendix Table A2 demonstrate that the gender asymmetry in AI wage effects is mediated by occupational structure and task complementarity. Women experience stronger wage gains when AI enhances productivity in service and operational roles, while men experience relatively larger gains in occupations where AI complements analytical and technical expertise. This interpretation aligns with the broader literature on technological change, which emphasizes that the distributional consequences of new technologies depend less on formal education per se than on the specific mix of tasks that remain complementary to, or are displaced by, automation (Acemoglu and Restrepo, 2018, Autor, 2013).

---

evolve gradually over time. Geographic mobility responses are also likely to be limited given migration frictions in China, including housing costs, family considerations, and the hukou system. Moreover, because the specification includes city and year fixed effects, the estimates are identified from within-city changes over time, which further mitigates concerns that the results reflect compositional differences across locations.

[9] While the results indicate that women experience stronger wage gains than men in several occupations when local AI labor demand increases, these estimates should not necessarily be interpreted as implying that AI diffusion eliminates the overall gender wage gap. The regressions capture gender differences in wage responses to changes in local AI-related labor demand rather than direct changes in the aggregate gender wage gap. Although stronger wage growth for women in certain occupations may contribute to narrowing gender wage disparities at the margin, Table 1 indicates that a substantial gender wage gap remains in the sample, and the effects are not uniform across occupations. More broadly, the overall gender wage gap reflects structural and institutional factors—including persistent occupational segregation, differences in career trajectories and labor-force participation, and other labor-market frictions—that tend to evolve gradually over time.

In sum, the occupation-specific results indicate that the gender asymmetry in AI-related wage effects is closely linked to differences in task composition across occupations. While our analysis does not directly observe task content at the micro level, the systematic variation across service-oriented, production, and analytically intensive occupations is consistent with a task-complementarity interpretation. The stronger wage response of women to local AI labor demand therefore appears to reflect the interaction between AI diffusion and occupational task structure. The gender-specific effects documented in Section 6.1 likely operate through structural differences in task exposure across occupations rather than through uniform shifts in returns to skill.

8 Conclusion

This paper presents new evidence on the impact of the rise of AI on labor market outcomes in China. By linking a large dataset of AI-related job postings with microdata from the China Family Panel Studies, we show that local AI labor demand is positively associated with individual wages. The effects are consistent across three complementary measures—AI job postings, posting firms, and job titles—each capturing a distinct dimension of AI adoption. A one-unit increase in AI labor demand (1,000 postings, firms, or job titles) is associated with wage gains of roughly 0.2–0.3 percent, a magnitude that is both statistically robust and economically meaningful.

The analysis also reveals substantial heterogeneity. First, women benefit disproportionately from rising AI labor demand: for women, a one-unit increase in AI hiring is associated with wage gains of 0.5–0.7 percent, whereas the corresponding effects for men are negligible. These results suggest that AI diffusion may help narrow gender wage gaps by raising returns to interpersonal and coordination tasks in which women are more strongly represented. Second, the wage effects also vary sharply across space. The Western provinces exhibit the largest and most robust gains, while the Eastern region shows modest effects, and the Central and Northeastern regions display none. Third, the positive wage impacts are concentrated in China's major AI-cluster cities— the Jing–Jin–Ji region, Yangtze River Delta, Pearl River Delta, and the Southwest–Central corridor—where complementary production factors such as skilled labor, digital infrastructure, and innovative firms are well developed. Outside these clusters, AI labor demand shows no statistically detectable influence on wages.

We further examine why female workers gain more from AI diffusion. The evidence suggests that occupational skill complementarity is the key mechanism: women's wage gains primarily arise in service-oriented and less skill-intensive occupations, such as clerical, sales, and production jobs, where AI automates routine tasks but enhances the value of human interaction, communication, and adaptability. In contrast, in high-skill professional or educational occupations, where AI complements analytical or informational tasks, the relative advantage shifts in favor of men. These results underscore that the gender asymmetry in AI wage effects reflects differences in task composition rather than formal education alone.

Several limitations should be acknowledged. First, the observation window of this study (2016–2022) ends just before the launch of large language models such as ChatGPT in late 2022, which marked the beginning of a new wave of rapid growth in AI applications—particularly generative AI—across industries. Consequently, our results do not capture the post-2023 acceleration in AI adoption and its potential impact on the labor market. Second, the online job postings data may not fully reflect all forms of AI adoption, especially among smaller or informal firms that are underrepresented on digital platforms. Despite these limitations, the findings reveal an apparent empirical regularity: where AI-related hiring expands, wages rise—

but the gains are distributed unevenly across gender, region, and occupation. Future research should explore how these dynamics evolve in the era of generative AI, investigating the roles of occupational re-sorting, firm upgrading, and digital infrastructure in shaping who benefits from technological change and the magnitudes of the effects. As China continues to scale up its AI capacity, understanding these distributional channels will be crucial for designing policies that promote both technological progress and inclusive growth.

References

Acemoglu, Daron. 2025. “The Simple Macroeconomics of Artificial Intelligence.” *NBER Working Paper 30734.*

Acemoglu, Daron, Erhan Dinlersoz, and Robert Seamans. 2024. “Artificial Intelligence and Business Dynamics.” *U.S. Census Center for Economic Studies Working Paper 24-16.*

Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review* 108 (6): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." Journal of Political Economy 128 (6): 2188–2244.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, eds. 2019. *The Economics of Artificial Intelligence: Implications for the Future of Work.* Chicago: University of Chicago Press.

Alderucci, Paolo, Benjamin Jones, Nicholas Bloom, and Erik Brynjolfsson. 2020. “Artificial Intelligence as a General Purpose Technology: Firm-Level Evidence from U.S. Patents.” *MIT–NBER Conference Paper.*

Autor, David H. 2013. "The 'Task Approach' to Labor Markets: An Overview." Journal for Labor Market Research 46 (3): 185–199.

Autor, David H., David Mindell, and Elisabeth Reynolds. 2024. *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines.* Cambridge, MA: MIT Press.

Brynjolfsson, Erik, Dan Rock, and Chad Syverson. 2023. “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies.” *American Economic Journal: Macroeconomics* 15 (1): 333–367.

Cazzaniga, Mauro; Augustus Panton; Longji Li; Carlo Pizzinelli and Marina M. Tavares. 2025. "A Gender Lens on Labor Market Exposure to Ai." *AEA Papers and Proceedings*, 115, 56–61.

Copestake, James, Rupa Chanda, and Srinivas Thiruvadanthai. 2024. “The Effect of Artificial Intelligence on Job Postings and Wages in India.” *Economics Letters* 234 (1): 110–124.

Damioli, Giuseppe, Valerio Van Roy, and Daniel Vertesy. 2021. “The Impact of Artificial Intelligence on Labor Productivity.” *Industrial and Corporate Change* 30 (4): 1091–1115.

Demombynes, Gabriel; Jorg Gero Langbein and Michael Weber. 2025. The World Bank,

Dinlersoz, Erhan, Lucia Foster, C. J. Krizan, and Roberto Samaniego. 2024. “AI, Firm Dynamics, and Productivity in the United States.” *Center for Economic Studies Discussion Paper CES-WP-24-16.*

Glaeser, Edward L., Hedi D. Kallal, Jose A. Scheinkman, and Andrei Shleifer. 1992. "Growth in Cities." Journal of Political Economy 100 (6): 1126–1152.

Josifidis, Kosta, and Jelena Supic. 2024. "Artificial Intelligence and the Productivity–Pay Gap in the USA: Industrial Insights and the Revival of Heterodox Ideas." *Structural Change and Economic Dynamics* 68 (1): 540–552.

Li, Ying, and Qian Li. 2023. "Artificial Intelligence, Automation, and Labor Market Upgrading in Chinese Manufacturing." *China Economic Review* 80 (December): 102051.

Microsoft and LinkedIn. 2024. "2024 Work Trend Index Annual Report: Ai at Work Is Here—Now Comes the Hard Part," Microsoft,

Moretti, Enrico. 2010. "Local Multipliers." American Economic Review 100 (2): 373–377.

Pew Research Center. 2023. "Which U.S. Workers Are Exposed to AI in Their Jobs?" Washington, DC.

Pizzinelli, Carlo; Augustus J Panton; Marina Mendes Tavares; Mauro Cazzaniga and Longji Li. 2023. "Labor Market Exposure to Ai: Cross-Country Differences and Distributio Nal Implications," *IMF Working Papers.* 1.

Septiandria, Fajar, Efthymios Constantinides, and Daniele Quercia. 2024. "Artificial Intelligence Exposure Across Occupations: A Semantic Approach." *PNAS Nexus* 3 (2): 1–13.

Sun, Xiaoyu, Yifan Zhang, and Ying Fan. 2024. "AI Adoption, Firm Productivity, and Employment Structure: Evidence from Chinese Listed Companies." *Journal of Asian Economics* 89 (2): 101625.

Sytsma, Donald, and Jonathan Sousa. 2024. *Artificial Intelligence and the Labor Force.* Santa Monica, CA: RAND Corporation.

Webb, Michael. 2020. "The Impact of Artificial Intelligence on the Labor Market." *Stanford University Job Market Paper.*

Xie, Yu, and Jingwei Hu. 2014. "An Introduction to the China Family Panel Studies (CFPS)." Chinese Sociological Review 47 (1): 3–29.

Zhaopin. 2022. "Report on Talent Demand and Development Environment of Listed Companies," Beijing: Zhaopin,

Table 1 Summary Statistics, 2016–2022 (overall and by gender)

Variables:	All	Male (1)	Female (2)	Diff. (1)−(2)
Wage Income (Annual, 2022 RMB)	48398.38	55137.15	40346.55	14790.60***
	[36499.47]	[37765.34]	[33183.93]	(863.88)
Log Wage Income (Annual, 2022 RMB)	10.56	10.73	10.35	0.38***
	[0.78]	[0.71]	[0.81]	(0.02)
AI Job Postings (city-year)	4138.10	3927.28	4390.01	-462.73
	[12429.16]	[12272.51]	[12611.11]	(300.31)
AI Posting Firms (city-year)	918.63	863.13	984.94	-121.81**
	[2339.59]	[2277.92]	[2409.90]	(56.52)
AI Job Titles (city-year)	2642.35	2523.24	2784.67	-261.43
	[7382.02]	[7328.79]	[7443.79]	(178.36)
Age	41.21	42.07	40.18	1.89***
	[10.12]	[10.35]	[9.74]	(0.24)
Gender (1=male, 0=female)	0.54	1.00	0.00	1.00
	[0.50]	[0.00]	[0.00]	-
Married (1=yes, 0=no)	0.83	0.82	0.84	-0.02**
	[0.37]	[0.38]	[0.36]	(0.01)
Years of Schooling	11.13	11.08	11.18	-0.10
	[3.87]	[3.71]	[4.05]	(0.09)
Agricultural Hukou (1=yes, 0=no)	0.66	0.69	0.63	0.05***
	[0.47]	[0.46]	[0.48]	(0.01)
Workplace in Same City (1=yes, 0=no)	0.88	0.84	0.92	-0.08***
	[0.33]	[0.37]	[0.27]	(0.01)
Hours Worked Per Week	50.96	52.58	49.01	3.57***
	[15.79]	[15.89]	[15.45]	(0.38)
Health Status (1=Poor, 5=Excellent)	3.16	3.23	3.08	0.16***
	[1.04]	[1.06]	[1.01]	(0.03)
Internet Use (1=yes, 0=no)	0.83	0.82	0.84	-0.02*
	[0.38]	[0.38]	[0.37]	(0.01)
N	6905	3759	3146	

Notes: Wage income is measured annually, after tax, and is CPI-adjusted to 2022 RMB. Log wage income is the natural logarithm of this measure. AI job postings, AI posting firms, and AI job titles are measured at the city–year level. "AI posting firms" refers to the number of distinct firms with AI-related postings, and "AI job titles" refers to the number of distinct AI-related job titles. Internet use denotes access via a mobile or computer. Summary statistics are based on individual-level observations. Monetary values are expressed in constant 2022 RMB. Brackets report standard deviations and parentheses are standard errors. * p<0.10, ** p<0.05, *** p<0.01.

Table 2 Effects of AI Job Postings, Firms, and Titles on Wage Outcomes, 2016–2022

Dependent variable: log wage			
AI job postings	0.0019**		
	(0.0007)		
AI posting firms		0.0124**	
		(0.0057)	
AI job titles			0.0034**
			(0.0015)
Age	0.0590***	0.0590***	0.0590***
	(0.0077)	(0.0077)	(0.0078)
Age squared	-0.0007***	-0.0007***	-0.0007***
	(0.0001)	(0.0001)	(0.0001)
Years of schooling	0.0589***	0.0589***	0.0589***
	(0.0034)	(0.0034)	(0.0034)
Gender (1=male)	0.3959***	0.3960***	0.3959***
	(0.0313)	(0.0313)	(0.0313)
Married (1=yes)	0.0467	0.0469	0.0466
	(0.0292)	(0.0291)	(0.0292)
Agricultural hukou (1=yes)	-0.0535**	-0.0534**	-0.0538**
	(0.0267)	(0.0267)	(0.0267)
Workplace in same city (1=yes)	-0.1847***	-0.1847***	-0.1846***
	(0.0309)	(0.0309)	(0.0309)
Log hours wrked per week	0.1383***	0.1383***	0.1383***
	(0.0274)	(0.0274)	(0.0274)
Health status (ref. = Poor)			
Fair	0.0479	0.0480	0.0479
	(0.0372)	(0.0372)	(0.0372)
Good	0.1062***	0.1062***	0.1063***
	(0.0304)	(0.0304)	(0.0304)
Very good	0.1357***	0.1357***	0.1358***
	(0.0314)	(0.0314)	(0.0313)
Excellent	0.0917**	0.0917**	0.0918**
	(0.0385)	(0.0385)	(0.0385)
Internet use (1=yes)	0.1350***	0.1349***	0.1348***
	(0.0232)	(0.0232)	(0.0232)
City & year fixed effects	Yes	Yes	Yes
Region-year fixed effects	Yes	Yes	Yes
Province time trends	Yes	Yes	Yes
City-level controls	Yes	Yes	Yes
Within R^2	0.246	0.246	0.246
N	6905	6905	6905

Note: AI job postings, AI posting firms, and AI job titles are measured at the city–year level and expressed in units of 1,000. Wage income is measured annually, after tax, and is CPI-adjusted to 2022 RMB. Log wage income is the natural logarithm of this measure. “AI posting firms” refers to the number of distinct firms with AI-related postings, and “AI job titles” refers to the number of distinct AI-related job titles. Internet use denotes access via a mobile or computer. All regressions control for include industry dummies and a set of city-level characteristics, including GDP per capita, fiscal expenditure as a share of GDP, secondary and tertiary industry value added, industrial structure index, higher education student ratio, and post and telecommunications as a share of GDP. Standard errors are in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table 3 Effects of AI Job Postings, Posting Firms, and Job Tittles on Wage Outcomes: Interaction vs. Gender-Specific Models

Dependent variable: log wage	Pooled Sample + Interaction			Male Sample			Female Sample		
	AI posts	AI firms	AI titles	AI posts	AI firms	AI titles	AI posts	AI firms	AI titles
AI job postings	0.0049***			-0.0004			0.0046***		
	(0.0009)			(0.0013)			(0.0013)		
Gender (1=male)	0.4203***	0.4274***	0.4229***						
	(0.0234)	(0.0234)	(0.0233)						
Gender*AI job postings	-0.0053***								
	(0.0012)								
AI posting firms		0.0292***			-0.0066			0.0312***	
		(0.0060)			(0.0095)			(0.0102)	
Gender*AI posting firms		-0.0309***							
		(0.0043)							
AI job titles			0.0088***			-0.0004			0.0080***
			(0.0016)			(0.0026)			(0.0026)
Gender*AI job titles			-0.0092***						
			(0.0019)						
Within R2	0.248	0.249	0.248	0.180	0.180	0.180	0.214	0.214	0.214
N	6905	6905	6905	3759	3759	3759	3146	3146	3146

Note: AI job postings, AI posting firms, and AI job titles are measured at the city–year level and expressed in units of 1,000. Wage income is measured annually, after tax, and is CPI-adjusted to 2022 RMB. Log wage income is the natural logarithm of this measure. "AI posting firms" refers to the number of distinct firms with AI-related postings, and "AI job titles" refers to the number of distinct AI-related job titles. Internet use denotes access via a mobile device or computer. All regressions include individual-level controls (age, age squared, years of schooling, marital status, agricultural hukou, workplace in same city, log hours worked per week, health status, and internet use), as well as city and year fixed effects, region–year fixed effects, province time trends, and city-level controls (GDP per capita, fiscal expenditure as a share of GDP, secondary and tertiary industry value added, industrial structure index, higher education student ratio, and post and telecommunications as a share of GDP). Standard errors are in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Figure 1 Total AI-Related Job Postings by Year

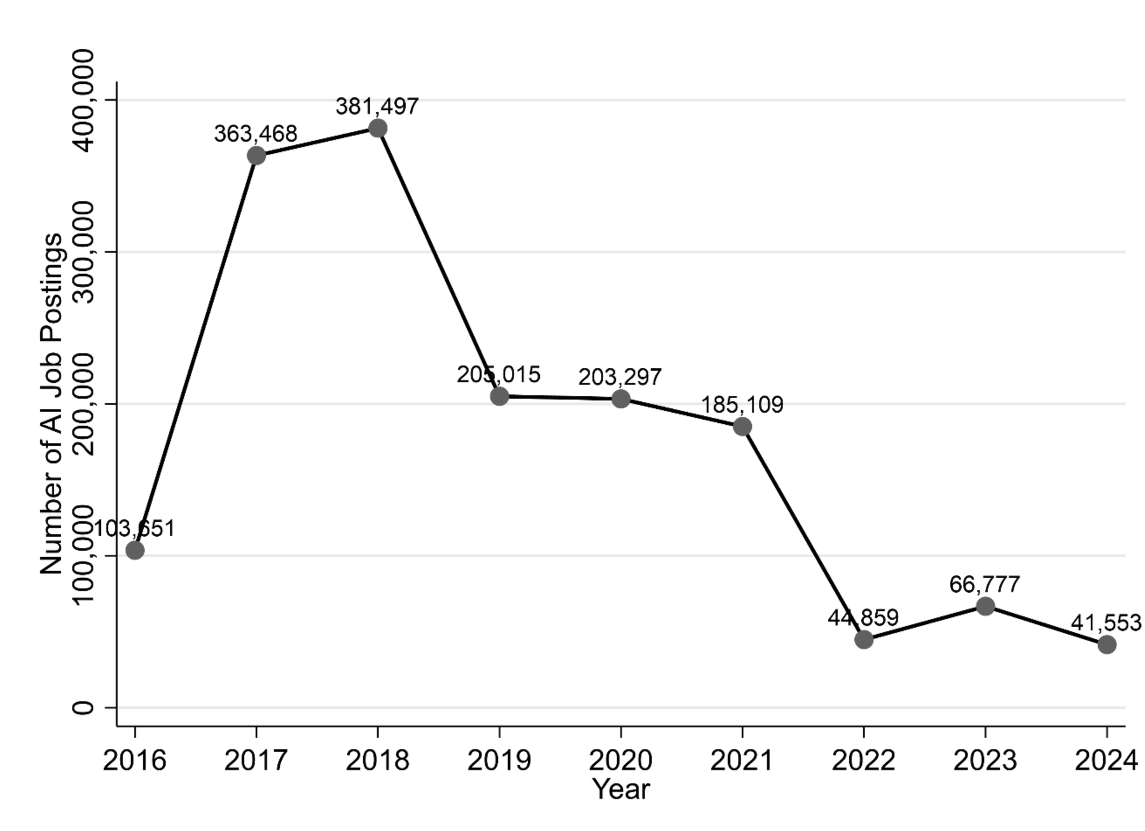

Note: The figure plots the total number of AI-related job postings in China from 2016 to 2024, based on postings collected from major online recruitment platforms.

Source: Data collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58同城), and Ganji.com (赶集网). AI-related postings are identified using a curated list of 73 keywords (e.g., artificial intelligence, deep learning, facial recognition, autonomous driving, neural networks).

Figure 2 Number of Firms Posting AI Jobs by Year

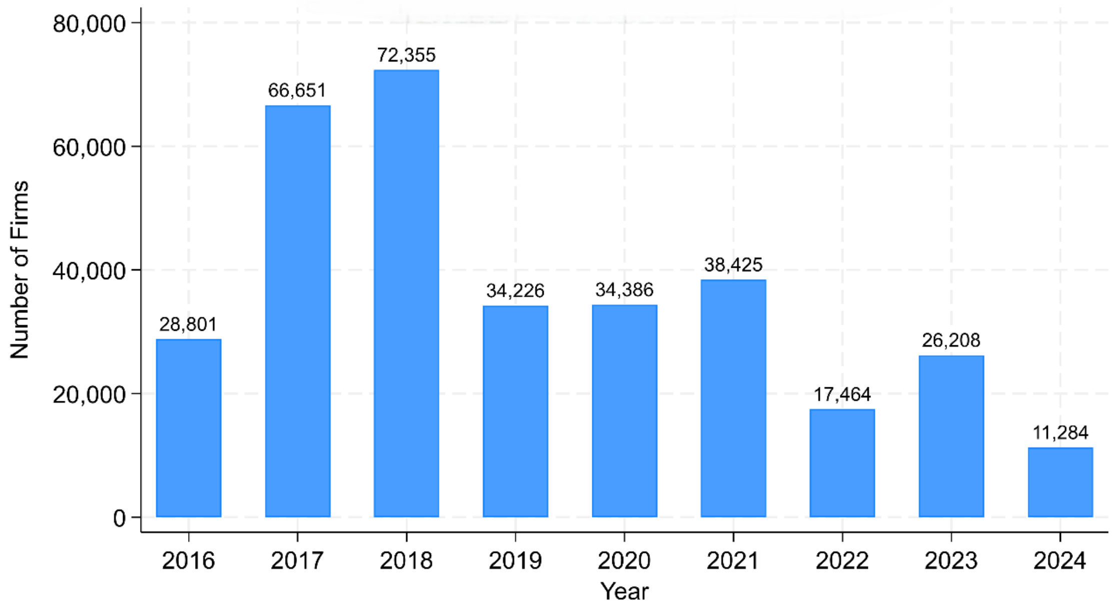

Note: The figure shows the number of distinct firms posting AI-related jobs in China between 2016 and 2024.

Source: Data collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58同城), and Ganji.com (赶集网). AI-related postings are identified using a curated list of 73 keywords (e.g., artificial intelligence, deep learning, facial recognition, autonomous driving, neural networks).

Figure 3 Top AI Clusters in China

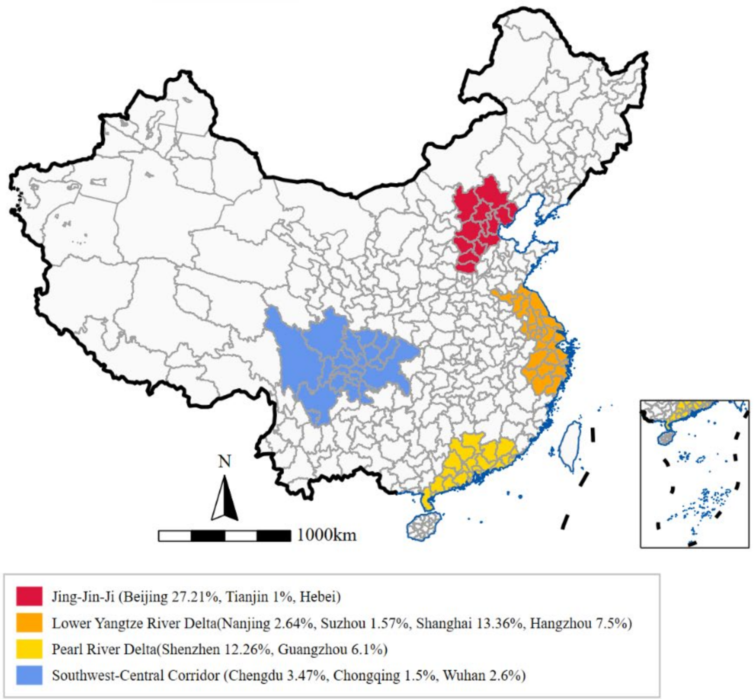

Source: Data collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58同城), and Ganji.com (赶集网). AI-related postings are identified using a curated list of 73 keywords (e.g., artificial intelligence, deep learning, facial recognition, autonomous driving, neural networks).

Figure 4 Effects of AI Job Postings, Posting Firms, and Job Tittles on Wage Outcomes: By Region

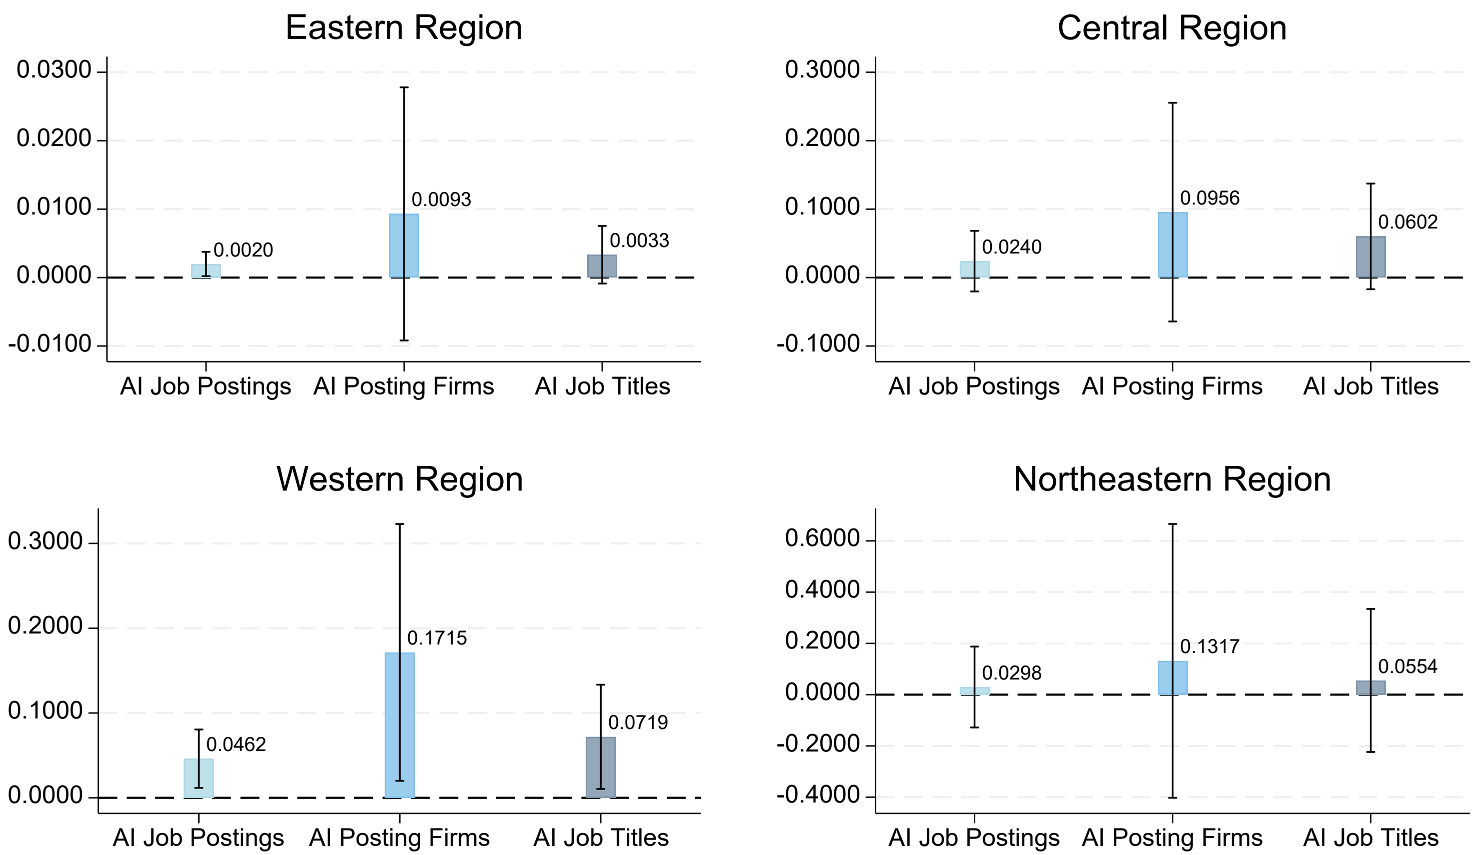

Note: AI job postings, AI posting firms, and AI job titles are measured at the city–year level and expressed in units of 1,000. Range bars indicate 95 percent confidence intervals. Following the classification of the National Bureau of Statistics of China, provinces are grouped into four major economic regions: Eastern (Beijing, Tianjin, Hebei, Shanghai, Jiangsu, Zhejiang, Fujian, Shandong, Guangdong, Hainan), Central (Shanxi, Anhui, Jiangxi, Henan, Hubei, Hunan), Western (Inner Mongolia, Guangxi, Chongqing, Sichuan, Guizhou, Yunnan, Tibet, Shaanxi, Gansu, Qinghai, Ningxia, Xinjiang), and Northeastern (Liaoning, Jilin, Heilongjiang).

Figure 5 Wage effects of AI labor demand by AI-cluster and non-cluster cities, 2016–2022

Note: AI job postings, AI posting firms, and AI job titles are measured at the city–year level and expressed in units of 1,000. Range bars indicate 95 percent confidence intervals. AI cluster cities are defined as four major regional innovation hubs: (1) Jing-Jin-Ji: Beijing, Tianjin, and Hebei in the north, (2) Yangtze River Delta: Shanghai, Nanjing, Hangzhou, Suzhou in the East, (3) Pearl River Delta: Shenzhen and Guangzhou in the South, and (4) Southwest-Central China: Chengdu, Chongqing, and Wuhan.

Figure 6 Differential Effects of AI Labor Demand on Female Relative to Male Wages, by Occupation

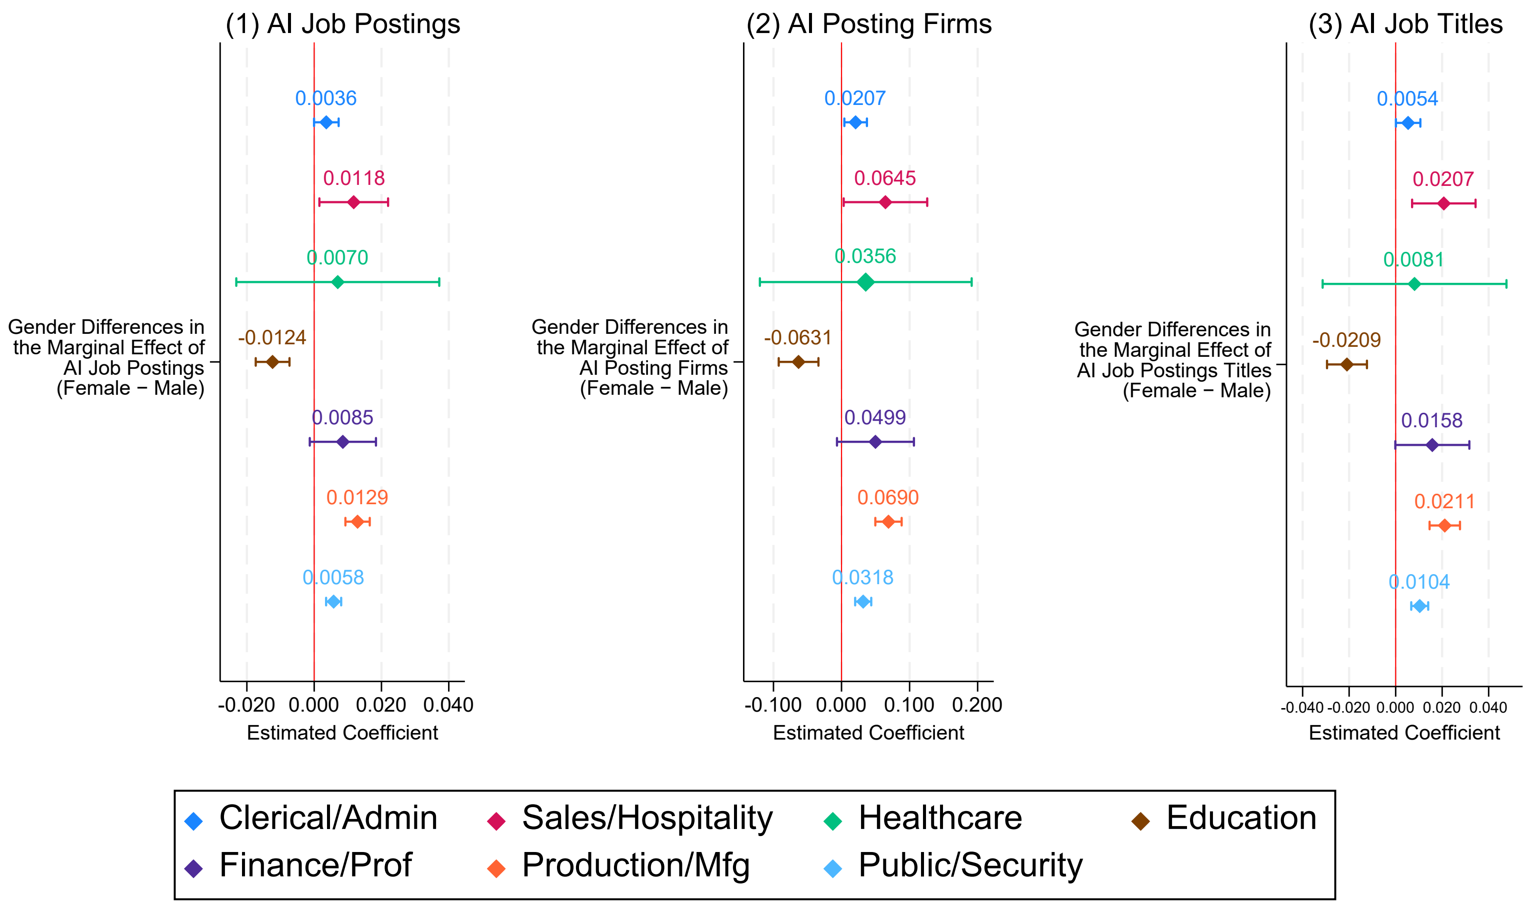

Note: Coefficients represent the estimated interaction term (Female × AI) from occupation-specific regressions of log wages on AI labor demand. AI variables (AI job postings, AI posting firms, and AI job titles) are measured at the city–year level in units of 1,000. Coefficients therefore reflect percentage-point differences in the marginal wage response to AI labor demand between women and men. The baseline effect for men (β_1) and the implied total effect for women ($\beta_1 + \beta_3$) are reported in Appendix Table A2. Standard errors are clustered at the city level; range bars indicate 95 percent confidence intervals. The seven occupational groups are aggregated from the CFPS occupational classification as follows: (1) Clerical and Administrative: administrative officers, clerks, accountants, and related office staff. (2) Sales and Hospitality: salespersons, service workers, restaurant and hotel staff, and related occupations. (3) Healthcare: physicians, nurses, pharmacists, and other health professionals. (4) Education: primary, secondary, and higher education teachers and other education staff. (5) Finance and Professions: financial officers, legal professionals, engineers, and other technical specialists. (6) Production and Manufacturing: workers in manufacturing, crafts, assembly, and related industries. (7) Public, Security, and Other Services: government and public-sector employees, police, security, and other uncategorized occupations.

Appendix Table A1 Keywords Used to Identify AI-related Jobs

Category	Keywords
Core AI Technologies	Artificial Intelligence, Machine Learning, Deep Learning, Neural Networks, Recurrent Neural Networks, Deep Neural Networks, Generative Adversarial Networks (GANs), Reinforcement Learning, Feature Extraction, Feature Recognition, Pattern Recognition, Knowledge Representation, Knowledge Graph, Knowledge Graph Representation, Support Vector Machine (SVM), Long Short-Term Memory (LSTM)
Natural Language & Speech	Machine Translation, Natural Language Processing (NLP), Question-Answering Systems, Speech Recognition, Speech Synthesis, Voice Interaction, Voiceprint Recognition, Speech Recognition Products
Computer Vision & Recognition	Computer Vision, Image Recognition, Facial Recognition, Biological Recognition, Human-Computer Interaction (duplicate removed), Human-Robot Collaboration
Big Data, Computing & Infrastructure	Big Data Analytics, Big Data Processing, Big Data Storage, Big Data Management, Big Data Platforms, Big Data Control, Business Intelligence, Distributed Computing, Edge Computing, Cloud Computing, Intelligent Computing, Automation
Applications & Smart Living Services	AI Products, AI Chips, Intelligent Chips, Intelligent Sensing, Intelligent Regulation, Intelligent Education, Intelligent Customer Service, Intelligent Healthcare, Intelligent Investment Advisory, Intelligent Insurance, Intelligent Services, Intelligent Search, Intelligent Wearables, Intelligent Elderly Care, Intelligent Environment, Smart Living, Smart Homes, Smart Speakers, Smart Agriculture, Smart Travel, Intelligent Transportation, Intelligent Networked Vehicles, Driverless Cars, Autonomous Driving, Robot Workflow, Augmented Reality, Virtual Reality, Augmented Intelligence

Note: The identification of AI-related keywords in this study is based on authoritative policy and academic sources to ensure enhance accuracy and coverage. Key references include the *Guidelines for Building a National New Generation Artificial Intelligence Standard System* (Standardization Administration of China and four other ministries), the *Three-Year Action Plan for Promoting the Development of a New Generation Artificial Intelligence Industry (2018–2020)* (Ministry of Industry and Information Technology), the *New Generation Artificial Intelligence Development Plan* (State Council, 2017), and the *Shenzhen Action Plan for the Development of a New Generation of Artificial Intelligence (2019–2023)*. Additional references include Stanford University's *Artificial Intelligence Index Report 2021* and Acemoglu et al. (2020a, 2020b). The intersection of these policies and academic sources informs the judgment used to define and classify AI-related job keywords.

Appendix Table A2 Marginal Effects of AI Labor Demand on Log Wages by Gender and Occupation

Dependent variable: log wage	AI labor demand		
	AI posts	AI firms	AI titles
(1) Clerical/Admin			
Marginal Effect of AI on Log Wages (Men, β_1)	-0.0014	-0.0091	-0.0003
	(0.0024)	(0.0176)	(0.0043)
Marginal Effect of AI on Log Wages (Women, $\beta_1 + \beta_3$)	0.0022	0.0116	0.0051
	(0.0022)	(0.0159)	(0.0039)
Gender Difference in Marginal Effect (Female – Male, β_3)	0.0036**	0.0207**	0.0054**
	(0.0018)	(0.0084)	(0.0027)
R2	0.249	0.249	0.249
N	1,225	1,225	1,225
(2) Sales/Hospitality			
Marginal Effect of AI on Log Wages (Men, β_1)	-0.0053	-0.0137	-0.0124
	(0.0069)	(0.0519)	(0.0116)
Marginal Effect of AI on Log Wages (Women, $\beta_1 + \beta_3$)	0.0065	0.0508	0.0083
	(0.0070)	(0.0571)	(0.0126)
Gender Difference in Marginal Effect (Female – Male, β_3)	0.0118**	0.0645**	0.0207***
	(0.0051)	(0.0308)	(0.0068)
R2	0.361	0.361	0.362
N	383	383	383
(3) Healthcare			
Marginal Effect of AI on Log Wages (Men, β_1)	0.0742	0.3618	0.2367
	(0.0734)	(0.3765)	(0.2620)
Marginal Effect of AI on Log Wages (Women, $\beta_1 + \beta_3$)	0.0811	0.3974	0.2448
	(0.0678)	(0.3460)	(0.2529)
Gender Difference in Marginal Effect (Female – Male, β_3)	0.0070	0.0356	0.0081
	(0.0145)	(0.0750)	(0.0191)
R2	0.775	0.775	0.773
N	75	75	75
(4) Education			
Marginal Effect of AI on Log Wages (Men, β_1)	0.0213***	0.1025**	0.0501***
	(0.0077)	(0.0470)	(0.0126)
Marginal Effect of AI on Log Wages (Women, $\beta_1 + \beta_3$)	0.0089	0.0393	0.0292**
	(0.0085)	(0.0526)	(0.0137)
Gender Difference in Marginal Effect (Female – Male, β_3)	-0.0124***	-0.0631***	-0.0209***
	(0.0025)	(0.0146)	(0.0043)
R2	0.336	0.334	0.341

N	356	356	356
(5) Finance/Professions			
Marginal Effect of AI on Log Wages	0.0022	0.0199	0.0043
(Men, β_1)	(0.0101)	(0.0742)	(0.0167)
Marginal Effect of AI on Log Wages	0.0107	0.0698	0.0201
(Women, $\beta_1 + \beta_3$)	(0.0072)	(0.0617)	(0.0133)
Gender Difference in Marginal Effect	0.0085*	0.0499*	0.0158*
(Female – Male, β_3)	(0.0049)	(0.0281)	(0.0079)
R2	0.317	0.315	0.317
N	212	212	212
(6) Production/Manufacturing			
Marginal Effect of AI on Log Wages	-0.0032	-0.0240	-0.0067
(Men, β_1)	(0.0023)	(0.0148)	(0.0041)
Marginal Effect of AI on Log Wages	0.0097***	0.0450***	0.0145***
(Women, $\beta_1 + \beta_3$)	(0.0029)	(0.0167)	(0.0046)
Gender Difference in Marginal Effect	0.0129***	0.0690***	0.0211***
(Female – Male, β_3)	(0.0018)	(0.0098)	(0.0033)
R2	0.221	0.222	0.221
N	2,072	2,072	2,072
(7) Public/Security/Other			
Marginal Effect of AI on Log Wages	0.0012	0.0091	0.0027
(Men, β_1)	(0.0016)	(0.0116)	(0.0033)
Marginal Effect of AI on Log Wages	0.0070***	0.0409***	0.0130***
(Women, $\beta_1 + \beta_3$)	(0.0020)	(0.0131)	(0.0037)
Gender Difference in Marginal Effect	0.0058***	0.0318***	0.0104***
(Female – Male, β_3)	(0.0011)	(0.0060)	(0.0018)
R2	0.280	0.280	0.280
N	2,479	2,479	2,479

Note: I labor demand variables (AI job postings, AI posting firms, and AI job titles) are measured at the city–year level in units of 1,000. "Marginal Effect of AI on Log Wages (Men, β_1)" reports the baseline marginal effect of AI labor demand for male workers (Female = 0). "Marginal Effect of AI on Log Wages (Women, $\beta_1 + \beta_3$)" reports the implied marginal effect for female workers. "Gender Difference in Marginal Effect (Female − Male, β_3)" reports the estimated interaction coefficient (Female × AI), capturing the percentage-point difference in wage responses to AI labor demand between women and men. All regressions include city and year fixed effects and the full set of individual-level controls. Standard errors (in parentheses) are clustered at the city level.

Appendix Table A3 Lagged Effects of AI Job Postings, Firms, and Titles on Wage Outcomes, 2016–2022

Dependent variable: log wage			
Lagged AI job postings (t-1)	0.0032		
	(0.0023)		
Lagged AI posting firms (t-1)		0.0140	
		(0.0218)	
Lagged AI job titles (t-1)			0.0040
			(0.0034)
Age	0.0627***	0.0627***	0.0627***
	(0.0088)	(0.0088)	(0.0088)
Age squared	-0.0008***	-0.0008***	-0.0008***
	(0.0001)	(0.0001)	(0.0001)
Years of schooling	0.0602***	0.0602***	0.0602***
	(0.0034)	(0.0034)	(0.0034)
Gender (1=male)	0.3990***	0.3990***	0.3990***
	(0.0332)	(0.0332)	(0.0332)
Married (1=yes)	0.0509	0.0509	0.0509
	(0.0319)	(0.0320)	(0.0319)
Agricultural hukou (1=yes)	-0.0461*	-0.0463*	-0.0462*
	(0.0243)	(0.0243)	(0.0243)
Workplace in same city (1=yes)	-0.1775***	-0.1776***	-0.1775***
	(0.0318)	(0.0318)	(0.0318)
Log hours wrked per week	0.1352***	0.1351***	0.1352***
	(0.0277)	(0.0277)	(0.0277)
Health status (ref. = Poor)			
Fair	0.0320	0.0321	0.0320
	(0.0386)	(0.0386)	(0.0386)
Good	0.1141***	0.1143***	0.1142***
	(0.0315)	(0.0315)	(0.0315)
Very good	0.1460***	0.1462***	0.1460***
	(0.0324)	(0.0324)	(0.0324)
Excellent	0.0893**	0.0896**	0.0893**
	(0.0414)	(0.0414)	(0.0414)
Internet use (1=yes)	0.1327***	0.1326***	0.1326***
	(0.0270)	(0.0271)	(0.0270)
City & year fixed effects	Yes	Yes	Yes
Region-year fixed effects	Yes	Yes	Yes
Province time trends	Yes	Yes	Yes
City-level controls	Yes	Yes	Yes
Within R^2	0.253	0.253	0.253
N	6149	6149	6149

Note: Lagged AI job postings, lagged AI posting firms, and lagged AI job titles are measured at the city–year level and expressed in units of 1,000 from the previous year. Wage income is measured annually, after tax, and is CPI-adjusted to 2022 RMB. Log wage income is the natural logarithm of this measure. "AI posting firms" refers to the number of distinct firms with AI-related postings, and "AI job titles" refers to the number of distinct AI-related job titles. Internet use denotes access via a mobile or computer. All regressions control for include industry dummies and a set of city-level characteristics, including GDP per capita, fiscal expenditure as a share of GDP, secondary and tertiary industry value added, industrial structure index, higher education student ratio, and post and telecommunications as a share of GDP. Standard errors are in parentheses. * p<0.10, ** p<0.05, *** p<0.01

Appendix Table A4 Lagged Effects of AI Job Postings, Posting Firms, and Job Tittles on Wage Outcomes: Interaction vs. Gender-Specific Models

Dependent variable: log wage	Pooled Sample + Interaction			Male Sample			Female Sample		
	AI posts	AI firms	AI titles	AI posts	AI firms	AI titles	AI posts	AI firms	AI titles
AI job postings (t-1)	0.0067***			-0.0002			0.0098*		
	(0.0030)			(0.003)			(0.005)		
Gender (1=male)	0.4317***	0.4390***	0.4315***						
	(0.0250)	(0.0253)	(0.0249)						
Gender*AI job postings (t-1)	-0.0060***								
	(0.0020)								
AI posting firms (t-1)		0.0327			-0.0146			0.0548	
		(0.023)			(0.0027)			(0.039)	
Gender*AI posting firms (t-1)		-0.0354***							
		(0.006)							
AI job titles (t-1)			0.0091***			-0.0011			0.0127*
			(0.004)			(0.005)			(0.007)
Gender* AI job titles (t-1)			-0.0087***						
			(0.0019)						
Within R2	0.256	0.256	0.256	0.188	0.188	0.188	0.225	0.225	0.225
N	6149	6149	6149	3356	3356	3356	2793	2793	2793

Note: Lagged AI job postings, AI posting firms, and AI job titles are measured at the city–year level and expressed in units of 1,000 from the previous year. Wage income is measured annually, after tax, and is CPI-adjusted to 2022 RMB. Log wage income is the natural logarithm of this measure. "AI posting firms" refers to the number of distinct firms with AI-related postings, and "AI job titles" refers to the number of distinct AI-related job titles. Internet use denotes access via a mobile device or computer. All regressions include individual-level controls (age, age squared, years of schooling, marital status, agricultural hukou, workplace in same city, log hours worked per week, health status, and internet use), as well as city and year fixed effects, region–year fixed effects, province time trends, and city-level controls (GDP per capita, fiscal expenditure as a share of GDP, secondary and tertiary industry value added, industrial structure index, higher education student ratio, and post and telecommunications as a share of GDP). Standard errors are in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Appendix Figure A1 Top 20 AI-related Keywords in Job Postings

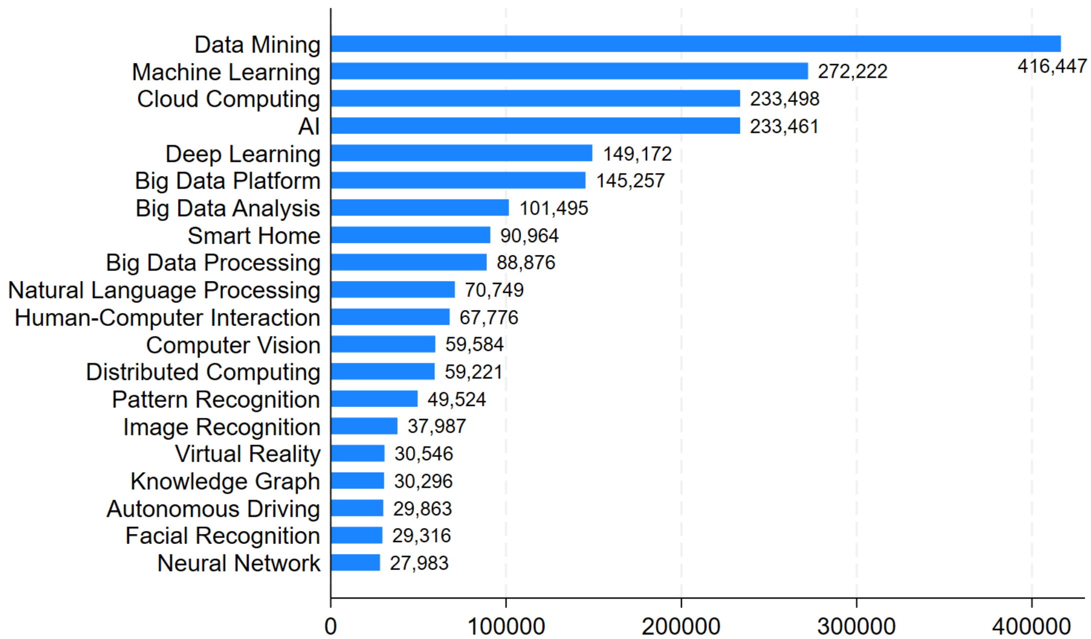

Source: Data collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS 直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58 同城), and Ganji.com (赶集网). AI-related postings are identified using a curated list of 73 keywords (e.g., artificial intelligence, deep learning, facial recognition, autonomous driving, neural networks).

Appendix Figure A2 Top 20 Cities with Most AI-Related Job Postings

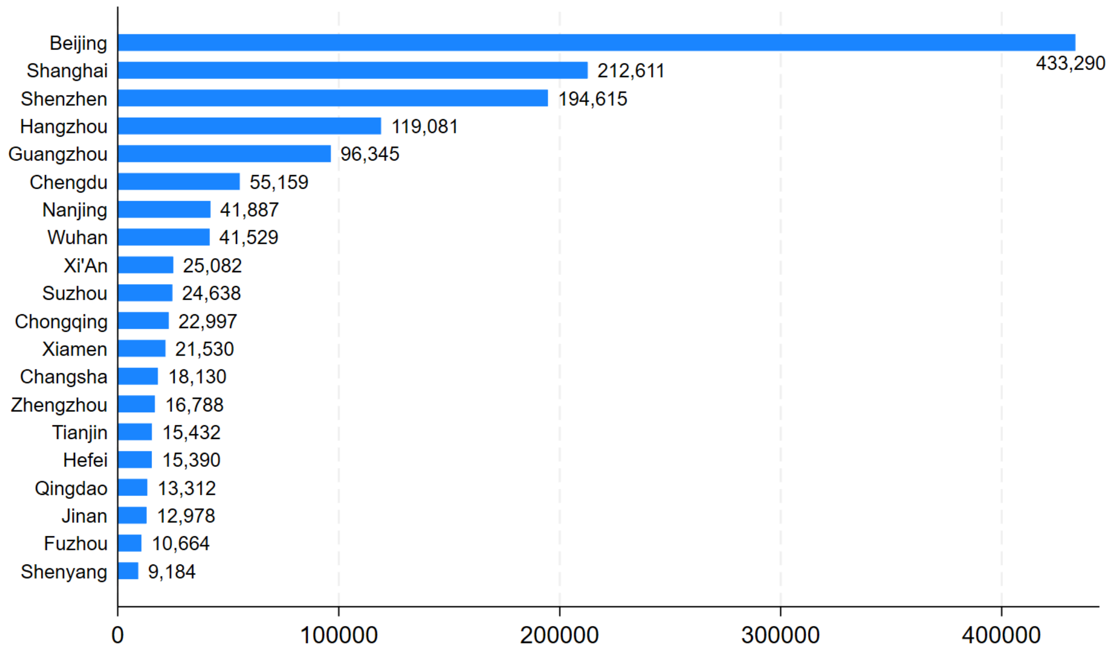

Source: Data collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58同城), and Ganji.com (赶集网). AI-related postings are identified using a curated list of 73 keywords (e.g., artificial intelligence, deep learning, facial recognition, autonomous driving, neural networks).

Appendix Figure A3 Top 10 AI-Related Job Positions

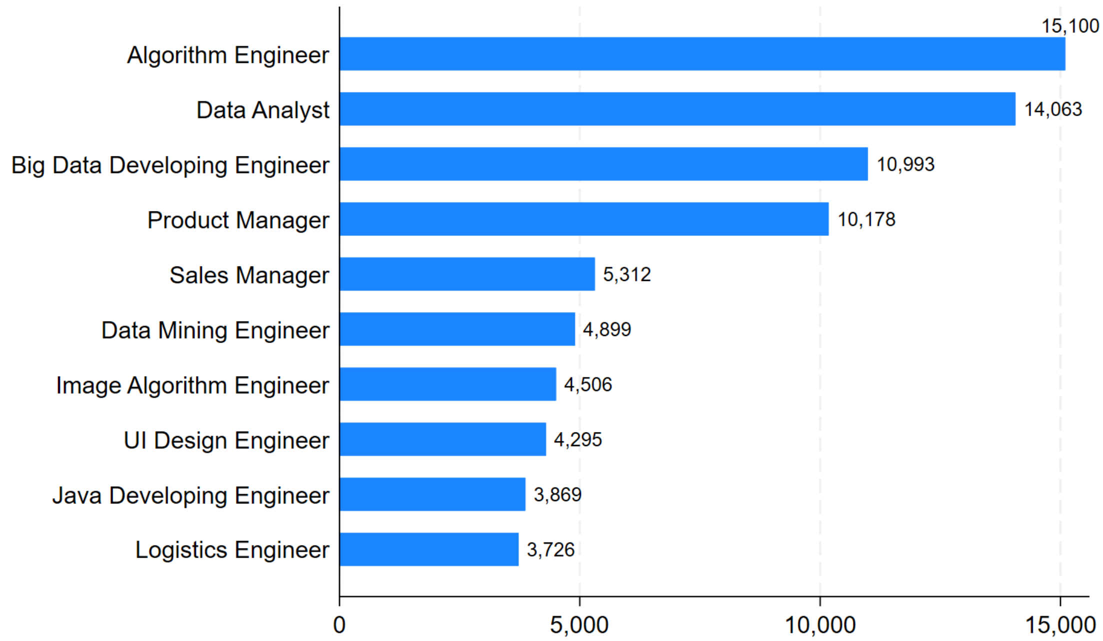

Source: Data collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58同城), and Ganji.com (赶集网). AI-related postings are identified using a curated list of 73 keywords (e.g., artificial intelligence, deep learning, facial recognition, autonomous driving, neural networks).

Appendix Figure A4 Education Requirement for AI-Related Jobs, 2024

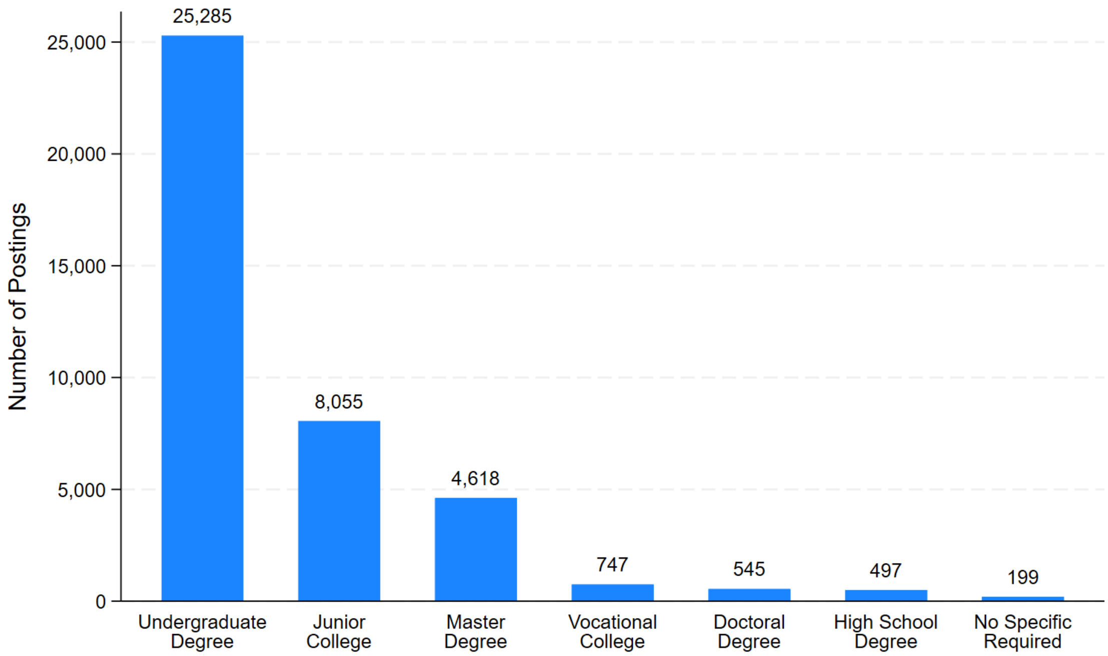

Source: Data collected from major online recruitment platforms in China, including 51job (前程无忧), BOSS Zhipin (BOSS直聘), Zhaopin (智联招聘), Liepin (猎聘网), Lagou (拉勾网), Kanzhun (看准网), 58.com (58同城), and Ganji.com (赶集网). AI-related postings are identified using a curated list of 73 keywords (e.g., artificial intelligence, deep learning, facial recognition, autonomous driving, neural networks).