# Generative AI and Pre-Hire Labor Market: Evidence from a Chinese Job Matching Platform*

Hanming Fang, Jiayin Hu, Miao Yu, Jiaqi Zhao

September 5, 2026

**Abstract**

This paper investigates the impact of Large Language Models on the pre-hire stage of the labor market. Using granular job posting and application records from a major online job platform in China, we document that employers adjust wages, time-related amenities, and callback rates simultaneously in response to the AI shock. These demand-side shocks transmit unevenly to job seekers: while junior applicants suffer substantial declines in callback rates, whereas senior applicants m face harsher time-related amenities. Meanwhile, the demand-side wage decline did not transmit to the applicant side. Furthermore, applicants exhibit limited occupational mobility to escape these shocks; instead, regional labor protections act as an institutional buffer against job-quality deterioration. These findings underscore the high elasticity of non-monetary job characteristics to technological shocks in developing countries, and highlight the role of institutional environments in mitigating labor market disruptions.

**Keywords**: Large Language Models; Non-monetary Amenities; Job Search and Matching; Labor Market Institutions
**JEL Codes**: O33, J23, J24, J88

*Preliminary and all comments welcome. Fang: Department of Economics, University of Pennsylvania (hanming.fang@econ.upenn.edu). Hu: National School of Development, Peking University (jyhu@nsd.pku.edu.cn). Yu: National School of Development, Peking University (myu2021@nsd.pku.edu.cn). Zhao: National School of Development, Peking University (jqzhao2017@nsd.pku.edu.cn).

# 1 Introduction

The proliferation of large language models (LLMs) raises significant policy concerns regarding their labor market implications. Functioning as a general-purpose technology, LLMs are increasingly widespread across industries (Eloundou et al., 2024) and demonstrate a distinct capacity to substitute for entry-level experience (Hosseini Maasoum and Lichtinger, 2025). The current debate centers on two critical questions: how technology-induced shifts in labor demand ultimately transmit to workers; and which specific groups are affected along which dimensions. To date, the literature has predominantly examined the firm-side response to AI integration, focusing on within-firm productivity gains (e.g., Brynjolfsson, Li and Raymond, 2025; Noy and Zhang, 2023; Dell' Acqua et al., 2026; Humlum and Vestergaard, 2025) and subsequent adjustments in labor demand (e.g., Eisfeldt et al., 2023; Brynjolfsson, Chandar and Chen, 2025; Hosseini Maasoum and Lichtinger, 2025; Liu, Wang and Yu, 2025).

Despite these empirical advances, less is known about the pre-hire stage of the labor market. At this juncture, the transmission of LLM shocks is shaped not only by the dynamic interactions between employers and job seekers, but also by the prevailing institutional environment. Furthermore, existing discussions primarily focus on wage and quantity adjustments, largely neglecting non-monetary dimensions of the employment contract. These non-monetary margins are especially critical in developing economies, where relatively weak regulations provide firms with greater elasticity to alter job characteristics (Breza and Kaur, 2025). Among these, time-related amenities represent a primary, yet critically underexplored, margin through which firms might shift the burden of LLM-induced shocks during the recruitment process.

To address this gap, this paper focuses on the online job market in China, where a substantial proportion of pre-hire matching currently occurs. We ask three primary questions: First, how do LLM shocks affect both the monetary and non-monetary dimensions of labor demand in the market? Second, how are these shocks transmitted to applicants? Third, can regional institutional environments mitigate the impact of LLM shocks?

China provides an ideal empirical setting for several reasons. First, it has experienced rapid evolution and deployment of LLMs since the release of ChatGPT (See Section 2.2). Meanwhile, due to China's vast diversity in industrial structure, there is substantial variation in occupational exposure to LLMs. Second, China is a developing economy characterized by prevalent overtime work, which offers a unique context to examine adjustments along non-monetary margins. According to its National Bureau of Statistics, average weekly working hours increased from 47.8 in 2021 to 48.6 in 2025, a level relatively higher than in other countries with similar income levels (Gethin and Saez, 2025). Finally, China's regionally

decentralized government structure creates significant spatial variation in institutional enforcement (Xu, 2011), including the stringency of labor regulations. This allows us to exploit regional variation to explore the mitigating role of regulation.

Specifically, we utilize daily-level records from a leading online recruitment platform in China, spanning from 2021 to 2024, which accounts for more than 30% of the country's online recruitment market share. We construct two analysis samples. The first is a job-posting sample, which randomly selects approximately 0.8 million vacancies alongside the universe of their corresponding applications. The second is an applicant-side sample, which tracks the complete application records of twenty thousand job seekers randomly drawn each year. For each sample, we focus on three primary outcome variables related to job quality: monthly salary; time-related amenities, defined as an indicator of whether the job allows for weekends off; and callback status, defined as an employer-initiated positive response within 24 hours of application submission. Crucially, time-related amenities and callbacks capture non-monetary margins of the employment contract and the matching process.

To identify the impact of LLM shocks, we employ a difference-in-differences framework that interacts predetermined LLM exposure with a post-ChatGPT indicator. We measure occupational exposure to LLMs using the index developed by Zhang et al. (2025), which follows the task-based approach of Felten, Raj and Seamans (2021) and Eloundou et al. (2024). The index quantifies the degree to which each occupation' s core tasks are exposed to LLM capabilities and is constructed using job postings from the same platform. Furthermore, we define applicants' exposure based on the occupation of their most recent job. The specification controls for year-by-month fixed effects, occupation fixed effects, and rich interactions of applicants' human capital characteristics.

Our first result shows that, on the demand side, high-LLM-exposure occupations experience a deterioration in both job quality and quantity following the shock. A one-standard-deviation increase in LLM exposure is associated with a 2.8 percent decline in posted wages, a 2.3 percentage-point decline in time-related amenities, and a 1.9 percentage-point decline in callback rates. The deterioration of highly exposed jobs is associated with both a contraction in their share of total job postings and a surge in requirements for education and work experience. This pattern suggests that the LLM shock increases employers' selectivity and shifts demand toward workers with stronger human capital signals. However, although jobs requiring more experience exhibit slightly smaller wage declines, they suffer from larger declines in time-related amenities.

Second, demand-side pressures are transmitted unevenly to job seekers. The demand-side wage decline did not transmit to the applicant side, likely due to the rigidity in applicants' wage expectations. In terms of non-monetary dimensions, junior and less-educated applicants

experience substantial declines in callback rates. Senior applicants and applicants from elite universities, by contrast, largely maintain their callback rates but face a deterioration in time-related amenities. In other words, they preserve matching prospects partly by shifting to jobs with longer or more rigid working hours. Therefore, the LLM shock does not simply affect some workers more than others; rather, it channels the burden through different margins depending on a worker's human capital.

Furthermore, we examine whether the impact of the LLM shock can be mitigated by individual adjustments and institutional environments. On the one hand, applicant adjustments are constrained by limited mobility: they do not systematically transition toward occupations with lower LLM exposure or switch to other occupations. This friction occurs either because applicants are reluctant to change their occupation due to job search costs or because they are unable to identify jobs with lower LLM exposure. On the other hand, we identify the mitigating effect of regional institutional environments. Using province-level labor dispute data to proxy for the strength of labor protection, we find that stronger protection prevents the deterioration of posted wages and time-related amenities. These findings suggest that, given limited applicant-side adjustments, pro-labor institutional environments can act as a buffer against job-quality deterioration under technological changes.

This paper contributes to several strands of the literature.

First and foremost, we contribute to the emerging literature on the labor market impacts of generative AI. Existing studies largely focus on within-firm productivity (e.g., Brynjolfsson, Li and Raymond, 2025; Noy and Zhang, 2023; Dell' Acqua et al., 2026; Humlum and Vestergaard, 2025) and subsequent labor demand adjustments. These adjustments occur along two margins: effects on incumbent workers, which tend to be relatively limited (Humlum and Vestergaard, 2025), and effects on job postings and new hires, where the literature highlights a contraction in quantity (Eisfeldt et al., 2023; Hui, Reshef and Zhou, 2024; Liu, Wang and Yu, 2025). However, the distributional effects across human capital remain mixed: while some research indicates that generative AI complements senior expertise (e.g., Auer, Köpfer and Švéda, 2024; Brynjolfsson, Chandar and Chen, 2025; Hosseini Maasoum and Lichtinger, 2025), other studies find that it disproportionately empowers junior workers (Brynjolfsson, Li and Raymond, 2025; Liu, Wang and Yu, 2025).

We advance this literature along three dimensions. First, we shift the focus to the pre-hire stage of the labor market, providing direct evidence on how AI-induced demand shocks transmit through search and matching frictions to affect applicant outcomes. Second, we highlight working-time amenities as a highly elastic, yet underexplored, adjustment margin in developing countries characterized by weak labor regulations. By incorporating this non-monetary dimension, we reconcile the mixed evidence on LLM's distributional effect: the

shock penalizes junior applicants through reduced wages and callback rates, while senior applicants bear the cost through deteriorating time-related amenities. Our focus on the time dimension echoes Jiang et al. (2025), which documents AI-induced changes in the workday using survey data. We depart from Jiang et al. (2025) by detailing interactions in the pre-hire stage, which allows us to identify exactly which applicants are affected on the time margin and how. Finally, we underscore the interaction between technological shocks and institutional environments, demonstrating that while applicants exhibit low mobility to escape LLM shocks, regional labor protections serve as a crucial institutional buffer.

Second, we contribute to the literature on the determinants of working time and non-monetary job amenities. Despite early optimism that technological progress would significantly reduce work hours (Keynes, 1930), aggregate market hours have not declined proportionately, particularly for high-wage workers (Costa, 2000; Aguiar and Hurst, 2007). While existing studies establish that technological change alters task composition and subsequent working hours (Autor, Levy and Murnane, 2003; Graetz and Michaels, 2018; Jiang et al., 2025), we provide novel evidence on how firms adjust along non-monetary margins during the pre-hire stage. We demonstrate that in a developing country context characterized by weak institutional regulation, the diffusion of large language models prompts firms to adversely shift time-related amenities. By highlighting the high elasticity of non-wage job characteristics, we show that firms utilize time demands as a primary margin of adjustment to technological shocks.

Third, our findings regarding regional heterogeneity speak to the literature evaluating the impact of labor market regulation. A substantial body of work explores how labor regulations shape unemployment (Botero et al., 2004), wages (Almeida and Carneiro, 2012), industrial performance (Besley and Burgess, 2004), and the transmission of macroeconomic and trade shocks (Ponczek and Ulyssea, 2022). Our micro-empirical evidence on the mitigating role of local labor protections echoes Acemoglu and Johnson (2023), who argue that the consequences of technological advancement depend on prevailing institutional frameworks. We show that when applicants exhibit low mobility to circumvent technological shocks, regional institutional protections act as a critical buffer dictating how demand-side shocks are ultimately transmitted to workers.

The remainder of the paper is organized as follows. Section 2 introduces the labor market and the evolution of large language models in China. Section 3 describes the online job platform data used in the analysis. Section 4 presents the empirical strategy used to identify the impact of LLM shocks on the pre-hire labor market. Section 5 utilizes job postings to document how large language models affect the demand side of the labor market. Section 6 explores how this demand-side shock is transmitted to the applicant side. Section 7 examines

the role of regional labor protections in mitigating LLM shocks. Section 8 concludes.

## 2 Institutional Background

### 2.1 China's Labor Market

Since the market oriented reforms in the mid-1990s, China's labor market has been characterized by informality (Park and Cai, 2011): employment relationships are governed not solely by legally binding contracts but also by a flexible set of job amenities that are less regulated. With the recent growth of digital labor platforms, informality has intensified as companies increasingly adopt crowdsourcing and franchised cooperation to legally circumvent formal labor relations (Zhou, 2020). According to a policy report in 2025, the number of flexible workers with no stable employment relationship reaches 300 million (China Research Center for New Employment Forms, 2026).

[Figure 1 about here]

One consequence of informality is long working hours. As shown in Figure 1, the average weekly working hours in China have exceeded 46 hours since 2017. This trend is reflected in the expansion of the so-called"996" schedule (9 a.m. to 9 p.m., six days a week) since the 2010s. Although China's Supreme People's Court and the Ministry of Human Resources and Social Security jointly declared the 996 schedule illegal in August 2021, [1], actual enforcement remains limited, with extended working hours persisting and even increasing to 48 hours in 2026. This trend is partly attributed to the recent waves of digitization, which blur the boundaries between work and personal time, enabling constant connectivity and remote task completion that effectively extend the workday beyond its formal limits.

As a continuous attempt to increase the formality of the labor market, China's labor regulatory framework has evolved substantially over the past three decades. The 1994 Labor Law established the foundational legal regime for contract-based employment, which codified basic protections covering labor contracts, working hours, minimum wages, and occupational safety. The 2008 Labor Contract Law significantly strengthened these protections by mandating written contracts for all employees, restricting the repeated use of fixed-term contracts, and increasing severance requirements (Gallagher et al., 2015). Despite this legislative progress, de-facto enforcement remains highly heterogeneous across regions. China's local governance is characterized by regional decentralization (Xu, 2011), where local officials under horizontal competition face strong career incentives to prioritize economic growth over

[1] See https://www.mohrss.gov.cn/SYrlzyhshbzb/laodongguanxi_/zcwj/diaojiezhongcai/202108/t20210825_421600.html

strict regulatory compliance. As a result, de facto labor protection varies considerably across regions (Gallagher et al., 2015).

### 2.2 The Evolution of Large Language Models in China

The diffusion of large language models in China has been extraordinarily rapid. By June 2025, the number of generative AI users in China had reached 515 million, doubling within six months.(China Internet Network Information Center, 2025). On the enterprise side, adoption takes two distinct forms. Technology-intensive firms integrate LLM capabilities directly into production pipelines through application programming interfaces (APIs), embedding models into workflows spanning automated content generation, code review, and customer service optimization. Meanwhile, a broader set of firms and individual workers access LLM capabilities through consumer-facing chat interfaces, using general-purpose platforms for tasks such as document drafting, translation, and data summarization. Both channels represent a substantive shift in how cognitive labor is organized within Chinese enterprises.

The technological shock can be traced to the release of ChatGPT by OpenAI on November 30, 2022. Although ChatGPT was never officially available in mainland China—due to both regulatory restrictions and platform-level access controls—the model's outstanding capabilities generated immediate and widespread attention among Chinese firms, technology developers, and policymakers.[2] The evolution of domestic LLM was remarkably swift: Baidu launched ERNIE Bot (Wenxin Yiyan) for invited testing on March 16, 2023, followed by Alibaba's Qwen on April 7, 2023, and ByteDance's Doubao on August 17, 2023. On August 31, 2023, the Cyberspace Administration of China approved the first batch of eight generative AI products—including ERNIE Bot and ByteDance's Doubao—for unrestricted public use. [3] The pace of China's LLM development intensified further in January 2025 with the release of DeepSeek-R1, an open-source reasoning model that attracted over 30 million daily active users within 20 days of launch. This concentrated timeline constitutes a sudden, large-scale technological shock to the Chinese market.

The rapid diffusion of LLM capabilities has altered firms' labor demand side through two primary channels. The first is direct production restructuring. As LLMs demonstrate increasing proficiency in routine cognitive tasks, firms reorganize their internal production processes by substituting or augmenting these tasks with model-based outputs. The second

[2] During the initial months following ChatGPT's release (approximately 2022Q4–2023Q1), a subset of Chinese users accessed the service through virtual private networks and third-party mirrors, while domestic technology firms began intensive development of competitive alternatives.

[3] By the end of 2023, 64 generative AI services had completed regulatory filing with the Cyberspace Administration of China; this number surged to 302 by December 2024. See https://www.cac.gov.cn/2025-01/08/c_1738034725920930.htm

channel operates through anticipatory adjustment. When managers perceive that certain tasks will soon become automatable,[4], they may preemptively adjust headcount, raise qualification requirements, or restructure job descriptions. This expectation channel implies that labor demand response may be observed on job postings even during the early stages of LLM diffusion.

# 3 Data

## 3.1 Exposure to Large Language Models

To measure occupational exposure to large language models, we use the AI Exposure Index constructed by Zhang et al. (2025). This index follows the task-based approach developed by Eloundou et al. (2024), which evaluates whether access to an LLM or an LLM-powered system can substantially reduce the time required to complete a task while maintaining the same level of quality. Zhang et al. (2025) adapts this framework to the Chinese labor market using online job-posting data. Their measure is constructed from job descriptions on the same online recruitment platform used in this paper, but it is based on a larger sample of approximately 1.25 million postings.

The construction proceeds in three steps. First, detailed job descriptions are matched to O*NET task descriptions using LLM-assisted classification. Second, each task is assigned an LLM exposure score by combining LLM-based evaluation with expert assessment. Third, task-level exposure scores are aggregated to the occupation level using task weights implied by the job-description to task mapping. Specifically, for each second-level occupation $J$, the task-based LLM exposure index is calculated as $S_J^d = \frac{1}{n_J} \sum_{j \in J} \sum_y s_y w_{jy}$, where $n_J$ denotes the number of job postings in occupation $J$, $s_y$ is the LLM exposure score of task $y$, and $w_{jy}$ is the weight of task $y$ in job posting $j$. A key advantage of this measure is that it is constructed from the actual task content of Chinese job postings rather than from a direct crosswalk between Chinese and U.S. occupational titles. This feature is particularly important in our setting because occupations with similar names may involve different task bundles across countries.

The exposure index is defined at the second-level occupation level in the platform's occupational classification.[5] Because occupational task content and skill requirements may

[4] For example, an executive at a listed technology company is quoted as stating: "It will take a few years, but future companies will be made of small units and a significant portion of the operational cost will be spent on AI." See The Wire China, "Balancing AI Pains and Gains," June 7, 2026, https://www.thewirechina.com/2026/06/07/balancing-ai-pains-and-gains/

[5] For example, within the first-level category 'Education/Training/Research," a second-level occupation may correspond to 'Early Childhood Education," under which more detailed job titles such as "Children's Education Teacher" are nested.

themselves respond to the LLM shock, we use a pre-shock exposure measure in our empirical analysis. Specifically, our baseline measure is the average LLM exposure of each second-level occupation over 2021 and 2022. We use this two-year average for two reasons. First, ChatGPT was released at the end of 2022, so occupational task content in 2022 was still largely predetermined with respect to the shock. Second, averaging over two years increases the number of postings used to construct the exposure measure and reduces the influence of extreme values in any single year. In the robustness analysis, we also show that the results are similar when exposure is measured using only 2021 data.

The pre-shock LLM exposure index varies considerably across second-level occupations. As shown in Table 1, the mean exposure value in our analysis sample is 0.589, with a standard deviation of 0.169. Occupations located near the sample mean include operations management, animation and motion design, and primary and secondary school curriculum tutoring, all of which have exposure values within 0.005 of the sample mean. A one-standard-deviation difference is economically sizable: it is roughly comparable to the gap between hotel services and compliance, risk control, legal affairs, and lawyers. Appendix Table A1 reports the LLM exposure index for all second-level occupations. Occupations with high LLM exposure include finance, editors and writers, customer service, and software development, while occupations with low exposure are concentrated among food services, general workers and technicians, domestic services and repair, and delivery and sorting.

### 3.2 Online Job Platform

Our empirical analysis is based on large-scale online recruitment data from one of the largest online job-search platforms in China, covering the period from 2021 to 2024. The data contain two linked data structures: job postings and application records. Job postings provide information on employer demand and job characteristics, while application records link applicants to specific vacancies and record employer responses. This data structure allows us to examine both sides of the pre-match labor market: how employers adjust job demand, job quality, and job requirements, and how applicants respond through their search and application choices.

#### 3.2.1 Job Posting Data

For the posting-level analysis, we draw a random sample of job advertisements from the platform. The sampling is stratified by year, with 200,000 postings randomly drawn in each year from 2021 to 2024. After data cleaning, the final posting-level sample contains 777,718

valid job postings, covering 183 industries and 93 second-level occupation categories.[6] For each posting, we observe the job title, posting location, industry and occupation categories, advertised wage range, required qualifications, and benefit descriptions. Required qualifications include education and work experience requirements. Benefit descriptions are recorded as text, from which we extract information on job amenities. Because job postings report wage ranges rather than point wages, we use the midpoint of the advertised wage range as our wage measure.

Panel B and Panel C of Table 1 present descriptive statistics from the job posting sample. On the employer side, we examine three dimensions of job posting outcomes: job quantity, job quality, and job requirements. Because the job posting sample is drawn with a fixed sampling size in each year, raw posting counts cannot be directly used to measure changes in occupational demand. We therefore construct occupation-year-level ratios to capture relative changes in job quantity across occupations. Job Posting Share is defined as the number of job postings in an occupation-year cell divided by the total number of sampled job postings in that year. Vacancy Share is defined analogously using the reported number of vacancies in each posting.[7]

For job quality, we focus on both monetary compensation and non-wage time-related amenities. The average advertised monthly wage is 9,137 RMB, which is markedly higher than the national average and reflects the platform's focus on white-collar job seekers. A key non-wage outcome in our analysis is the presence of time-related amenities. We define a dummy variable equal to one if the benefit description explicitly mentions core time-related benefits, such as a regular five day workweek, weekends off, or no overtime.[8] We also use stricter and broader definitions of time-related amenities and report the corresponding results in the robustness section. This measure should be interpreted as an indicator that a job explicitly advertises favorable time arrangements, rather than as a complete measure of whether the job actually provides such amenities. In particular, postings coded as zero do not necessarily lack time-related amenities. However, relative to postings that explicitly mention these benefits, they are less likely to offer clearly stated protections against overtime or irregular schedules. In the job posting sample, about 19% of postings explicitly mention

[6] Industries and occupations are defined according to the platform's internal classification system. For example, industries include categories such as petroleum, petrochemical, and chemical industries, while second-level occupations include categories such as e-commerce operations.

[7] Because reported vacancy counts may be imprecise or inflated in some postings, we construct Vacancy Share using only postings with reported vacancies no greater than 20 and treat Job Posting Share as the primary measure of occupational demand.

[8] The core keywords used to identify time-related amenities are: two days off on weekends (周末双休), fixed two day weekends (固定双休), fixed weekends off (周末固定双休), two day weekends (双休), five workdays and two days off (上五休二), five days on and two days off (做五休二), weekends off (周末休息), no overtime (不加班), no mandatory overtime (不强制加班), and rejection of overtime (拒绝加班).

time-related amenities.

Finally, we measure job requirements using education and work experience thresholds. Specifically, we construct indicators for whether a posting requires at least a bachelor's degree and whether it requires at least three years of work experience. In the full job posting sample, 20% of postings require at least a bachelor's degree, and 19% require at least three years of work experience, again underscoring the platform's white-collar orientation.

### 3.2.2 Application Record Data

We use application record data to study employer responses and applicants' search behavior at the pre-hire stage. Each observation in the application data corresponds to an application–job pair that links information about the applicant, the job, and the employer's response. For applicants, we observe the application date, gender, age, education, work experience, and city of residence. Applicants are also required to report desired job characteristics, including expected wage, preferred work location, and desired occupation. For jobs, we observe the same set of posting characteristics described above, including job title, location, industry, occupation, wage range, qualification requirements, and benefit descriptions. Employer responses are captured by two indicators: whether the employer sends any reply within 24 hours and whether the reply is positive. A positive response is defined as an interview invitation, phone call, contact exchange, or follow-up conversation about the job. Detailed examples of positive responses are provided in Appendix Figure A1.

Our main callback measure is an indicator of whether the employer gives a positive response within 24 hours after the application is submitted. This 24-hour window reflects the platform's standard reporting horizon and therefore corresponds to the response margin directly observable in our data. Although this measure captures immediate rather than all eventual responses, it provides a credible proxy for employer interest at the early stage of job matching.

We draw random sub-samples from the full database for empirical analysis. Specifically, we design different sampling strategies for the employer-side and applicant-side analyses, resulting in two separate application samples. For the employer-side analysis, we randomly select 20,000 job postings per year and include all application records received by these postings during the same year. This job-centered sampling strategy yields 2,635,033 application-level observations. Because multiple applicants apply to the same vacancy, this sample allows us to observe employer callback behavior while holding fixed the job being applied to. As shown in Panel D of Table 1, the average probability of receiving a positive reply within 24 hours is 22%, suggesting that employers on this platform are relatively responsive.

For the applicant-side analysis, we randomly select 20,000 job seekers per year and collect

their complete application histories during that year. The final analysis sample contains 58,833 applicants and 1,574,232 application-level observations [9]. It allows us to examine how job seekers adjust their search behavior in response to the LLM shock, including the types of jobs they apply to, the wage levels and time amenities of applied jobs, the number of applications submitted, and the duration of the search process. We also observe employer responses in this sample. Although the definition of a positive callback is the same as in the employer-side sample, the interpretation differs. In the employer-side sample, callback outcomes primarily capture employers' screening behavior across applicants and jobs. In the applicant-side sample, callback outcomes more directly reflect applicants' job-search difficulty and their attractiveness in the labor market.

In the baseline applicant-side analysis, we measure an applicant's LLM exposure using the exposure of the applicant's last occupation. This choice captures the extent to which the applicant's existing occupational background is exposed to LLM technologies. Accordingly, the baseline applicant-side regressions exclude individuals with no previous work experience. In the robustness analysis, we alternatively measure applicant exposure using the exposure of the applicant's desired occupation, which allows us to include applicants without prior work experience.

Table 2 reports summary statistics for the applicant-side sample. Relative to applicants' last jobs, 32% of applications are directed toward jobs with lower LLM exposure, while 23% are directed toward jobs with higher LLM exposure. Panel B further summarizes annual search behavior at the applicant level. On average, each applicant submits 26.93 applications per year and applies across 8.16 industries, indicating a broad search for job opportunities. Among applicants who receive at least one positive callback, the average time from the first application to the first positive callback is 3.11 days.

Taken together, the posting-level and application-level samples enable us to study both sides of the pre-match labor market: how employers adjust job demand, job characteristics, and hiring preferences, and how applicants respond through their search and application choices.

### 3.3 Supplementary Data

We supplement the online recruitment data with province-level information from provincial labor statistical yearbooks. These data provide measures of labor dispute resolution, including the number of cases accepted in the current period, the number of settled cases, the number of cases won by workers, the number of cases with partial worker and employer

[9] Our analysis sample excludes applicants with no application history, no previous work experience record, invalid occupation classification, or missing demographic characteristics.

wins or other outcomes, and the number of cases won by employers. Because the outbreak of COVID-19 after 2020 may have generated abnormal fluctuations in labor disputes and dispute resolution, we use pre-pandemic province-level measures from 2019.

We also use province-level employment counts from the 2020 Population Census to normalize labor dispute outcomes by the size of the working population. Based on these data, we construct two measures of regional labor protection. The first is the number of labor dispute cases per worker, defined as the number of cases accepted in 2019 divided by the total number of employed workers in the province. The second is the worker win rate, defined as the number of cases won by workers divided by the number of settled cases in 2019. Appendix Table A2 reports the summary statistics for these province-level measures.

These two measures capture complementary dimensions of the regional labor protection environment. A lower number of cases per worker indicates fewer observed labor disputes relative to the size of the workforce, which may reflect a less conflictual or more effectively protected labor environment. A higher worker win rate indicates that workers are more likely to prevail, conditional on disputes entering the formal resolution process. In the heterogeneity analysis, we therefore interpret provinces with lower cases per worker and higher worker win rates as having stronger regional labor protection.

# 4 Empirical Strategy

Our empirical strategy exploits the release of ChatGPT as a shock that increased the economic relevance of LLM-related task exposure. The core idea is to compare labor market outcomes before and after the shock across occupations with different predetermined levels of LLM exposure.

We mainly estimate three sets of specifications corresponding to the three main empirical margins studied in this paper. First, using posting-level data, we examine whether employers adjust job demand, job quality, and job requirements in more exposed occupations after the LLM shock. Second, using application records from the employer-sample, we study whether applications to more exposed jobs receive differentially lower callback rates. Third, using records from the applicant-sample, we examine whether job seekers whose previous occupations are more exposed adjust their search behavior toward different types of jobs.

## 4.1 Employer-Side: Changes in Job Postings

We first examine how the LLM shock affects job postings on the employer side. As discussed in Section 3.2, we consider three dimensions of job posting outcomes: job quantity, job quality, and job requirements.

In all regressions, let $g$ denote a second-level occupation category, and $LLMExposure_g$ denote the pre-shock LLM exposure index of occupation $g$. For a job posting $j$, $g(j)$ denotes the second-level occupation to which the job belongs, while $o(j)$ denotes a more detailed occupation-type cell defined by the interaction of industry, second-level occupation, and third-level occupation. The indicator $Post_t$ equals one for observations after the ChatGPT shock, which is defined as December 2022 onward.

We use occupation-year-level posting shares to capture relative changes in job quantity across occupations. We estimate

$$Y_{gy} = \alpha + \beta Post_y \times LLMExposure_g + \lambda_y + \mu_g + \varepsilon_{gy}, \tag{1}$$

where $Y_{gy}$ denotes the job posting share or vacancy share of occupation $g$ in the year $y$. Job posting share is defined as the number of postings in occupation $g$ divided by the total number of sampled postings in year $y$. Vacancy share is defined analogously using the reported number of vacancies in each posting. The year fixed effects $\lambda_y$ absorb aggregate shocks common to all occupations, including macroeconomic conditions and platform-wide changes in hiring activity. The occupation fixed effects $\mu_g$ absorb time-invariant differences across second-level occupations.

For job quality and job requirements, the unit of observation is an individual job posting. For job posting $j$ posted in calendar month $t$, we estimate

$$Y_{jt} = \alpha + \beta Post_t \times LLMExposure_{g(j)} + \tau_t + \eta_{c(j)} + \theta_{o(j)} + \varepsilon_{jt}. \tag{2}$$

The outcome $Y_{jt}$ denotes the characteristics of job posting $j$, including the logarithm of the advertised wage midpoint, the time-related amenity indicator, education requirements, and experience requirements. The term $LLMExposure_{g(j)}$ denotes the pre-shock LLM exposure index of the second-level occupation to which job $j$ belongs. The model includes calendar year-by-month fixed effects, $\tau_t$, which absorb aggregate shocks common to all occupations, including macroeconomic conditions, platform-wide demand fluctuations, and seasonality. We also include job-city fixed effects, $\eta_{c(j)}$, to control for persistent differences across local labor markets. In addition, $\theta_{o(j)}$ denotes detailed occupation-type fixed effects.

Noted that the occupation-type fixed effect $\theta_{o(j)}$ is defined at a more granular level than the second-level occupation.[10] This structure allows us to measure exposure at the occupation level while flexibly absorbing persistent differences across finer job types. Standard errors

[10] For example, a job may belong to the first-level industry category 'Internet/Telecommunications/Hardware," the second-level occupation category 'Artificial Intelligence,"and the third-level occupation category "Algorithm Engineer."

are clustered at the occupation-type level.

### 4.2 Employer-Side: Changes in Callbacks

We next examine whether the LLM shock affects employer responses to applications. Using the job-centered application sample, each observation corresponds to an application submitted by applicant $i$ to job $j$ in calendar month $t$. We estimate

$$Y_{ijt} = \alpha + \beta Post_t \times LLMExposure_{g(j)} + \tau_t + \eta_{c(j)} + \theta_{o(j)} + X_i'\gamma + W_{ijt}'\rho + \varepsilon_{ijt}. \quad (3)$$

The outcome $Y_{ijt}$ is an indicator equal to one if applicant $i$ receives a positive response from the employer within 24 hours after applying to job $j$. As described in Section 3.2, a positive response includes an interview invitation, phone call, contact exchange, or follow-up conversation about the job. The term $LLMExposure_{g(j)}$ denotes the pre-shock LLM exposure index of the second-level occupation associated with job $j$.

The model includes calendar year-by-month fixed effects, job-city fixed effects, and detailed occupation-type fixed effects. The occupation-type fixed effects are defined in the same way as in Equation (2). We further control for applicant characteristics, $X_i$, including interactions of gender, age group, education level, and work experience group.[11] Since callback outcomes are defined based on whether a positive response occurs within 24 hours, response rates may vary with the day of application; therefore, we include workday controls, $W_{ijt}$, indicating whether the application date or the following day is a workday. Standard errors are clustered at the occupation-type level.

### 4.3 Applicant-Side: Changes in Application Behavior

Finally, we examine how the demand-side LLM shock transmits to job seekers during the pre-hire stage. Using the application record from applicant-side sample, we estimate

$$Y_{ijt} = \alpha + \beta Post_t \times LLMExposure_{g_i^0} + X_i'\gamma + \tau_t + \mu_{g_i^0} + \varepsilon_{ijt}. \quad (4)$$

The outcome $Y_{ijt}$ captures the characteristics of the job to which applicant $i$ applies, including the posted monthly wage, the time-related amenity indicator, and the LLM exposure of the applied job. For an applicant $i$, $g_i^0$ denotes the applicant's most recent occupation before the

[11] Education is classified into six categories based on the applicant's most recent degree: (1) middle school or below; (2) vocational or technical secondary school; (3) high school; (4) junior college (associate degree); (5) bachelor's degree; and (6) master's degree. Work experience is grouped into seven categories: (1) no experience; (2) less than 1 year; (3) 1–3 years; (4) 3–5 years; (5) 5–10 years; (6) 10–20 years; and (7) more than 20 years.

current job search. The term $LLMExposure_{g_i^0}$ represents applicant $i$'s exposure to the LLM shock, defined as the LLM exposure of the applicant's most recent occupation $g_i^0$ prior to the current job search. In the robustness analysis, we use alternative definitions of applicant-level exposure, including the exposure of the applicant's desired occupation.

The specification includes three sets of controls. First, the vector $X_i$ contains applicant-level controls, including interactions of gender, age group, education level, and work experience group. Second, calendar year-by-month fixed effects, $\tau_t$. Third, fixed effects for the applicant's last occupation, $\mu_{g_i^0}$, absorb time-invariant heterogeneity across previous occupations, including the main effect of $LLMExposure_{g_i^0}$. Standard errors are clustered at the applicant's last-occupation level to account for potential error correlation among applicants originating from the same occupational background.

We also estimate effects on applicants' annual search behavior using applicant-year-level outcomes. For applicant $i$ observed in year $y$, we estimate

$$Y_{iy} = \alpha + \beta Post_y \times LLMExposure_{g_i^0} + X_i'\gamma + \lambda_y + \mu_{g_i^0} + \varepsilon_{iy}. \tag{5}$$

The outcome $Y_{iy}$ denotes annual search behavior of applicant $i$ in year $y$, including the number of applications submitted, the number of industries applied to, and the number of days from the first application to the first positive callback. The exposure measure $LLMExposure_{g_i^0}$, applicant controls $X_i$, and last-occupation fixed effects $\mu_{g_i^0}$ follow the same definitions as in Equation (4). The year fixed effects $\lambda_y$ absorb aggregate annual shocks. The annual post indicator $Post_y$ equals one for years after 2022. In this specification, $\beta$ captures the effect of the LLM shock on applicants' aggregate search intensity, search breadth, and search duration.

# 5 Demand Side Results

## 5.1 Changes in Job Postings: Market Share, Wage, Time amenities and Requirement

### 5.1.1 Baseline Results

We first examine how LLM exposure affects job postings on the demand side of the labor market. This analysis asks whether more exposed occupations experience differential changes in posting shares, posted wages, time-related amenities, and qualification requirements after the ChatGPT shock. We use two complementary levels of analysis. Panel A of Table 3 aggregates postings to the occupation-category $\times year level and examines changes in the relative share of job postings and$

[Table 3 about here]

The results in Table 3 show that occupations with higher LLM exposure experience a broad decline in labor demand and posted job quality after the shock. In Panel A, the coefficients are negative and statistically significant for both the job ratio and the position ratio, suggesting that more exposed occupations account for a smaller share of postings and advertised positions after the LLM shock. In Panel B, more exposed occupations also exhibit higher qualification requirements. A one-standard-deviation increase (0.17) in LLM exposure is associated with a 0.47 percentage point increase in the probability that a posting requires at least three years of experience and a 0.63 percentage point increase in the probability that it requires a bachelor's degree or above.

At the same time, posted wages and time-related amenities decline significantly. A one-standard-deviation increase in exposure is associated with approximately a 2.0 percent decrease in wages. It is also associated with a 1.7 percentage point decline in the probability that a job explicitly advertises time-related amenities, which represents a decline of about 6.5 percent relative to the pre-shock mean.

These patterns suggest that employers in more LLM-exposed occupations adjust along multiple margins. On the extensive margin, the relative demand for these occupations declines, as reflected in lower posting and position shares. On the intensive margin, among jobs that continue to be posted, employers demand stronger human-capital signals while offering lower wages and fewer explicitly stated time-related amenities. This combination is consistent with a deterioration in the opportunity set faced by workers in more exposed occupations: remaining vacancies become more selective, but not necessarily more attractive.

### 5.1.2 Dynamic Results

Figure 2 presents the corresponding dynamic estimates for the posting-level outcomes in Panel B of Table 3. The event time is defined at the quarter level, with 2022Q3 omitted as the reference period. The vertical dashed line marks the beginning of the post period, and the light-blue dashed lines indicate the launch periods of major Chinese LLM tools. The dynamic patterns are broadly consistent with the baseline results. Before the ChatGPT shock, the estimated coefficients fluctuate around zero and do not display a clear differential trend for log wages or qualification requirements. The time-related amenity estimates are slightly negative in the pre-period, but they are relatively stable and do not show a sharp pre-existing decline. After the shock, the coefficients move increasingly in the direction implied by Table 3. These dynamic patterns suggest that the adjustment was gradual rather than immediate, and that the effects strengthened as domestic LLM tools became more salient in China.

[Figure 2 about here]

### 5.1.3 Heterogeneity Results

We further examine whether the deterioration in posted job quality differs by job requirements. Figure 3 presents the dynamic heterogeneous effects of LLM exposure on demand-side wage and time-related amenities. Table A7 presents the corresponding regression estimates.

[Figure 3 about here]

As shown in Figure 3, wage declines are somewhat smaller among jobs with higher education or experience requirements. For jobs requiring at least three years of experience, a one-standard-deviation increase in LLM exposure is associated with a wage decline of approximately 1.52 percent, compared with about 1.99 percent for jobs requiring less than three years of experience. Similarly, the implied wage decline is about 1.35 percent among jobs requiring a bachelor's degree or above, compared with about 1.96 percent among jobs with lower education requirements.

The pattern for time-related amenities differs from the wage results. The decline in time-related amenities is substantially larger among jobs with higher experience requirements. For jobs requiring at least three years of experience, a one-standard-deviation increase in LLM exposure is associated with a 3.45 percentage point decline in the probability of explicitly advertised time-related amenities. For jobs requiring less than three years of experience, the corresponding decline is 1.27 percentage points. Panel (b) of Figure 3 confirms this divergence: after the shock, the time-amenity effect becomes markedly more negative for high-experience jobs, especially after the launch periods of major Chinese LLM tools. By contrast, heterogeneity by education requirement is much less pronounced. The decline is about 1.49 percentage points among jobs requiring a bachelor's degree or above and about 1.39 percentage points among jobs with lower education requirements, and Panel (d) of Figure 3 shows broadly similar post-shock dynamics across the two education groups.

Taken together, the heterogeneity analysis suggests that human-capital requirements do not uniformly protect job quality across all margins. Higher-requirement jobs experience somewhat smaller wage declines. However, jobs requiring more experience face a sharper deterioration in time-related amenities. This distinction is important because it suggests that the burden of the LLM shock may be transmitted through different job-quality margins for different segments of the labor market. For lower-requirement jobs, the adjustment is more visible in wages. For higher-experience jobs, the adjustment appears more strongly through non-wage amenities, especially time-related working conditions.

### 5.1.4 Robustness Checks

We conduct a series of robustness checks to assess whether the results on job characteristics are sensitive to alternative specifications and variable definitions. Appendix Table A3 shows that the main findings are robust to alternative control structures. In Panel A, we add city-by-time fixed effects to absorb local labor market shocks that vary over time. All results remain robust, and the magnitude changes little. In Panel B, we further include industry-specific time trends. The negative effects on posted wages and time-related amenities remain statistically significant, and the effect on education requirements remains positive and significant. The coefficient on experience requirements becomes smaller and statistically insignificant, suggesting that the evidence on rising experience requirements is somewhat more sensitive to controlling for industry-specific trends.

Appendix Table A4 addresses the concern that treating each posting equally may underweight job advertisements that correspond to multiple vacancies. We therefore re-estimate the posting-level regressions using frequency weights based on the reported number of vacancies, excluding postings with more than 20 reported vacancies to reduce the influence of extreme or potentially inflated vacancy counts. The results remain qualitatively consistent with the baseline. Although the magnitudes are smaller than in the unweighted baseline specification, the signs and statistical significance of all four outcomes are preserved.

We also examine the robustness of the results to alternative ways of measuring LLM exposure and time-related amenities. Appendix Table A5 replaces the continuous exposure index with an indicator of whether an occupation's LLM exposure is above the median. The estimates show the same pattern as the baseline results. Finally, Appendix Table A6 shows that the deterioration in time-related amenities is not driven by the specific keyword definition used in the baseline analysis. The estimated effect remains negative and statistically significant when we use both a stricter definition, based only on explicit two-day weekend and no-overtime phrases, and a broader definition that also includes regular working hours and eight-hour work schedules. Taken together, these robustness checks support the interpretation that the LLM shock reduced job quality and increased hiring selectivity in more exposed occupations.

## 5.2 Changes in Employers' Callback

We next examine whether the demand-side adjustment documented above is also reflected in employers' callback decisions. While the previous subsection focuses on posted vacancies and job characteristics, this subsection turns to employer responses at the application stage. The estimates are based on the employer-side application sample.

Panel (a) of Figure 4 reports the dynamic effects of LLM exposure on callback rates. Before the shock, there is no clear pre-existing downward trend. After the shock, the estimates become more negative, especially around and after the launch periods of major Chinese LLM tools. The post-period pattern points to a decline in employers' willingness to provide positive callbacks for applications to more LLM-exposed occupations.

[Figure 4 about here]

Panel (b) of Figure 4 further shows that the callback penalty is heterogeneous by applicant human capital. Applicants with lower experience or lower education face larger negative effects, while stronger human-capital signals partly mitigate the decline. For applicants with no work experience, the estimated coefficient is -0.131; for those with 0–5 years of experience, it is -0.148. By contrast, the coefficient is smaller in magnitude for applicants with 5–10 years of experience (-0.097) and further attenuates for those with at least 10 years of experience (-0.057).

A similar gradient appears with education. The estimated effect is -0.130 for applicants with less than a bachelor's degree, -0.082 for applicants with a bachelor's degree or above, and only -0.028 and statistically insignificant, for applicants with both a bachelor's degree and an elite-university background.

These results suggest that the LLM shock affects not only the composition and quality of posted vacancies but also employers' screening behavior at the application stage. The overall decline in callback rates indicates that applicants to more exposed occupations face weaker employer demand after the shock. At the same time, the heterogeneity patterns show that stronger human-capital signals partially buffer applicants from this decline. Experience, higher education, and elite-university background appear to reduce the callback penalty, consistent with employers placing greater value on observable signals of applicant quality when uncertainty or adjustment pressure rises in more exposed occupations. However, the mitigation is only partial: except for the elite-university group, the estimates remain negative across all applicant groups, indicating that human capital attenuates but does not fully eliminate the adverse callback effect associated with LLM exposure.

## 6 Applicant Side Results

Having established that the LLM shock induces significant firm-side adjustments in both monetary and non-monetary job characteristics, we now turn to the applicant side to examine the extent to which these demand-side shocks transmit to job seekers. We first show that the monetary shock does not transmit to the applicant side, likely due to the rigidity in

applicants' wage expectations. We then document significant heterogeneity in how the burden along non-monetary margins is distributed across applicant profiles: junior and less-educated applicants primarily bear the shock through a substantial decline in callback rates, while senior applicants and those from elite universities, instead, face a pronounced deterioration in time-related amenities. Finally, we provide suggestive evidence explaining why job seekers cannot fully offset these adverse demand shifts: they lack sufficient mobility to reallocate their search efforts toward less LLM-exposed positions.

### 6.1 Transmission from Demand to Supply

We begin by examining the overall effect of LLM shocks on the characteristics of jobs targeted by job seekers. As shown in the first column of Table 4 and visualized in Figure 5, applicants with higher exposure to LLMs experience a significant negative impact on the time-related amenities of the positions they apply for, although the magnitude is smaller than in Table 3. Meanwhile, we observe a null overall effect on the wage and callback rate of the applied jobs. Taken together, the evidence suggests that on average, job seekers partially mitigate the shock from the demand-side, but they still applied for lower-paying jobs with worse time-related amenities.

[Table 4 about here]

[Figures 5 and 6 about here]

This aggregate effect, however, masks substantial heterogeneity across different applicant profiles, particularly along non-monetary dimensions. The first heterogeneity originates from seniority, as shown in columns (2) and (3) of Table 4. For junior applicants, the LLM shock translates into a greater negative impact on callback rates. By contrast, senior applicants even see a slightly higher callback rate. Yet, this apparent resilience comes at a cost: senior applicants bear additional burdens along the non-monetary dimension, applying for jobs with substantially lower time-related amenities. Therefore, although generative AI might appear to be a seniority-biased technology that favors more experienced workers (Brynjolfsson, Chandar and Chen, 2025; Hosseini Maasoum and Lichtinger, 2025), senior applicants are not pure winners once we account for the hidden, non-monetary dimensions of job applications. Similar heterogeneity emerges when we examine applicants by their educational attainment (columns (4) and (5)). Applicants from elite universities apply for positions that demand longer or more rigid working hours. Conversely, low-education applicants absorb the shock primarily through greater declines in applied wages, although this effect is not statistically

significant. The dynamic pattern of the LLM shock is shown in Figure 6, which suggests a growing impact as original LLM models in China become accessible.

It is noteworthy that the null effect on applicant-side wages is likely driven by the rigidity of applicants' wage expectations. First, job seekers exhibit sticky reservation wages that do not adjust to demand-side technological shocks. As demonstrated in Appendix Table A10, the LLM shock does not induce a statistically significant change in applicants' self-reported expected wages (Panel B). Consistently, we observe no significant shift in the propensity of applicants to target positions that offer wages below their stated expectations (Panel A). Second, we consistently observe a downward wage rigidity over the course of individuals' careers: job seekers are highly averse to accepting new positions that pay less than their previous employment. Appendix Table A11 provides corroborating evidence for this mechanism, documenting a monotonic increase in wages across successive work experiences within the same applicant. Because job seekers maintain their reservation wages, the transmission of the shock along the monetary margin is not obvious; instead, the cost on the applicant-side is largely borne through non-monetary dimensions.

We conduct several robustness checks regarding variable definitions and specifications. First, we show that our results are robust to alternative definitions of time-related amenities, as presented in Appendix Figure A2. Second, we check the robustness of our LLM exposures. The results hold in Appendix Figure A3 when replacing the continuous exposure index with an indicator for whether an occupation's LLM exposure is above the median.We also employ an alternative definition of the LLM shock based on the exposure of applicants' expected occupations rather than their last-held occupations, yielding consistent results (Appendix Figure A4).

Finally, we address potential selection issues concerning whether the applicant-side results are driven by LLM-induced layoffs. If the LLM shock to the last occupation leads to more layoffs, these applicants might be systematically different. However, Appendix Figure A5 demonstrates no systematic difference in the distribution of LLM exposure among job seekers across years. This aligns with recent literature suggesting that there is no significant layoff effect, at least during the early stages of the LLM shock (Humlum and Vestergaard, 2025).

We further verify our heterogeneity results using alternative grouping criteria. Decomposing the sample by the interaction of experience and education level (Appendix Table A9) confirms that the deterioration in time-related amenities is concentrated among senior and highly educated applicants. Additionally, the decline in callback rates is largely driven by junior applicants, particularly those who graduated from elite universities. We also test alternative grouping criteria in Appendix Figure A6. Grouping experience by five-year intervals instead of three, or using age groups instead of work experience, produces similar heterogene-

ity in terms of seniority. We also use a bachelor's degree instead of elite-university status as the grouping variable for education: while the inter-group gap in time-related amenities narrows, the gap in applied wages persists.

In summary, although applicants experience relatively smaller shocks compared to the magnitude of the demand-side shock, they are still unable to fully mitigate its adverse effects. Different groups of job seekers bear the costs of this technological shift along different dimensions—either through reduced wages and matching prospects or through deteriorated working conditions. Consequently, there is no clear winner among applicants in the face of the LLM shock.

### 6.2 Lack of Mobility

The preceding analysis demonstrates that applicants are unable to fully offset the adverse effects of the demand-side LLM shock. Table 5 provides suggestive evidence explaining this phenomenon: job seekers lack sufficient cross-occupational mobility to circumvent the LLM shock. The average LLM exposure of applied jobs slightly declined after the shock (Column (1)). However, this is primarily driven by applicants giving up jobs with higher LLM exposure (Column (2)), instead of switching to jobs with lower LLM exposure (Column (3)).

The limited mobility probably stems from a combination of two underlying mechanisms. First, market frictions may prevent applicants from swiftly transitioning across occupational boundaries. As shown in column (4) of Table 5, the propensity to apply for occupations distinct from their previous employment remains unchanged after the shock. Second, even if applicants can switch across occupations, they may lack the necessary information to identify and target roles that are relatively insulated from the LLM shock.

[Table 5 about here]

Furthermore, Table 6 examines whether applicants compensate for the shock by intensifying their overall job search efforts. Job seekers decreased their search intensity, either in terms of the total annual number of applications submitted or the breadth of industries targeted (Panels (A) and (B)); this is primarily driven by senior and lower-educated applicants. Additionally, Panel C reveals a decline in the number of days since receiving positive feedback from employers in these groups with diminished search intensity.

[Table 6 about here]

## 7 The Role of Regional Labor Protection

Previous sections have revealed that the LLM shock induces significant adjustments on the demand side along both monetary and non-monetary margins. These shocks are subsequently transmitted to applicants, with different demographic groups bearing the costs along distinct dimensions. Furthermore, as Section 6.2 underscores, individual responses, such as cross-occupational mobility, are insufficient to offset the adverse effects of the technological shock. Given the limited capacity of individual applicants to evade these shocks, a natural follow-up question is whether institutional safeguards can act as an effective buffer. This section explores this question by examining regional heterogeneity on labor protection in applicant-side analysis

Specifically, we utilize provincial-level measures of labor protection to investigate whether the impact of the LLM shock is attenuated in regions with stronger institutional safeguards. Data on regional labor protection are sourced from the *China Labour Statistical Yearbook*. We measure the strength of labor protection by the the share of labor dispute cases won by workers, reflecting the degree of institutional support for workers conditional on a dispute occurring. We then divide provinces equally based on the median value of 2019.[12] we divide the sample based on the applicants' locations.

[Table 7 about here]

Table 7 presents the heterogeneity results. The adverse impacts of the LLM shock on both wages and time-related amenities are smaller in regions with stronger institutional safeguards, while the heterogeneity in wage is particularly pronounced. Therefore, regional labor protection seems more effective in buffering shocks along the monetary dimension (columns (1) to (2)), compared to the time-amenity dimension (columns (3) to (4)). This discrepancy likely arises because wages are highly observable and directly governed by labor contracts, whereas time-related amenities are often implicit, less verifiable, and consequently harder to regulate.

Taken together, these findings highlight the critical role of regional labor protection under technological shocks, particularly when individual mobility is constrained by market frictions. Besides, while pro-labor institutional environments can effectively safeguard workers against the deterioration of observable job quality, such as wages, their ability to regulate hidden dimensions inherently faces more challenges.

[12] We use the median value of 2019 to isolate the interruption of COVID-19 Pandemic.

# 8 Conclusion

This paper investigates the impact of Large Language Models on the pre-hire stage of the labor market. Leveraging job postings and application records from a major online job platform in China, we document several key findings. First, on the labor demand side, we observe a significant deterioration in both job wages and time-related amenities. Second, the transmission of this demand-side shock to applicants is highly heterogeneous: junior and less-educated applicants face steep declines in callback rates, whereas senior and more-educated applicants bear the cost primarily through worsened time-related amenities. The demand-side shock is not transmitted to the applicant-side wage expectations yet, likely due to the rigidity in applicants' wage expectations. Finally, individual responses are insufficient to offset the LLM shock because of limited mobility. Instead, regional labor protection institutions act as a crucial buffer in mitigating these negative impacts.

Our findings yield several important policy implications. First and foremost, we highlight non-monetary margins—specifically time-related amenities—during the transmission of technological shocks, which is highly elastic in the context of developing countries. Second, because applicant response itself is insufficient to circumvent these shocks, government action is essential. Moreover, the heterogeneous nature of the shock's transmission implies that different applicant groups bear the costs along different dimensions. Therefore, policy interventions tailored to specific demographic profiles will be more effective than uniform labor market regulations.

These discussions also shed light on several avenues for future research. A primary extension involves disentangling the mechanisms that govern the substitution between different adjustment margins. Specifically, it remains an open question why the LLM shock manifests on monetary margins for some occupations and applicants, while on time-related margins for others. Additionally, our empirical analysis assumes that the underlying matching efficiency of the pre-hire stage remains constant. As online job platforms increasingly integrate LLMs to facilitate screening and matching, such as algorithmic recommendation, future work should explore how this adoption interacts with labor demand shocks. Finally, while we emphasize the mitigating benefits of regional labor protections, regulation inevitably involves trade-offs between firms and workers. Given sub-national variations in institutional stringency, firms may strategically migrate vacancies to regions with weaker enforcement, which prompts a potential "race to the bottom." Searching for empirical evidence on this spatial equilibrium response remains a critical step for evaluating the welfare implications of AI-induced disruptions on the labor market and local governance.

## References


**Acemoglu, Daron, and Simon Johnson.** 2023. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity.* New York: PublicAffairs.

**Aguiar, Mark, and Erik Hurst.** 2007. "Measuring trends in leisure: the allocation of time over five decades." *The quarterly journal of economics*, 122(3): 969–1006.

**Almeida, Rita, and Pedro Carneiro.** 2012. "Enforcement of labor regulation and informality." *American Economic Journal: Applied Economics*, 4(3): 64–89.

**Auer, Raphael, David Köpfer, and Josef Švéda.** 2024. "The rise of generative AI: Modelling exposure, substitution, and inequality effects on the US labour market." CESifo Working Paper.

**Autor, David H, Frank Levy, and Richard J Murnane.** 2003. "The skill content of recent technological change: An empirical exploration." *The Quarterly journal of economics*, 118(4): 1279–1333.

**Besley, Timothy, and Robin Burgess.** 2004. "Can labor regulation hinder economic performance? Evidence from India." *The Quarterly journal of economics*, 119(1): 91–134.

**Botero, Juan C, Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer.** 2004. "The regulation of labor." *The Quarterly Journal of Economics*, 119(4): 1339–1382.

**Breza, Emily, and Supreet Kaur.** 2025. "Labor markets in developing countries." *Annual Review of Economics*, 17.

**Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen.** 2025. "Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence."

**Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond.** 2025. "Generative AI at work." *The Quarterly Journal of Economics*, 140(2): 889–942.

**China Internet Network Information Center.** 2025. "Generative Artificial Intelligence Application Development Report (2025)." China Internet Network Information Center Research Report.

**China Research Center for New Employment Forms.** 2026. "2025 Chinese Blue-Collar Employment Research Report: Transitioning from Survival-Oriented to Development-Oriented Employment." China Research Center for New Employment Forms Research Report.

**Costa, Dora L.** 2000. "The Wage and the Length of the Work Day: From the 1890s to 1991." *Journal of Labor Economics*, 18(1): 156–181.

**Dell' Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz, Katherine C Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R Lakhani.** 2026. "Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality." *Organization Science*, 37(2): 403–423.

**Eisfeldt, Andrea L., Gregor Schubert, Miao Ben Zhang, and Bledi Taska.** 2023. "Generative AI and Firm Values."

**Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock.** 2024. "GPTs are GPTs: Labor market impact potential of LLMs." *Science*, 384(6702): 1306–1308.

**Felten, Edward, Manav Raj, and Robert Seamans.** 2021. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." *Strategic management journal*, 42(12): 2195–2217.

**Gallagher, Mary, John Giles, Albert Park, and Meiyan Wang.** 2015. "China's 2008 Labor Contract Law: Implementation and Implications for China's Workers." *Human Relations*, 68(2): 197–235.

**Gethin, Amory, and Emmanuel Saez.** 2025. "Global working hours." National Bureau of Economic Research.

**Graetz, Georg, and Guy Michaels.** 2018. "Robots at Work." *The Review of Economics and Statistics*, 100(5): 753–768.

**Hosseini Maasoum, Seyed Mahdi, and Guy Lichtinger.** 2025. "Generative AI as seniority-biased technological change: Evidence from US résumé and job posting data." *Available at SSRN 5425555*.

**Hui, Xiang, Oren Reshef, and Luofeng Zhou.** 2024. "The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market." *Organization Science*, 35(6): 1977–1989.

**Humlum, Anders, and Emilie Vestergaard.** 2025. *Large language models, small labor market effects.* Vol. 33777, National Bureau of Economic Research Cambridge, MA.

**Jiang, Wei, Junyoung Park, Jiqiu Xiao, and Shen Zhang.** 2025. *AI and the extended workday: Productivity, contracting efficiency, and distribution of rents.* National Bureau of Economic Research.

**Keynes, John Maynard.** 1930. "Economic Possibilities for Our Grandchildren." In *Essays in Persuasion*. 321–332. New York: W.W. Norton.

**Liu, Yan, He Wang, and Shu Yu.** 2025. "Labor Demand in the Shadow of Generative AI: Evidence from the U.S. Job Posting Data."

**Noy, Shakked, and Whitney Zhang.** 2023. "Experimental evidence on the productivity effects of generative artificial intelligence." *Science*, 381(6654): 187–192.

**Park, Albert, and Fang Cai.** 2011. "The Informalization of the Chinese Labor Market." In *From Iron Rice Bowl to Informalization: Markets, State and Workers in a Changing China.* , ed. Sarosh Kuruvilla, Mary Gallagher and Ching-kwan Lee, 17–35. Ithaca, NY: Cornell University Press.

**Ponczek, Vladimir, and Gabriel Ulyssea.** 2022. "Enforcement of labour regulation and the labour market effects of trade: Evidence from Brazil." *The Economic Journal*, 132(641): 361–390.

**Xu, Chenggang.** 2011. "The fundamental institutions of China's reforms and development." *Journal of economic literature*, 49(4): 1076–1151.

**Zhang, Dandan, Hang Yu, Lixing Li, Jiayin Hu, Yiqing Mo, and Hongbo Li.** 2025. "The Measurement of AI Exposure and Its Impact on Labor Demand in China: Evidence from Large Language Models." *Management World*, , (7): 59–72.

**Zhou, Irene.** 2020. *Digital labour platforms and labour protection in China.* ILO Working Paper.

# Tables

Table 1. Summary Statistics: LLM Exposure Index, Job Postings and Callback Rates

| Variables | Obs | Mean | Std. |
|---|---|---|---|
| **Panel A: LLM Exposure Index** | | | |
| LLM Exposure Index | 93 | 0.589 | 0.169 |
| **Panel B: Occupation-Year Posting Shares** | | | |
| Job Posting Share (%) | 5,192 | 0.077 | 0.376 |
| Vacancy Share (%) | 5,188 | 0.077 | 0.433 |
| **Panel C: Job Posting Characteristics** | | | |
| Wage | 777,718 | 9,137.234 | 4,862.374 |
| Time-related Amenity (Yes = 1) | 777,718 | 0.191 | 0.393 |
| Education Requirement $\geq$ Bachelor's Degree | 777,718 | 0.203 | 0.402 |
| Experience Requirement $\geq$ 3 Years | 777,718 | 0.191 | 0.393 |
| **Panel D: Callback Outcomes from Employer-side Sample** | | | |
| Positive Reply (Yes = 1) | 2,635,033 | 0.222 | 0.416 |

*Notes:* This table reports summary statistics for key variables used in the employer-side analysis. Panel A summarizes the pre-shock LLM exposure index at the second-level occupation level. Panel B reports occupation-year-level posting shares. Job Posting Share is defined as the number of job postings in an occupation-year cell divided by the total number of sampled job postings in that year. Vacancy Share is defined analogously using the reported number of vacancies in each posting. Because reported vacancy counts may be imprecise or inflated in some postings, Vacancy Share is constructed using only postings with reported vacancies no greater than 20. Panel C reports job posting-level characteristics. Posted wage is measured by the midpoint of the advertised wage range. The time-related amenity indicator equals one if the posting explicitly mentions core time-related benefits, such as weekends off or no overtime. Panel D reports application-level callback outcomes, where Positive Callback within 24 Hours equals one if the employer gives a positive response within 24 hours.

Table 2. Summary Statistics of Outcome Variables: Applicant Side

| Variable | Obs | Mean | Std. |
|---|---|---|---|
| **Panel A: Characteristics of Jobs Applied** | | | |
| Wage | 1,574,232 | 9,570.031 | 14,377.790 |
| Time-related amenity | 1,574,232 | 0.232 | 0.422 |
| Positive reply (Callback) | 1,573,996 | 0.284 | 0.451 |
| Occupation Class: last job $\neq$ applied | 1,574,232 | 0.559 | 0.496 |
| LLM exposure of applied job | 1,538,684 | 0.579 | 0.200 |
| LLM exposure of applied job $<$ last job | 1,538,684 | 0.324 | 0.468 |
| LLM exposure of applied job $>$ last job | 1,538,684 | 0.225 | 0.417 |
| **Panel B: Applicant Annual Search Behavior** | | | |
| # of Applications | 58,833 | 26.929 | 37.688 |
| # of Industries Applied | 58,833 | 8.164 | 7.816 |
| Days to receive positive callback since first application | 45,807 | 3.111 | 7.753 |

*Notes:* This table reports summary statistics for key outcome variables used in the analysis. Panel A summarizes the characteristics of jobs that applicants apply to. Panel B summarizes applicant-level annual search behavior.

Table 3. The Impact of LLM Exposure on Job Posting Outcomes

**Panel A: Occupation Category-Year Aggregated Outcomes**

| VARIABLES | (1) Job Posting Share (%) | (2) Vacancy Share (%) |
|---|---|---|
| Post × LLM Exposure | -0.217** | -0.229*** |
| | (0.090) | (0.067) |
| Observations | 5,059 | 5,059 |
| Adjusted $R^2$ | 0.794 | 0.854 |
| Pre-policy Mean | 0.0770 | 0.0771 |

**Panel B: Individual Job Posting-Level Outcomes**

| VARIABLES | (3) Exp. Requirement $\geq 3$ | (4) Edu. Requirement $\geq$ Bachelor | (5) Log Wage | (6) Time-related Amenity |
|---|---|---|---|---|
| Post × LLM Exposure | 0.028*** | 0.037*** | -0.121*** | -0.101*** |
| | (0.008) | (0.008) | (0.045) | (0.019) |
| Observations | 777,687 | 777,687 | 777,687 | 777,687 |
| Adjusted $R^2$ | 0.347 | 0.310 | 0.443 | 0.108 |
| Pre-policy Mean | 0.227 | 0.203 | 8.989 | 0.263 |

*Notes:* This table reports the estimated effects of LLM exposure on job posting outcomes using the job posting data at different levels of aggregation. Panel A aggregates the original job posting data to the occupation category-year level. Panel B uses individual job postings as the unit of observation. In panel A, all specifications control for year fixed effects and occupation classification fixed effects. In Panel B, all specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Standard errors clustered at the occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4. Impact of LLM Shock on Applied Job Characteristics and Employer Response

| Group | (1) Full Sample | (2) Senior Applicants (Exp. $\geq 3$ years) | (3) Junior Applicants (Exp. $< 3$ years) | (4) Graduated from Elite Universities | (5) Graduated from Other Universities and Below |
|---|---|---|---|---|---|
| **Panel A: Outcome = Log Wage** | | | | | |
| Post × LLM Exposure | -0.020 | -0.020 | -0.006 | 0.050 | -0.026 |
| | (0.016) | (0.020) | (0.021) | (0.058) | (0.016) |
| Observations | 1,574,232 | 1,077,996 | 496,236 | 158,019 | 1,416,213 |
| Adjusted $R^2$ | 0.305 | 0.329 | 0.182 | 0.385 | 0.287 |
| Pre-treatment Mean | 9.02 | 9.09 | 8.86 | 9.24 | 8.99 |
| **Panel B: Outcome = Time-related Amenity** | | | | | |
| Post × LLM Exposure | -0.067*** | -0.081*** | -0.020 | -0.122*** | -0.065*** |
| | (0.012) | (0.014) | (0.016) | (0.029) | (0.013) |
| Observations | 1,574,232 | 1,077,996 | 496,236 | 158,019 | 1,416,213 |
| Adjusted $R^2$ | 0.052 | 0.057 | 0.048 | 0.077 | 0.053 |
| Pre-treatment Mean | 0.29 | 0.27 | 0.31 | 0.29 | 0.29 |
| **Panel C: Outcome = Callback Rate** | | | | | |
| Post × LLM Exposure | -0.002 | 0.027* | -0.088*** | 0.025 | -0.002 |
| | (0.013) | (0.014) | (0.021) | (0.033) | (0.013) |
| Observations | 1,573,996 | 1,077,807 | 496,189 | 157,987 | 1,416,009 |
| Adjusted $R^2$ | 0.072 | 0.077 | 0.077 | 0.096 | 0.073 |
| Pre-treatment Mean | 0.29 | 0.27 | 0.32 | 0.22 | 0.30 |

*Notes:* This table reports the estimated effects of LLM exposure on applied job characteristics across different applicant subgroups. Panels A–C report effects on log posted wage, time-related amenity, and an indicator for receiving a callback, respectively. Columns (1)–(5) present results for the full sample, senior applicants (Exp. $\geq 3$ years), junior applicants (Exp. $< 3$ years), applicants graduated from elite universities, and applicants graduated from non-elite universities and below, respectively. The key regressor is the interaction between the post indicator and the LLM exposure measure of the applied job. Standard errors clustered at the applied-job occupation classification level are reported in parentheses. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5. LLM Shock and Job Seeker's Mobility Across Occupations

| Outcome | LLM Exposure of Applied Jobs | LLM Exposure of Applied Jobs, < Last Job | LLM Exposure of Applied Jobs, > Last Job | Occupation Class: last ≠ applied |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Post × LLM Exposure | -0.016** | 0.028 | -0.027* | 0.002 |
| | (0.008) | (0.028) | (0.015) | (0.024) |
| Observations | 1,538,684 | 1,538,684 | 1,538,684 | 1,574,232 |
| Adjusted $R^2$ | 0.522 | 0.277 | 0.231 | 0.257 |
| Pre-treatment Mean | 0.59 | 0.32 | 0.22 | 0.55 |

*Notes:* This table reports the estimated effects of LLM exposure on job seekers' occupational mobility. Each column reports occupational mobility outcomes: the LLM exposure level of applied jobs (column 1), whether the LLM exposure of applied jobs is lower or higher than that of the last job (columns 2–3), and whether the applied occupation class differs from the last job (column 4). The key regressor is the interaction between the post indicator and the LLM exposure measure. Standard errors clustered at the applied-job occupation classification level are reported in parentheses. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6. Applicant Level: Searching Effort

| Group | (1) Full Sample | (2) Senior Applicants (Exp. $\geq 3$ years) | (3) Junior Applicants (Exp. $< 3$ years) | (4) Graduated from Elite Universities | (5) Graduated from Other Universities and Below |
|---|---|---|---|---|---|
| **Panel A: Outcome = # of Applications** | | | | | |
| Post × LLM Exposure | -5.437*** | -7.217*** | -0.694 | -4.447 | -5.371*** |
| | (1.459) | (1.777) | (3.423) | (8.178) | (1.459) |
| Observations | 58,833 | 42,910 | 15,923 | 5,618 | 53,215 |
| Adjusted $R^2$ | 0.075 | 0.072 | 0.127 | 0.163 | 0.081 |
| Pre-treatment Mean | 25.36 | 24.15 | 28.60 | 27.31 | 25.15 |
| **Panel B: Outcome = # of Industries Applied** | | | | | |
| Post × LLM Exposure | -0.629* | -1.266*** | 1.102 | 0.280 | -0.645* |
| | (0.334) | (0.383) | (0.717) | (1.536) | (0.332) |
| Observations | 58,833 | 42,910 | 15,923 | 5,618 | 53,215 |
| Adjusted $R^2$ | 0.118 | 0.112 | 0.166 | 0.208 | 0.122 |
| Pre-treatment Mean | 7.87 | 7.59 | 8.60 | 8.32 | 7.82 |
| **Panel C: Outcome = Days to First Positive Callback** | | | | | |
| Post × LLM Exposure | -1.352*** | -1.776*** | -0.101 | -3.450* | -1.277*** |
| | (0.343) | (0.396) | (0.637) | (1.871) | (0.364) |
| Observations | 45,807 | 32,704 | 13,103 | 4,152 | 41,655 |
| Adjusted $R^2$ | 0.084 | 0.093 | 0.118 | 0.265 | 0.083 |
| Pre-treatment Mean | 3.22 | 3.47 | 2.61 | 4.75 | 3.06 |

*Notes:* This table reports the estimated effects of LLM exposure on applicants' annual search behavior across different applicant subgroups. Panels A–C report effects on the number of applications submitted, the number of industries applied to, and the number of days to receive the first positive callback after the first application, respectively. Columns (1)–(5) present results for the full sample, senior applicants (Exp. $\geq 3$ years), junior applicants (Exp. $< 3$ years), applicants graduated from elite universities, and applicants graduated from non-elite universities and below, respectively. The key regressor is the interaction between the post indicator and the LLM exposure measure. Standard errors clustered at the applied-job occupation classification level are reported in parentheses. All regressions include last-job occupation classification fixed effects, year fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

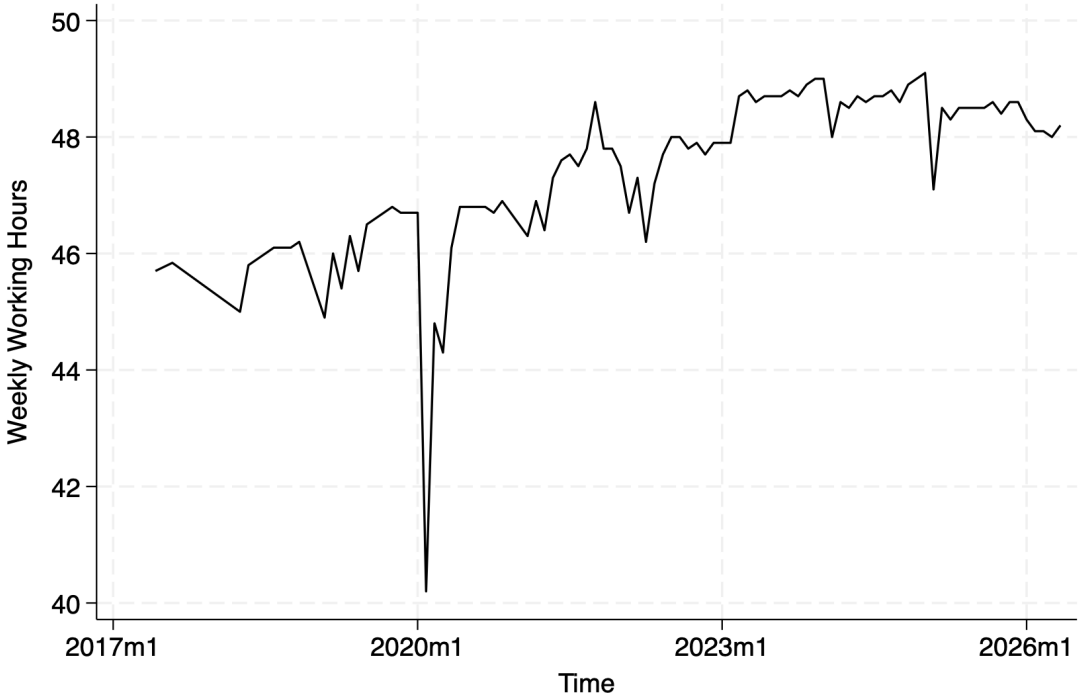
Table 7. LLM Exposure, Worker Win Rate in Labor Disputes, and Job Characteristics

| Outcome: | Log Wage | | Time Amenity | |
|---|---|---|---|---|
| | High | Low | High | Low |
| | (1) | (2) | (3) | (4) |
| Post × LLM Exposure | 0.010 | -0.039** | -0.065*** | -0.071*** |
| | (0.020) | (0.018) | (0.015) | (0.014) |
| Observations | 695,898 | 874,451 | 695,898 | 874,451 |
| Adjusted $R^2$ | 0.298 | 0.338 | 0.056 | 0.058 |
| Pre-policy Mean | 8.93 | 9.09 | 0.28 | 0.29 |

*Notes*: This table reports heterogeneity in the effects of the LLM shock on applied job characteristics by the regional worker win rate in labor disputes. Columns (1)–(2) report effects on log wage and columns (3)–(4) report effects on time amenity. Within each outcome group, the sample is split by whether the applicant's last-job province has a high or low worker win rate in labor disputes. The table uses applicant-side application record data, where outcomes are characteristics of jobs applied to by applicants. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). Standard errors clustered at the applied-job occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

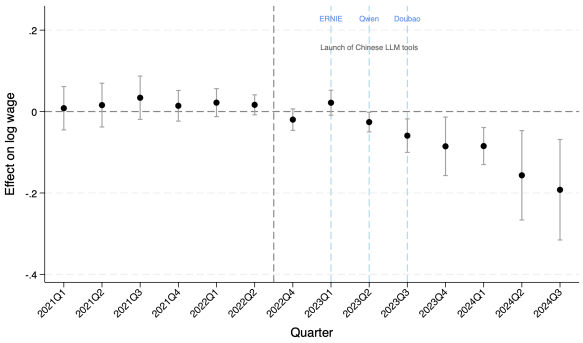
# Figures

Figure 1. Time Trend of Weekly Working Hours in China

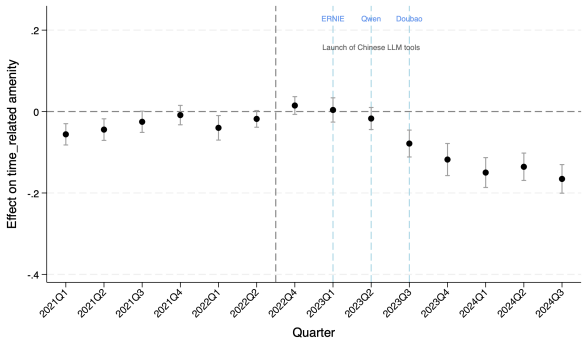



*Notes:* This figure presents the time trend of weekly working hours in China. Data is from the National Bureau of Statistics of China.

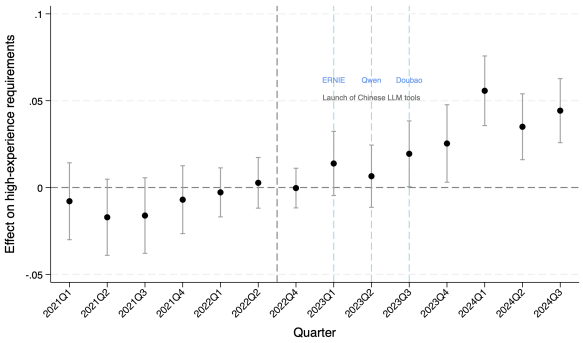
Figure 2. Dynamic Effects of LLM Exposure on Job Posting Outcomes



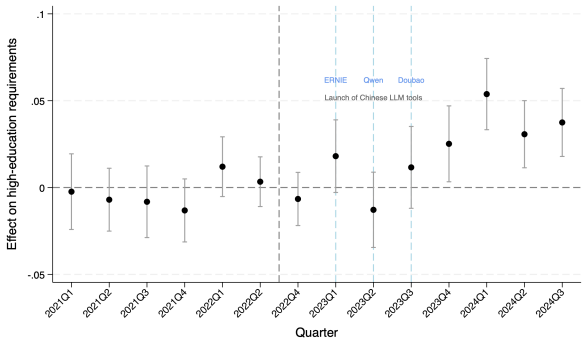

(a) Log Wage




(b) Time-related Amenity



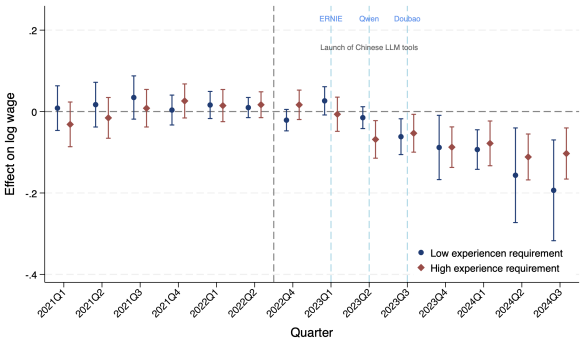

(c) Exp. Requirement $\geq 3$



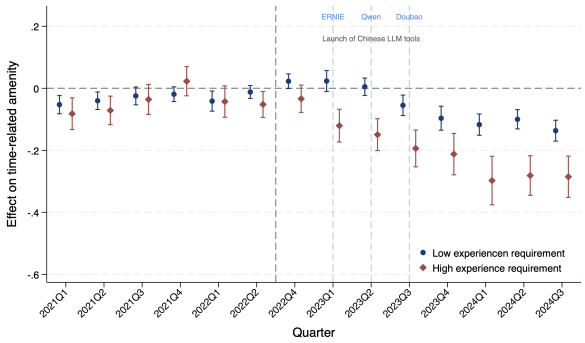

(d) Edu. Requirement $\geq$ Bachelor

*Notes:* This figure presents the dynamic effects of LLM exposure on job posting-level outcomes corresponding to Panel B of Table 3. The event time is defined at the quarter level, with 2022Q3 omitted as the reference period. Dots represent point estimates, and vertical bars report 95% confidence intervals. The vertical dashed line marks the beginning of the post period, and the light-blue dashed lines indicate the launch periods of major Chinese LLM tools. All specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Standard errors are clustered at the occupation classification level.

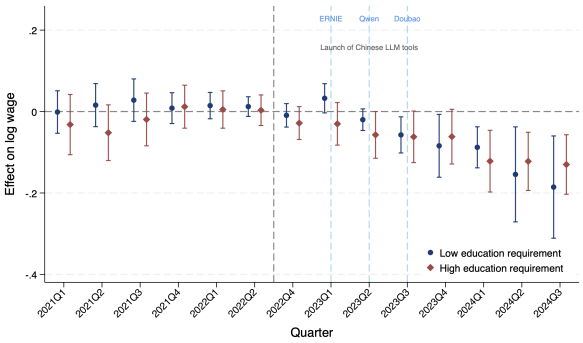
Figure 3. Event-study Plot: Demand-Side Wage and Amenity Heterogeneity



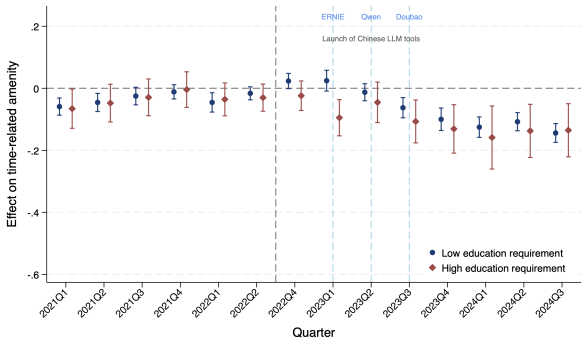

(a) Log Wage (By Exp)




(b) Time-related Amenity (By Exp)



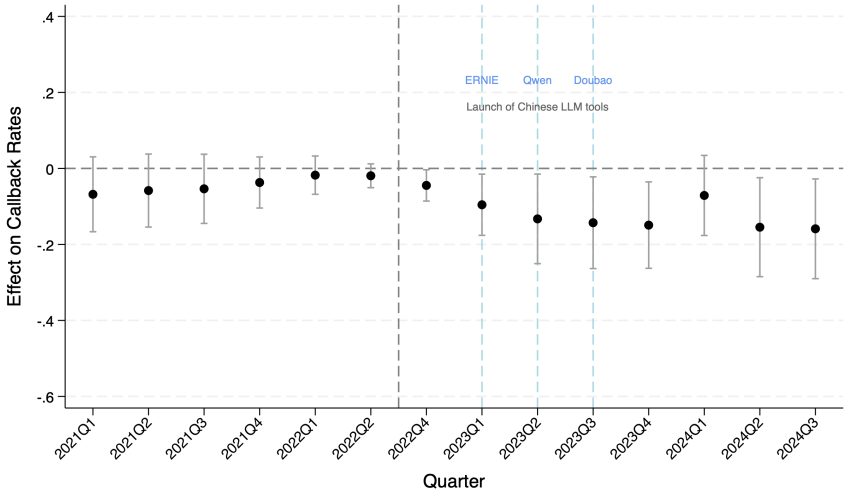

(c) Log Wage (By Edu)



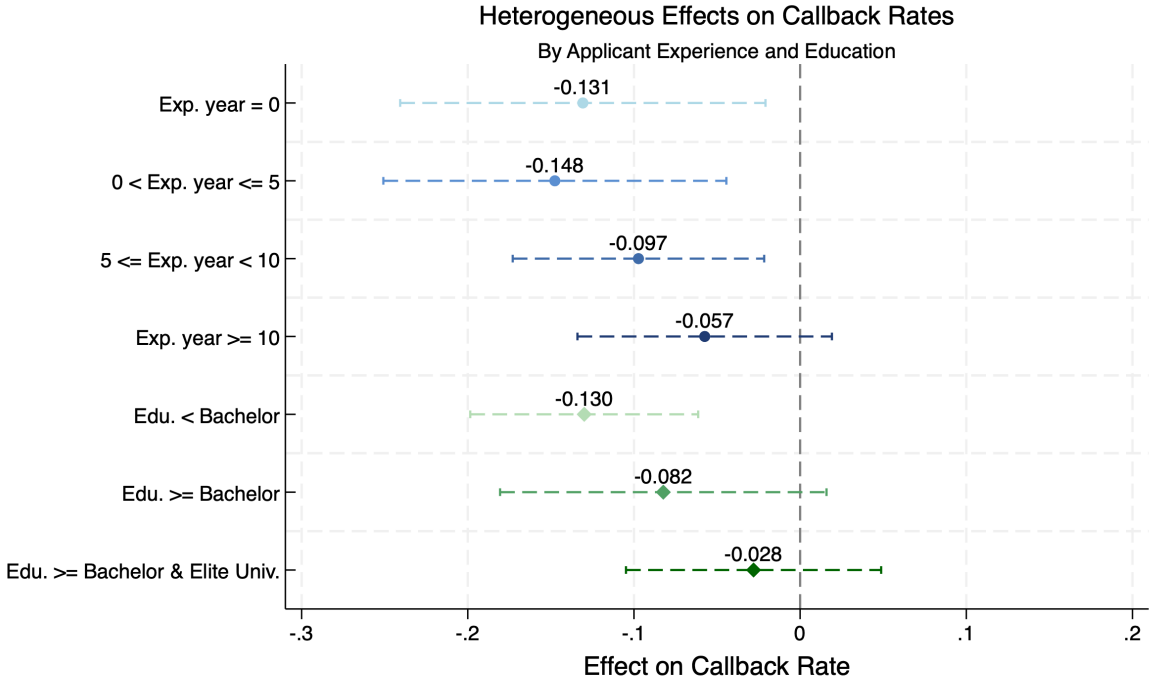

(d) Time-related Amenity (By Edu)

*Notes:* This figure presents the dynamic heterogeneous effects of LLM exposure on demand-side job posting outcomes. Panels (a) and (b) show heterogeneity by experience requirement, while panels (c) and (d) show heterogeneity by education requirement. The event time is defined at the quarter level, with 2022Q3 omitted as the reference period. Dots represent point estimates, and vertical bars report 95% confidence intervals. The vertical dashed line marks the beginning of the post period, and the light-blue dashed lines indicate the launch periods of major Chinese LLM tools. All specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Standard errors are clustered at the occupation classification level.

Figure 4. Effects of LLM Exposure on Callback Rates



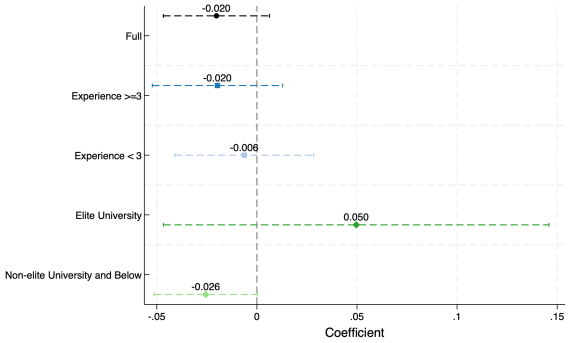

(a) Dynamic Effects on Callback Rates



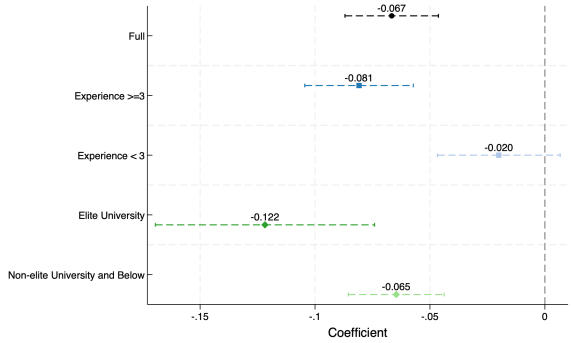

(b) Heterogeneous Effects by Applicant Experience and Education

*Notes:* This figure presents the effects of LLM exposure on callback rates on the employer side. Panel (a) reports the dynamic effects at the quarter level, with 2022Q3 omitted as the reference period. Dots represent point estimates, and vertical bars indicate 95% confidence intervals. The vertical dashed line marks the beginning of the post period, and the light-blue dashed lines indicate the launch periods of major Chinese LLM tools. Panel (b) reports heterogeneous effects by applicants' years of work experience, education level, and elite-university background. All specifications control for year-by-month fixed effects, job-city fixed effects, occupation classification fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). Standard errors are clustered at the occupation classification level.

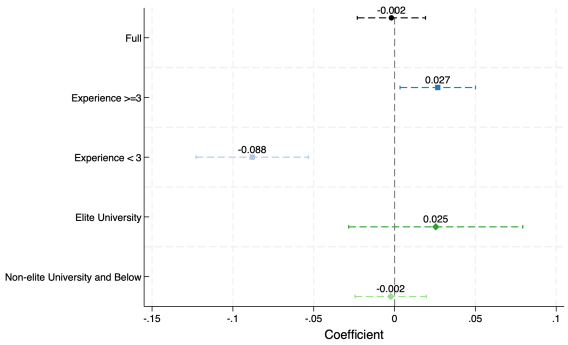
Figure 5. Impact of LLM Shock on Applied Job Characteristics and Employer Response




(a) Log Wage



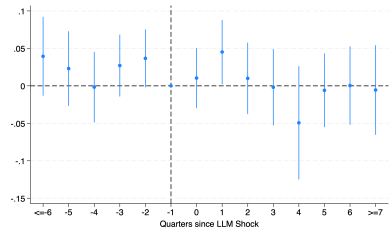

(b) Time-related Amenity



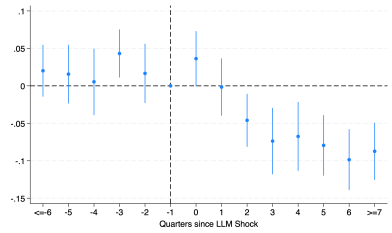

(c) Callback Rate

*Notes:* This figure summarizes the estimated impact of LLM exposure on applied job characteristics across subsamples. Each panel reports point estimates and 95% confidence intervals from separate regressions corresponding to Panels A–C of Table 4: Full sample, Experience $\geq 3$, Experience $< 3$, Elite University, and Others. Panel (a) uses log posted wage, panel (b) uses time-related amenity, and panel (c) uses an indicator for callback as the dependent variable. Standard errors are clustered at the applied-job occupation classification level. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels).

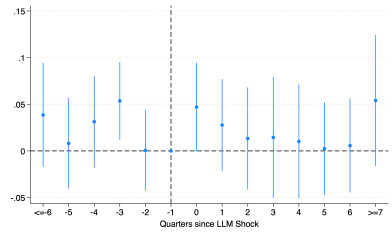
Figure 6. Event-study Plot: Applicant Side



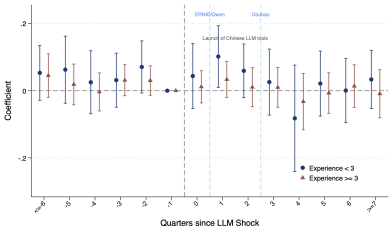

(a) Log Wage (Full)



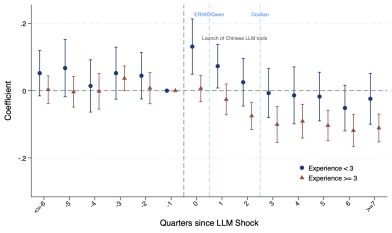

(b) Time-related Amenity (Full)



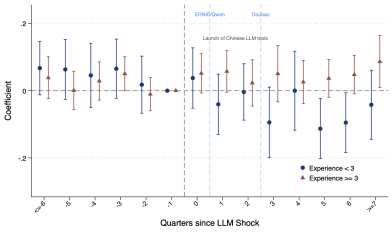

(c) Callback Rate (Full)



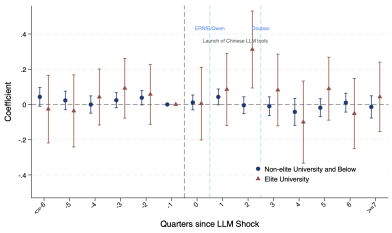

(d) Log Wage (By Exp)



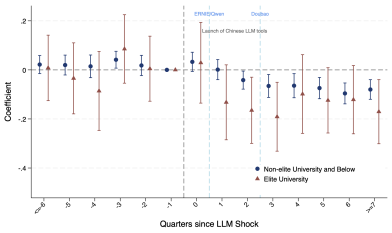

(e) Time-related Amenity (By Exp)



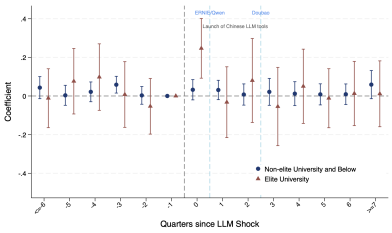

(f) Callback Rate (By Exp)




(g) Log Wage (By Edu)




(h) Time-related Amenity (By Edu)




(i) Callback Rate (By Edu)

*Notes:* This figure presents event-study estimates of the impact of LLM exposure on applied job characteristics and employer response. The top row reports full-sample estimates. The middle row reports heterogeneity by applicant experience, and the bottom row reports heterogeneity by applicant education. Columns correspond to log posted wage, time-related amenity, and an indicator for callback, respectively. Dots represent point estimates, and vertical bars report 95% confidence intervals. Standard errors are clustered at the applied-job occupation classification level. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels).

# Appendix

Table A1. LLM Exposure Index by Occupation Category

| First-level Occupation | Second-level Occupation | LLM Exposure Index |
|---|---|---|
| Internet/Telecommunications/Hardware | Artificial Intelligence | 0.7500 |
| Internet/Telecommunications/Hardware | Front-end Development | 0.8622 |
| Internet/Telecommunications/Hardware | Semiconductors/Chips | 0.5122 |
| Internet/Telecommunications/Hardware | Data Engineers | 0.7355 |
| Internet/Telecommunications/Hardware | Electronics/Electrical/Automation | 0.4558 |
| Internet/Telecommunications/Hardware | Mobile Development | 0.8596 |
| Internet/Telecommunications/Hardware | Software Development | 0.8248 |
| Internet/Telecommunications/Hardware | Telecommunications and Hardware R&D | 0.5346 |
| Product/Project/Senior Management | Product Managers | 0.6816 |
| Product/Project/Senior Management | Project Management | 0.5530 |
| Product/Project/Senior Management | Senior Management | 0.5606 |
| Human Resources/Administration/Finance/Legal Affairs | Human Resources | 0.5524 |
| Human Resources/Administration/Finance/Legal Affairs | Compliance, Risk Control, Legal Affairs, and Lawyers | 0.7083 |
| Human Resources/Administration/Finance/Legal Affairs | Clerks/Assistants | 0.5456 |
| Human Resources/Administration/Finance/Legal Affairs | Records Management | 0.7007 |
| Human Resources/Administration/Finance/Legal Affairs | Administration | 0.6029 |
| Human Resources/Administration/Finance/Legal Affairs | Finance | 0.8613 |
| Agriculture/Energy/Environmental Protection | Agriculture/Forestry/Animal Husbandry/Fishery | 0.4406 |
| Agriculture/Energy/Environmental Protection | Environmental Science/Environmental Protection | 0.5628 |
| Agriculture/Energy/Environmental Protection | Energy/Minerals/Geology | 0.5959 |

Continued on next page

Table A1 – continued from previous page

| First-level Occupation | Second-level Occupation | LLM Exposure Index |
|---|---|---|
| Healthcare/Medical Aesthetics/Medical Services | Clinical Research/Trials | 0.7155 |
| Healthcare/Medical Aesthetics/Medical Services | Physicians/Pharmacists | 0.3769 |
| Healthcare/Medical Aesthetics/Medical Services | Medical Devices | 0.7600 |
| Healthcare/Medical Aesthetics/Medical Services | Pharmaceutical Marketing/Sales | 0.7944 |
| Healthcare/Medical Aesthetics/Medical Services | Nurses/Medical Assistants | 0.3787 |
| Healthcare/Medical Aesthetics/Medical Services | Biomedical/Pharmaceutical R&D | 0.5716 |
| Business Services/Life Services | Health/Beauty | 0.5773 |
| Business Services/Life Services | Business Services | 0.6174 |
| Business Services/Life Services | Domestic Services/Repair | 0.1520 |
| Business Services/Life Services | Tourism Services | 0.6781 |
| Business Services/Life Services | Property Management/Security | 0.2613 |
| Business Services/Life Services | Sports and Fitness | 0.4965 |
| Business Services/Life Services | Hotel Services | 0.5419 |
| Business Services/Life Services | Retail and Department Stores | 0.6072 |
| Business Services/Life Services | Food Services | 0.2854 |
| Marketing/Branding/Public Relations | Public Relations and Media | 0.6847 |
| Marketing/Branding/Public Relations | Marketing/Brand Promotion | 0.7330 |
| Marketing/Branding/Public Relations | Market Research | 0.6295 |
| Marketing/Branding/Public Relations | Planning | 0.7658 |
| Film, Television, and Media | Advertising | 0.6194 |
| Film, Television, and Media | Film and Television Production | 0.4117 |
| Film, Television, and Media | Performers/Agents | 0.5958 |
| Real Estate/Engineering/Construction | Engineering Safety/Quality Inspection | 0.3647 |
| Real Estate/Engineering/Construction | Engineering Development Technicians | 0.4625 |
| Real Estate/Engineering/Construction | Engineering Management | 0.5100 |
| Real Estate/Engineering/Construction | Architecture/Interior Design | 0.5331 |
| Real Estate/Engineering/Construction | Real Estate Transactions | 0.7132 |
| Real Estate/Engineering/Construction | Real Estate Planning and Development | 0.6445 |
| Real Estate/Engineering/Construction | Construction Workers | 0.3101 |

Continued on next page

Table A1 – continued from previous page

| First-level Occupation | Second-level Occupation | LLM Exposure Index |
|---|---|---|
| Education/Training/Research | Primary and Secondary School Curriculum Tutoring | 0.5923 |
| Education/Training/Research | Arts and Talent Training | 0.5724 |
| Education/Training/Research | Academic Affairs Management | 0.6775 |
| Education/Training/Research | Early Childhood Education | 0.4618 |
| Education/Training/Research | Scientific/Academic Research | 0.5573 |
| Education/Training/Research | Vocational Training | 0.7450 |
| Education/Training/Research | Language Training | 0.6071 |
| Logistics/Procurement/Supply Chain | Transportation | 0.3064 |
| Logistics/Procurement/Supply Chain | Warehousing Management | 0.4575 |
| Logistics/Procurement/Supply Chain | Logistics Management | 0.5941 |
| Logistics/Procurement/Supply Chain | Delivery and Sorting | 0.2019 |
| Logistics/Procurement/Supply Chain | Procurement/Supply Chain/Materials Management | 0.6061 |
| Manufacturing/Operations Management | Chemicals | 0.4886 |
| Manufacturing/Operations Management | Printing and Packaging | 0.1052 |
| Manufacturing/Operations Management | General Workers/Technicians | 0.1627 |
| Manufacturing/Operations Management | Apparel/Textiles/Leather | 0.5123 |
| Manufacturing/Operations Management | Mechanical Equipment Maintenance | 0.2179 |
| Manufacturing/Operations Management | Mechanical Design/Manufacturing | 0.4992 |
| Manufacturing/Operations Management | Automobile Manufacturing | 0.5531 |
| Manufacturing/Operations Management | Automobile Sales and Services | 0.5410 |
| Manufacturing/Operations Management | Production Management | 0.4738 |
| Manufacturing/Operations Management | Production Quality Management | 0.4832 |
| Management Trainees/Non-corporate Occupations | Management Trainees/Reserve Cadres | 0.6280 |
| Editing/Journalism/Translation | Editors/Proofreaders/Writers | 0.8517 |
| Editing/Journalism/Translation | Translation | 0.7967 |
| Visual/Interaction/Design | Animation and Motion Design | 0.5908 |
| Visual/Interaction/Design | Industrial Design | 0.6287 |
| Visual/Interaction/Design | Visual/Interaction Design | 0.6577 |

Continued on next page

Table A1 – continued from previous page

| First-level Occupation | Second-level Occupation | LLM Exposure Index |
|---|---|---|
| Operations and Testing | Pre-sales and After-sales Engineers | 0.6960 |
| Operations and Testing | Test Engineers | 0.7028 |
| Operations and Testing | Operations Support | 0.6214 |
| Operations/Professional Analysis | Professional Analysis | 0.6738 |
| Operations/Professional Analysis | Business Operations | 0.6784 |
| Operations/Professional Analysis | Customer Service | 0.8509 |
| Operations/Professional Analysis | New Media Operations | 0.7663 |
| Operations/Professional Analysis | E-commerce Operations | 0.6871 |
| Operations/Professional Analysis | Operations Management | 0.5851 |
| Finance/Insurance | Insurance | 0.7910 |
| Finance/Insurance | Investment and Financing | 0.7320 |
| Finance/Insurance | Securities/Funds/Futures | 0.7557 |
| Finance/Insurance | Banking and Financial Services | 0.7361 |
| Sales/Business Development | Business Development | 0.7046 |
| Sales/Business Development | Sales Management | 0.7059 |
| Sales/Business Development | Sales Consultants | 0.7979 |

Table A2. Summary Statistics: Regional Labor Protection

| Variables | Mean | Std. Dev. |
|---|---|---|
| # of Labor Dispute Cases Accepted | 34,450.645 | 34,203.186 |
| # of Labor Dispute Cases Settled | 34,412.258 | 34,382.821 |
| # of Cases Worker Wins | 10,118.677 | 7,568.519 |
| # of Cases Employer Wins | 3,634.355 | 4,543.941 |
| # of Cases Partial Wins or Other Outcomes | 20,659.226 | 24,386.861 |
| Employed Population | 2,117,154.400 | 1,518,910.700 |
| Labor Dispute Cases per Worker | 0.019 | 0.021 |
| Worker Win Rate | 0.366 | 0.137 |

*Notes:* This table reports summary statistics for province-level measures of labor dispute resolution and regional labor protection. The labor dispute variables are measured in 2019 and are obtained from provincial labor statistical yearbooks. The employed population is measured at the province level using the 2020 Population Census. Labor Dispute Cases per Worker is defined as the number of labor dispute cases accepted in 2019 divided by the employed population from the 2020 Census. Worker Win Rate is defined as the number of cases won by workers divided by the number of settled labor dispute cases in 2019. The sample includes 31 provinces.

Table A3. Robustness Checks: Alternative Control Specifications

**Panel A: Including City-by-Time Fixed Effects**

| VARIABLES | (1) Log Wage | (2) Time-related Amenity | (3) Exp. Requirement $\geq 3$ | (4) Edu. Requirement $\geq$ Bachelor |
|---|---|---|---|---|
| Post × LLM Exposure | -0.112*** | -0.100*** | 0.031*** | 0.038*** |
| | (0.041) | (0.018) | (0.008) | (0.008) |
| Observations | 776,106 | 776,106 | 776,106 | 776,106 |
| Adjusted $R^2$ | 0.448 | 0.112 | 0.347 | 0.311 |
| Pre-policy Mean | 8.989 | 0.263 | 0.227 | 0.203 |

**Panel B: Including City-by-Time Fixed Effects and Industry-Specific Time Trends**

| VARIABLES | (1) Log Wage | (2) Time-related Amenity | (3) Exp. Requirement $\geq 3$ | (4) Edu. Requirement $\geq$ Bachelor |
|---|---|---|---|---|
| Post × LLM Exposure | -0.064*** | -0.107*** | 0.010 | 0.017*** |
| | (0.019) | (0.017) | (0.008) | (0.006) |
| Observations | 776,106 | 776,106 | 776,106 | 776,106 |
| Adjusted $R^2$ | 0.462 | 0.127 | 0.356 | 0.317 |
| Pre-policy Mean | 8.989 | 0.263 | 0.227 | 0.203 |

*Notes:* This table reports robustness checks using alternative control specifications. All specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Panel A adds city-by-time fixed effects. Panel B further controls for industry-specific time trends. Standard errors clustered at the occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A4. Robustness Check: Weighted Regressions by Number of Vacancies

| VARIABLES | (1) Log Wage | (2) Time-related Amenity | (3) Exp. Requirement $\geq 3$ | (4) Edu. Requirement $\geq$ Bachelor |
|---|---|---|---|---|
| Post × LLM Exposure | -0.058* | -0.044*** | 0.018** | 0.026*** |
| | (0.033) | (0.015) | (0.008) | (0.007) |
| Observations | 2,706,316 | 2,706,316 | 2,706,316 | 2,706,316 |
| Adjusted $R^2$ | 0.420 | 0.108 | 0.357 | 0.283 |
| Pre-policy Mean | 8.990 | 0.273 | 0.240 | 0.219 |

*Notes:* This table reports robustness checks using weighted regressions. Regressions are weighted by the number of vacancies in each job posting using frequency weights. To reduce the influence of outliers, we exclude job postings with more than 20 vacancies. All specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Standard errors clustered at the occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A5. Robustness Check: Using High LLM Exposure Dummy

| VARIABLES | (1) Log Wage | (2) Time-related Amenity | (3) Exp. Requirement $\geq 3$ | (4) Edu. Requirement $\geq$ Bachelor |
|---|---|---|---|---|
| Post × High LLM Exposure Dummy | -0.035** | -0.024** | 0.007* | 0.008* |
| | (0.017) | (0.010) | (0.004) | (0.004) |
| Observations | 777,687 | 777,687 | 777,687 | 777,687 |
| Adjusted $R^2$ | 0.443 | 0.107 | 0.346 | 0.310 |
| Pre-policy Mean | 8.989 | 0.263 | 0.227 | 0.203 |

*Notes:* This table reports robustness checks using a dummy measure of LLM exposure. High LLM exposure is defined as an indicator equal to one if an occupation's LLM exposure is above the median, and zero otherwise. All specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Standard errors clustered at the occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A6. Robustness Check: Alternative Definitions of Time-Related Amenity

| VARIABLES | (1) Strict Definition | (2) Broad Definition |
|---|---|---|
| Post × High LLM Exposure | -0.130*** | -0.114*** |
| | (0.019) | (0.027) |
| Observations | 777,687 | 777,687 |
| Adjusted $R^2$ | 0.115 | 0.098 |
| Pre-policy Mean | 0.261 | 0.264 |

*Notes:* This table reports robustness checks using alternative definitions of time-related amenity. Column (1) uses a stricter definition that only captures explicit two-day-weekend and no-overtime phrases. Column (2) uses a broader definition that further includes phrases related to regular working hours and eight-hour work schedules. The baseline measure used in the main analysis captures core time-friendly signals, including two-day weekends and no-overtime arrangements. All specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Standard errors clustered at the occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A7. Heterogeneous Effects by Job Requirements on the Demand Side

| | Log Wage | | | | Time-related Amenity | | | |
|---|---|---|---|---|---|---|---|---|
| | Exp. Requirement $\geq$ 3 Years | Exp. Requirement $<$ 3 Years | Edu. Requirement $\geq$ Bachelor | Edu. Requirement $<$ Bachelor | Exp. Requirement $\geq$ 3 Years | Exp. Requirement $<$ 3 Years | Edu. Requirement $\geq$ Bachelor | Edu. Requirement $<$ Bachelor |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Post × LLM Exposure | -0.090*** | -0.118** | -0.080*** | -0.116** | -0.204*** | -0.075*** | -0.088** | -0.082*** |
| | (0.023) | (0.046) | (0.029) | (0.046) | (0.027) | (0.018) | (0.037) | (0.016) |
| Observations | 148,354 | 629,272 | 157,993 | 619,591 | 148,354 | 629,272 | 157,993 | 619,591 |
| Adjusted $R^2$ | 0.488 | 0.396 | 0.423 | 0.414 | 0.089 | 0.116 | 0.069 | 0.121 |
| Pre-policy Mean | 9.346 | 8.898 | 9.260 | 8.909 | 0.286 | 0.257 | 0.315 | 0.248 |

*Notes:* This table reports heterogeneity in the effects of LLM exposure on job posting outcomes by job requirements. Columns (1)–(4) use log posted wage as the dependent variable, and columns (5)–(8) use time-related amenity as the dependent variable. The sample is split by experience requirements in columns (1), (2), (5), and (6), and by education requirements in columns (3), (4), (7), and (8). All specifications control for year-by-month fixed effects, job-city fixed effects, and occupation classification fixed effects. Standard errors clustered at the occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A8. Changes in Callback Rates on the Demand Side

| | All Sample | By Applicant Experience | | | | By Applicant Education | | |
|---|---|---|---|---|---|---|---|---|
| | Callback Rate | Exp. Year $= 0$ | $0 <$ Exp. Year $\leq 5$ | $5 \leq$ Exp. Year $< 10$ | Exp. Year $\geq 10$ | Edu. $<$ Bachelor | Edu. $\geq$ Bachelor | Edu. $\geq$ Bachelor & Elite Univ. |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Post $\times$ LLM Exposure | -0.110** | -0.131* | -0.148** | -0.097** | -0.057 | -0.130*** | -0.082 | -0.028 |
| | (0.045) | (0.067) | (0.063) | (0.046) | (0.047) | (0.042) | (0.060) | (0.047) |
| Observations | 2,634,686 | 492,964 | 615,868 | 564,344 | 961,357 | 1,124,336 | 1,256,752 | 304,389 |
| Adjusted $R^2$ | 0.145 | 0.182 | 0.146 | 0.139 | 0.133 | 0.161 | 0.117 | 0.121 |
| Pre-policy Mean | 0.261 | 0.306 | 0.314 | 0.267 | 0.192 | 0.314 | 0.228 | 0.187 |

*Notes:* This table reports the estimated effects of LLM exposure on callback rates. Column (1) uses the full sample. Columns (2)–(5) split the sample by applicants' years of work experience, and columns (6)–(8) split the sample by applicants' education level and elite-university background. All specifications control for year-by-month fixed effects, job-city fixed effects, occupation classification fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). Standard errors clustered at the occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A9. Robustness: Decomposing by Experience × Education Level

| Applicant Sample | (1) Senior & Elite Universities | (2) Senior & Other | (3) Junior & Elite Universities | (4) Junior & Other |
|---|---|---|---|---|
| **Panel A: Outcome = Log Wage** | | | | |
| Post × LLM Exposure | 0.086 | -0.025 | -0.101 | -0.012 |
| | (0.060) | (0.019) | (0.136) | (0.021) |
| Observations | 125,959 | 952,037 | 32,060 | 464,176 |
| Adjusted $R^2$ | 0.427 | 0.312 | 0.306 | 0.172 |
| Pre-treatment Mean | 9.29 | 9.06 | 9.06 | 8.85 |
| **Panel B: Outcome = Time-related Amenity** | | | | |
| Post × LLM Exposure | -0.139*** | -0.079*** | 0.025 | -0.019 |
| | (0.031) | (0.015) | (0.090) | (0.016) |
| Observations | 125,959 | 952,037 | 32,060 | 464,176 |
| Adjusted $R^2$ | 0.085 | 0.057 | 0.109 | 0.050 |
| Pre-treatment Mean | 0.29 | 0.27 | 0.30 | 0.31 |
| **Panel C: Outcome = Callback Rate** | | | | |
| Post × LLM Exposure | 0.046 | 0.027* | -0.024 | -0.093*** |
| | (0.036) | (0.015) | (0.084) | (0.022) |
| Observations | 125,936 | 951,871 | 32,051 | 464,138 |
| Adjusted $R^2$ | 0.103 | 0.078 | 0.149 | 0.078 |
| Pre-treatment Mean | 0.21 | 0.28 | 0.26 | 0.32 |

*Notes:* This table reports robustness checks decomposing the sample by the interaction between experience and education levels. Columns split the sample into four subsamples: senior applicants from elite universities, senior applicants from other universities or below bechalor's degree, junior applicants from elite universities, and junior applicants from other universities or below bechalor's degree. Panels A–C use log posted wage, time-related amenity, and an indicator for callback as the dependent variables, respectively. The key regressor is the interaction between the post indicator and the LLM exposure measure. Standard errors clustered at the applied-job occupation classification level are reported in parentheses. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A10. The Impact of LLM Shock on Applicants' Expected Wage

| Applicant Sample | (1) Full Sample | (2) Senior Applicants (Exp. $\geq$ 3 years) | (3) Junior Applicants (Exp. $<$ 3 years) | (4) Graduated from Elite Universities | (5) Graduated from Other Universities and Below |
|---|---|---|---|---|---|
| **Panel A: Outcome = Applied Wage $<$ Expected Wage** | | | | | |
| Post $\times$ LLM Exposure | -0.005 | -0.003 | -0.002 | -0.042 | -0.006 |
| | (0.015) | (0.018) | (0.019) | (0.051) | (0.016) |
| Observations | 1,431,629 | 943,470 | 488,159 | 124,386 | 1,307,243 |
| Adjusted $R^2$ | 0.027 | 0.028 | 0.040 | 0.087 | 0.028 |
| Pre-treatment Mean | 0.24 | 0.26 | 0.20 | 0.26 | 0.24 |
| **Panel B: Outcome = Expected Wage (Applicant Level)** | | | | | |
| Post $\times$ LLM Exposure | 0.008 | 0.006 | 0.027 | -0.066 | 0.014 |
| | (0.018) | (0.020) | (0.030) | (0.080) | (0.018) |
| Observations | 53,674 | 38,087 | 15,587 | 4,347 | 49,327 |
| Adjusted $R^2$ | 0.259 | 0.260 | 0.259 | 0.385 | 0.254 |
| Pre-treatment Mean | 8.87 | 8.91 | 8.76 | 9.00 | 8.86 |

*Notes:* This table reports the estimated effects of LLM exposure on applicants' expected wages. Panel A uses an indicator equal to one if the applied-job wage is below the applicant's expected wage. Panel B uses the applicant's expected wage as the dependent variable, with the applicant-year as the unit of observation. Columns (1)–(5) present results for the full sample, senior applicants (Exp. $\geq$ 3 years), junior applicants (Exp. $<$ 3 years), applicants graduated from elite universities, and applicants graduated from non-elite universities and below, respectively. The key regressor is the interaction between the post indicator and the LLM exposure measure. Standard errors clustered at the applied-job occupation classification level are reported in parentheses. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A11. The Monotonicity of Wages

| Outcome | Wage (Upper Bound) | | | Wage (Lower Bound) | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| 2nd Experience | 474*** | 1,077*** | 1,086*** | 397*** | 942*** | 949*** |
| | (52.855) | (78.170) | (78.579) | (38.025) | (54.893) | (55.371) |
| 3rd Experience | 1,317*** | 2,216*** | 2,227*** | 993*** | 1,834*** | 1,841*** |
| | (91.365) | (134.530) | (134.830) | (61.728) | (92.804) | (93.254) |
| 4th Experience | 2,245*** | 3,560*** | 3,576*** | 1,604*** | 2,814*** | 2,824*** |
| | (119.305) | (178.627) | (179.856) | (79.476) | (121.061) | (122.208) |
| ≥5th Experience | 3,498*** | 5,464*** | 5,494*** | 2,432*** | 4,187*** | 4,206*** |
| | (158.877) | (253.702) | (255.351) | (103.975) | (171.204) | (173.176) |
| Demographic Control | N | N | Y | N | N | Y |
| Experience Start Year FE | Y | Y | Y | Y | Y | Y |
| Experience Occupation FE | Y | Y | Y | Y | Y | Y |
| Resident City FE | Y | N | N | Y | N | N |
| Applicant FE | N | Y | Y | N | Y | Y |
| Observations | 172,309 | 172,309 | 169,722 | 172,309 | 172,309 | 169,722 |
| Adjusted $R^2$ | 0.260 | 0.772 | 0.772 | 0.265 | 0.779 | 0.779 |
| Outcome Mean | 8,012.91 | 8,012.91 | 8,013.07 | 5,350.14 | 5,350.14 | 5,350.38 |

*Notes:* This table reports applicant–experience-level regressions of wages on indicators for the order of work experience, with the first experience as the omitted category. Columns (1)–(3) use the upper bound of the wage range as the dependent variable, and columns (4)–(6) use the lower bound. Columns (2)–(3) and (5)–(6) include applicant fixed effects, so identification comes from within-person comparisons across successive jobs. Columns (3) and (6) additionally control for individual demographic characteristics. Standard errors clustered at the job occupation classification level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

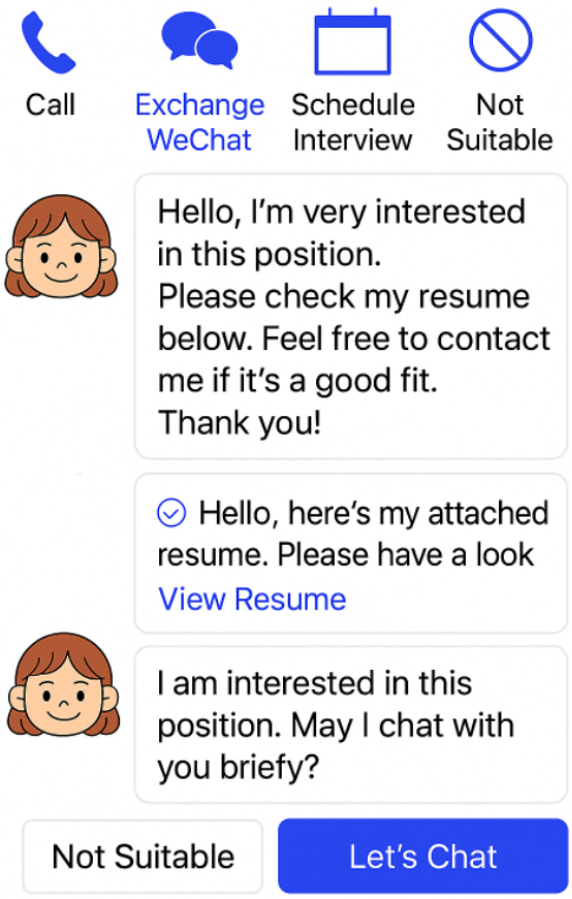
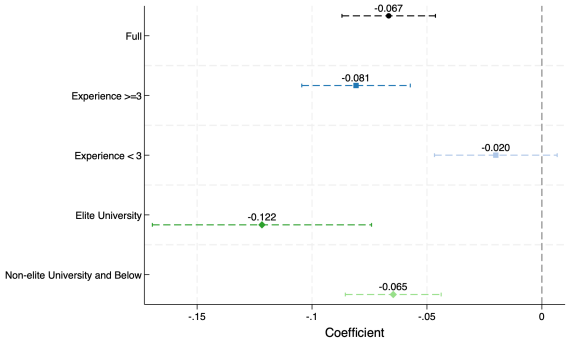


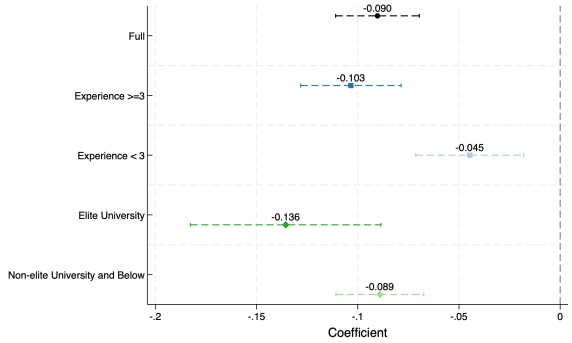

Figure A1. Example of Employer-side Chat Interface on Platform

*Note:* This figure displays the employer-side chat interface on the job platform (translated into English). A response is classified as positive if the employer takes any of the following actions: (1) initiates a phone call; (2) exchanges WeChat accounts (a widely used messaging app in China); (3) schedules an interview; or (4) replies "Let's chat" in response to an applicant's message expressing interest in the position.

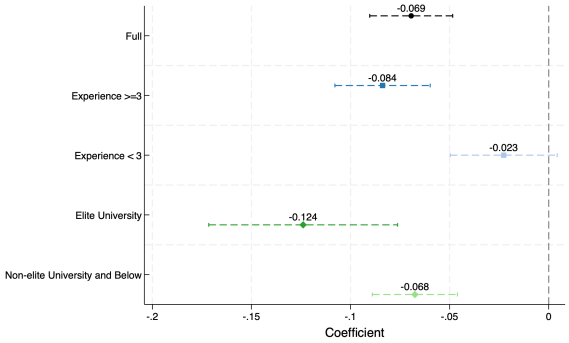
Figure A2. Robustness: Definitions of Time-related Amenity



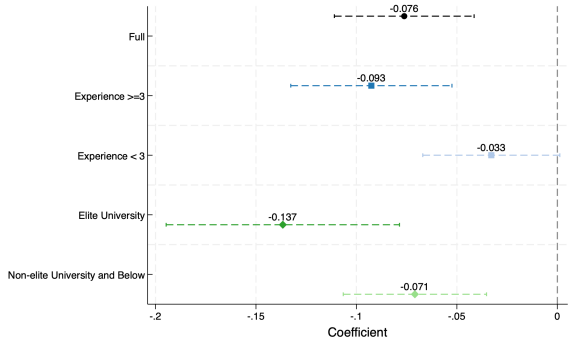

(a) Weekends Off




(b) Weekends Off (Narrow)



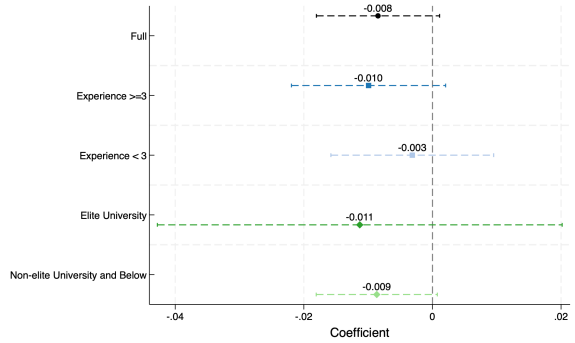

(c) Weekends Off + Eight-hour Workday



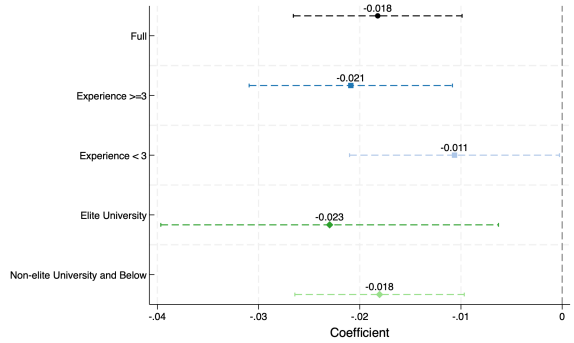

(d) Weekends Off + Eight-hour Workday + Other

*Notes:* This figure presents robustness checks using alternative definitions of time-related amenity as the dependent variable. Panel (a) uses a baseline definition of weekends-off arrangements; panel (b) uses a narrow definition that explicitly statets "weekends off". Panel (c) augments the baseline definition with eight-hour workday phrases. Panel (d) further adds vacation leaves and weaker amenity expressions. Each panel reports point estimates and 95% confidence intervals from separate regressions across the same subsamples as in Figure 5. Standard errors are clustered at the applied-job occupation classification level. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels).

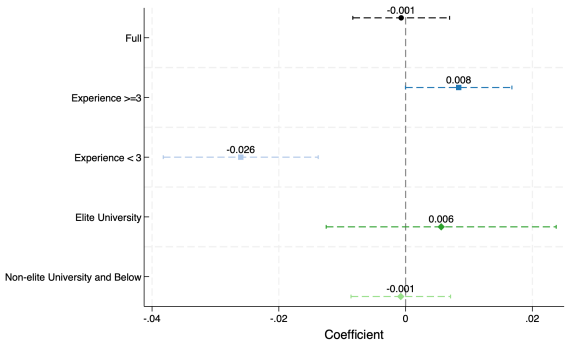
Figure A3. Robustness: Exposure Dummy (High/Low)




(a) Log Wage



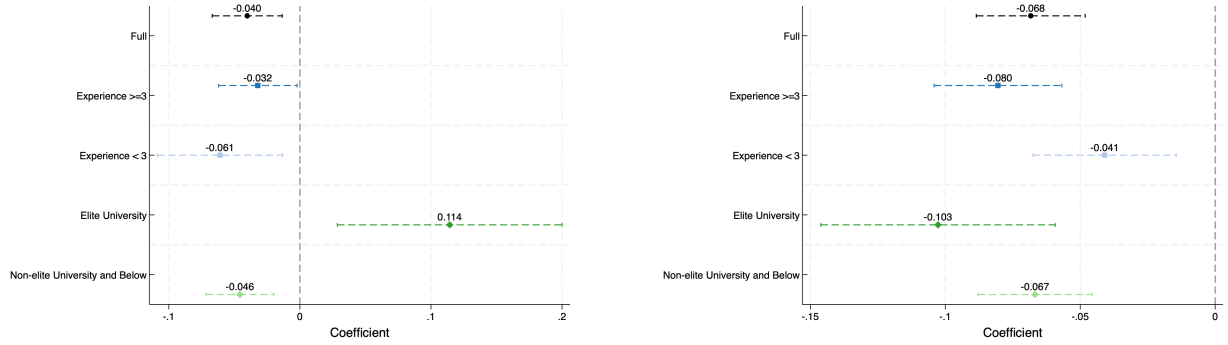

(b) Time-related Amenity




(c) Callback Rate

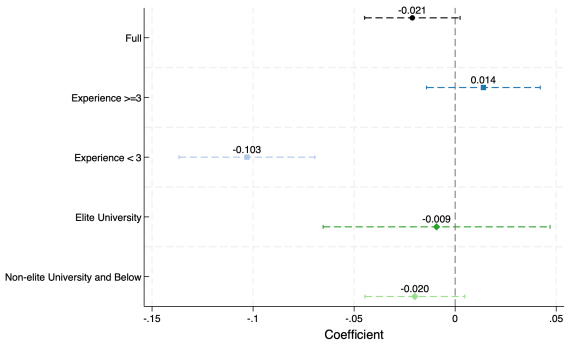
*Notes:* This figure presents robustness checks using a high/low LLM exposure dummy in place of the continuous exposure measure. High LLM exposure is defined as an indicator equal to one if an occupation's LLM exposure is above the median. The dependent variables are log posted wage in panel (a), time-related amenity in panel (b), and an indicator for callback in panel (c). Each panel reports point estimates and 95% confidence intervals across the same subsamples as in Figure 5. Standard errors are clustered at the applied-job occupation classification level. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels).

Figure A4. Robustness: Expected-job Exposure




(a) Log Wage

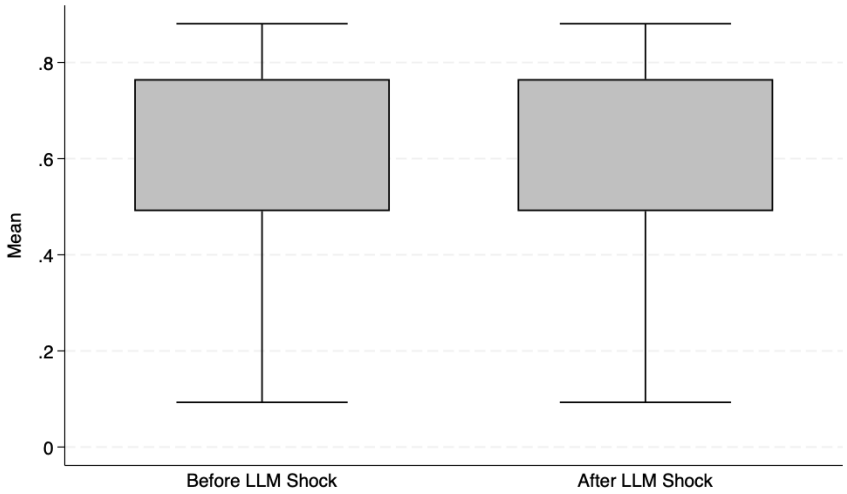
(b) Time-related Amenity




(c) Callback Rate

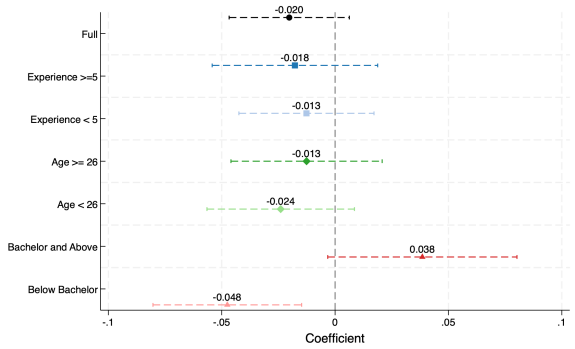
*Notes:* This figure presents robustness checks using the LLM exposure measure of the applicant's expected (rather than applied) job. The dependent variables are log posted wage in panel (a), time-related amenity in panel (b), and an indicator for callback in panel (c). Each panel reports point estimates and 95% confidence intervals across the same subsamples as in Figure 5. Standard errors are clustered at the applied-job occupation classification level. All regressions include expected-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels).

Figure A5. Distribution of LLM Shock Exposure Among Job Seekers

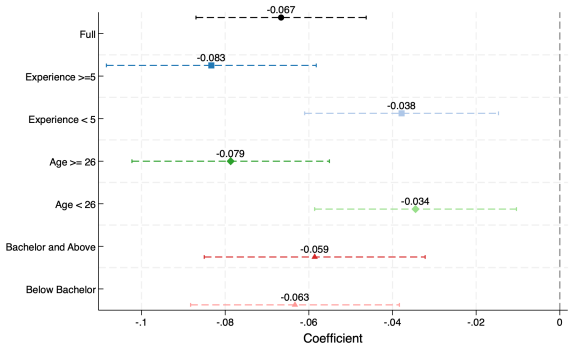



*Notes:* This figure compares the distribution of LLM-shock exposure among job seekers before and after the LLM shock using a box plot. The pre-period corresponds to 2022; the post-period corresponds to the periods following the launch of major Chinese LLM tools.

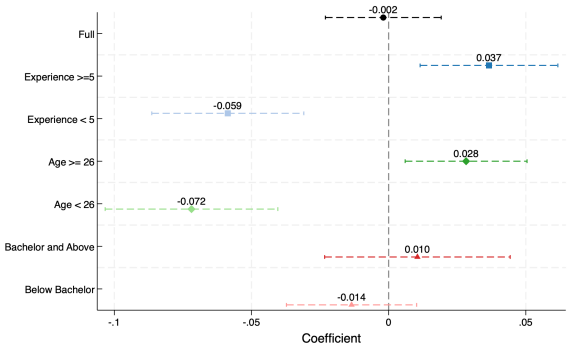
Figure A6. Robustness: Alternative Grouping Criteria




(a) Log Wage




(b) Time-related Amenity




(c) Callback Rate

*Notes:* This figure presents robustness checks using alternative grouping criteria for the subsamples. The dependent variables are log posted wage in panel (a), time-related amenity in panel (b), and an indicator for callback in panel (c). Each panel reports point estimates and 95% confidence intervals from separate regressions under the alternative groupings. Standard errors are clustered at the applied-job occupation classification level. All regressions include last-job occupation classification fixed effects, year-month fixed effects, and individual-level controls (interactions of gender, age, education, and experience levels).