Generative AI and the Pre-Match Labor Market: Evidence from Job Postings, Applications, and Callbacks

Hanming Fang, University of Pennsylvania & NBER
Jiayin Hu, Peking University
Yu Miao, Peking University
Jiaqi Zhao, Peking University

The rapid diffusion of large language models (LLMs) has raised widespread concerns about labor-market disruption. A growing literature documents how generative AI affects within-firm productivity (Brynjolfsson et al., 2025; Noy and Zhang, 2023; Dell'Acqua et al., 2026), employment and earnings (Hui, Reshef, and Zhou, 2024; Humlum and Vestergaard, 2025), and firm labor demand (Eisfeldt et al., 2026; Liu, Wang, and Yu, 2025; Berger et al., 2024). However, much of this work focuses on *realized* labor-market outcomes after employment relationships have been formed, or on firms' *labor demand* adjustments alone. Less is known about how LLM-related demand shocks are transmitted to job seekers before hiring occurs, and through which margins different workers absorb these shocks. Moreover, existing studies have paid limited attention to *non-wage* dimensions, such as time-related amenities, of job postings. This omission is particularly important in the labor markets of developing countries, where long working hours are a salient policy concern and formal protections against working-time deterioration may be weak.

This paper studies how the LLM shock reshapes the pre-match labor market in China. We examine changes in employer demand, job quality, applicant search behavior, and employer–applicant interaction by jointly observing job postings, applications, callbacks, and search dynamics. We focus not only on job quantity and wages, but also on time-related amenities, a hidden yet highly elastic dimension of job quality that captures whether employers advertise regular working hours, weekends off, and no-overtime policies.

Data and methodology. We use daily-level administrative data from one of China's largest online job boards. The platform covers predominantly white-collar occupations across nearly all industries and regions and includes job seekers with diverse levels of education and work experience. The sample spans January 2021 to December 2024, covering the period before and after the widespread deployment of major LLM tools. The data contain two linked components.

The first is a job-posting sample of approximately 0.8 million observations, which includes information on occupation, city, industry, posted wages, education and experience requirements, and unstructured benefit descriptions. We use natural language processing to construct an indicator for time-related amenities from these benefit descriptions. The second component consists of millions of real-time application records linking individual job seekers to specific job postings and employer callback decisions. Each observation in the application data corresponds to an applicant–job pair and contains detailed information on the job listing, applicant characteristics, and employer response.

We measure occupational exposure to LLMs using the index developed by Zhang et al. (2025), which follows the task-based approach of Felten et al. (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 1.25 million job postings from the same platform. Our empirical strategy uses a difference-in-differences framework that interacts predetermined continuous occupational LLM exposure with a post-ChatGPT indicator.

Main findings. We document four main findings. First, on the demand side, high-LLM-exposure occupations experience both quantity and quality deterioration after the shock. Their share of total job postings declines significantly. At the same time, employers increase human-capital requirements: exposed occupations become more likely to require at least three years of experience and a bachelor's degree. These changes suggest that employers demand more qualified candidates while offering less favorable job terms. We find that a one-standard-deviation increase in LLM exposure (about 0.23) is associated with a 2.8 percent decline in posted wages and a 2.3 percentage-point decline in time-related amenities. The adjustment is also heterogeneous across job requirements. Jobs requiring higher education or more experience show slightly smaller wage declines, but jobs with higher experience requirements exhibit larger declines in time-related amenities.

Second, employer callback behavior also changes after the shock. Overall callback rates decline in high-LLM-exposure occupations. However, human capital partially buffers this decline. Applicants with at least five years of experience or an elite-university background face substantially smaller reductions in callback rates than other applicants. This pattern suggests that

the LLM shock increases employers' selectivity and shifts demand toward workers with stronger human-capital signals.

Third, on the applicant side, demand-side pressures are transmitted unevenly to job seekers. We characterize this pattern as multidimensional distributional incidence. Junior applicants, defined as those with less than three years of work experience, bear the shock primarily through monetary and matching penalties: the wages of the jobs they apply to decline, and their callback rates fall substantially. These effects indicate that inexperienced workers are increasingly priced out of high-exposure segments. Senior applicants and applicants from elite universities, by contrast, largely maintain their applied wage levels and callback rates, but face a deterioration in time-related amenities. In other words, they preserve nominal compensation and matching prospects partly by accepting jobs with longer or more rigid working hours. The LLM shock therefore does not simply affect some workers more than others; rather, it channels the burden through different margins depending on workers' human capital.

Fourth, applicants adjust their search behavior, but these adjustments are constrained by limited mobility. Some applicants broaden their occupational search after the shock, but they do not substantially move toward occupations with lower LLM exposure. The magnitude of adjustment is therefore insufficient to offset demand-side deterioration. Senior applicants reduce the total number of applications and narrow their industry scope yet obtain callbacks more quickly. Junior applicants, in contrast, broaden their search across industries but face declining returns to this search intensification. These patterns suggest that applicant-side adjustment is costly and incomplete, especially for workers with weaker human-capital signals.

Finally, we examine whether labor protection moderates the effects of the LLM shock. Using province-level labor-dispute data to proxy for the strength of labor protection, we find that weaker protection amplifies job-quality deterioration. In regions with weaker labor protection, high-LLM-exposure jobs experience larger declines in posted wages and time-related amenities. Applicants in these regions are also pushed toward jobs with lower wages and worse working conditions. These findings suggest that institutional environments can buffer, or amplify, the job-quality costs of technological change.

This paper contributes to the literature in four ways. First, we provide direct evidence on the pre-match stage of labor-market adjustment under AI shocks. Existing studies mainly examine the effects of generative AI on productivity and task performance (Brynjolfsson et al., 2025; Noy and Zhang, 2023; Dell'Acqua et al., 2026), employment and earnings (Hui, Reshef, and Zhou, 2024; Humlum and Vestergaard, 2025), worker exposure and adoption (Eloundou et al., 2024; Bick, Blandin, and Deming, 2026), or firm-side labor demand and job postings (Acemoglu et al., 2022; Eisfeldt et al., 2026; Liu, Wang, and Yu, 2025; Berger et al., 2024; Chen, Srinivasan, and Zakerinia, 2025). We complement this literature by opening the matching process before employment relationships are formed. By jointly observing postings, applications, callbacks, and search behavior, we trace how LLM-related demand shocks reach job seekers and show that applicant-side adjustment is constrained by limited mobility, preventing workers from fully offsetting adverse demand-side shocks.

Second, we highlight working-time amenities as an underexplored margin of adjustment to the LLM shock. Posting-based studies of AI and labor demand have primarily focused on vacancies, skill requirements, wages, and firm hiring demand (Acemoglu et al., 2022; Alekseeva et al., 2021; Liu, Wang, and Yu, 2025; Berger et al., 2024; Chen, Srinivasan, and Zakerinia, 2025), while studies using administrative labor-market records mainly capture realized employment, earnings, or formal hours (Jiang et al., 2025). We show that time-related amenities in job postings deteriorate as a distinct margin of job quality. This finding suggests that standard labor-market outcomes such as wages may understate the welfare costs of LLM shocks, especially in settings where informal overtime and weak enforcement make working-time conditions difficult to measure in administrative data.

Third, we document multidimensional distributional incidence. Existing work often ranks workers or occupations by exposure, adoption, or productivity gains (Eloundou et al., 2024; Bick, Blandin, and Deming, 2026; Brynjolfsson et al., 2025; Humlum and Vestergaard, 2025). We instead show that a single technological shock can impose qualitatively different costs on different worker groups. Junior workers are penalized through wages and callbacks, while senior and highly educated workers are penalized through time-related amenities. This finding implies that the distributional consequences of LLM shocks cannot be fully captured by a single outcome such as

employment or wages. It also helps reconcile why some groups may appear protected in conventional outcomes while still experiencing deterioration in non-wage job quality.

Fourth, we provide policy-relevant evidence that labor protection mitigates job-quality deterioration. Existing studies have largely interpreted the labor-market effects of LLMs through technology, task exposure, or firm adoption. Our evidence suggests that the adverse consequences of LLM shocks are institutionally mediated rather than technologically determined alone. Stronger labor protection weakens the deterioration of wages and time-related amenities and reduces applicants' exposure to lower-quality jobs. This finding points to the importance of labor-market institutions in shaping the distributional and welfare consequences of technological change.

Overall, our findings show that the LLM shock reshapes the pre-match labor market by jointly affecting employer demand, job quality, applicant search behavior, and employer–applicant interaction. High-LLM-exposure occupations experience fewer postings, lower offered wages, and worse time-related amenities. Employers respond by raising human-capital requirements, and applicants with stronger human-capital signals face smaller callback penalties. However, the burden of the shock is transmitted unevenly: junior workers are affected mainly through wages and callbacks, while senior and highly educated workers are affected more through time amenities. Applicant-side adjustment is limited by mobility constraints and cannot fully offset the negative effects. Weaker protection amplifies job-quality deterioration and pushes applicants toward lower-wage jobs with poorer time amenities. These findings underscore the need for policies that safeguard non-monetary working conditions, especially time-related amenities, and that target different worker groups according to the specific margins through which they bear the costs of technological change.

References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics, 40*(S1), S293–S340. doi:10.1086/718327

Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. *Labour Economics, 71*, Article 102002. doi:10.1016/j.labeco.2021.102002

Berger, P. G., Cai, W., Qiu, L., & Shen, C. X. (2024). *Employer and employee responses to generative AI: Early evidence*. SSRN Working Paper. doi:10.2139/ssrn.4874061

Bick, A., Blandin, A., & Deming, D. J. (2026). The rapid adoption of generative AI. *Management Science*. Advance online publication. doi:10.1287/mnsc.2025.02523

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942. doi:10.1093/qje/qjae044

Chen, W. X., Srinivasan, S., & Zakerinia, S. (2025). *Displacement or complementarity? The labor market impact of generative AI* (Harvard Business School Working Paper No. 25-039). Harvard Business School.

Dell'Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Organization Science, 37*(2), 403–423. doi:10.1287/orsc.2025.21838

Eisfeldt, A. L., Schubert, G., Taska, B., & Zhang, M. B. (2026). Generative AI and firm values. *Journal of Finance*. Forthcoming.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science, 384*(6702), 1306–1308. doi:10.1126/science.adj0998

Felten, E. W., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal, 42*(12), 2195–2217. doi:10.1002/smj.3286

Hui, X., Reshef, O., & Zhou, L. (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science, 35*(6), 1977–1989. doi:10.1287/orsc.2023.18441

Humlum, A., & Vestergaard, E. (2025). *Large language models, small labor market effects* (NBER Working Paper No. 33777). National Bureau of Economic Research. doi:10.3386/w33777

Jiang, W., Park, J., Xiao, R., & Zhang, S. (2025). *AI and the extended workday: Productivity, contracting efficiency, and distribution of rents* (NBER Working Paper No. 33536). National Bureau of Economic Research. doi:10.3386/w33536

Liu, Y., Wang, H., & Yu, S. (2025). *Labor demand in the age of generative AI: Early evidence from the U.S. job posting data* (World Bank Policy Research Working Paper No. 11263). World Bank.

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187–192. doi:10.1126/science.adh2586

Zhang, D., Yu, H., Li, L., Hu, J., Mo, Y., & Li, H. (2025). 中国人工智能技术暴露度的测算及其对劳动需求的影响——基于大语言模型的新证据 [The measurement of AI exposure and its impact on labor demand in China: Evidence from large language models]. *管理世界 [Management World], 41*(7), 59–75.