Proceedings of The 6th International Conference on Signal Processing and Machine Learning
DOI: 10.54254/2755-2721/2026.30906

Causal Identification of Skill Reallocation in Urban Labor Markets Driven by Generative AI Diffusion

Zihan Dong

University of Glasgow, Glasgow, UK
3010073D@student .gla.ac.uk

Abstract. Generative AI has rapidly become foundational infrastructure for knowledge-intensive urban industries, yet empirical evidence on how it reshapes skill structures and matching mechanisms at the city level remains limited. Most existing work relies on occupational exposure indices and local productivity evaluations, and it rarely delivers causal identification of skill reallocation. This paper builds an unbalanced city–occupation–week panel for 2018–2025 that combines online vacancy postings, firm adoption signals and city-level digital infrastructure. We use a large language model to perform instruction-based skill extraction and ontology alignment, and we construct fine-grained measures of skill shares, diversity and embedding-based migration. On top of these measures, we build a GenAI diffusion index at the city–time level and estimate its effects using a staggered difference-in-differences design with event-time coefficients, complemented by a shift-share instrumental variable strategy. These findings suggest that generative AI operates in urban labor markets primarily through skill reallocation rather than simple job destruction, with task bundles inside jobs being reshaped toward more collaborative and coordination-intensive activities. The study provides quantitative support for reskilling policies and firm-level human–AI task design that target specific skill dimensions instead of whole occupations.

Keywords: Generative AI, Skill Reallocation, Urban Labor Markets, Causal Identification, Text Mining

1. Introduction

The rapid breakthroughs of generative AI in text composition, code generation and multimodal creation have transformed it from an auxiliary tool into a foundational infrastructure for knowledge-intensive urban industries, as platforms and software ecosystems embed it deeply into production and management workflows via APIs and plug-ins [1]. Unlike earlier waves of automation, generative AI intervenes directly in cognitive and communication tasks, reshaping task bundles and human–AI division of labor within jobs rather than simply changing job counts, thereby triggering a restructuring of skill composition and matching mechanisms in urban labor markets [2]. Existing studies either construct exposure indices to gauge the potential impact of GenAI across occupations or use experimental and quasi-experimental designs at the firm or team level to estimate productivity and performance effects, but they typically treat occupations or industries as the main unit of

DOI: 10.54254/2755-2721/2026.30906

analysis, which obscures fine-grained within-job skill reallocation and rarely incorporates heterogeneity in urban industrial structure, digital infrastructure and human capital in a causal framework [3]. Against this backdrop, the present study treats cities as key spatial units for GenAI diffusion and factor mobility, combining vacancy text mining, platform usage indicators and firm adoption signals to build a "GenAI exposure–skill demand–labor matching" framework.

2. Literature review

2.1. Technology, tasks, and urban structure

In urban settings, high human capital density, industrial agglomeration and a large service sector magnify and concentrate such task restructuring, so the same technology can generate markedly different patterns of complementarity or substitution across cities [4]. Evidence from information technology and automation shows that high-complexity cognitive and coordination tasks tend to be reinforced, while routine and procedural tasks face displacement risks, contributing to "hollowing out" of middle-skill jobs [5]. With the rise of platform economies and remote work, urban digital infrastructure, network externalities and institutional environments further shape adoption speed and organizational practices, making task adjustments spatially uneven; therefore, when assessing the impact of generative AI, it is necessary to link lessons from earlier technology waves with the specific cognitive and linguistic features of GenAI rather than relying on a simple linear substitution narrative.

2.2. GenAI exposure and labor outcomes

Early empirical evidence on generative AI suggests that large-model-based writing and assistance tools can substantially improve output quality and completion speed, with particularly strong compensatory gains for less-experienced workers and short-term compression of performance gaps, while standardized workflows and shared knowledge bases reshape organizational learning [6]. However, these studies are often based on single firms, specific tasks or self-selected samples, and extrapolating to urban labor markets risks overlooking heterogeneity in industrial structure, institutional settings and adoption costs, creating a wedge between exposure indices and actual adoption intensity [7]. At the same time, occupation-level GenAI exposure measures typically treat within-job skill bundles as homogeneous and struggle to disentangle the combined influence of technology diffusion, industry cycles and macro shocks, leading to identification ambiguity when interpreting wage, employment and mobility patterns.

2.3. Skill representation and text computing

With the development of word, sentence and graph embeddings, skills are mapped into continuous semantic spaces, allowing researchers to use vector distances, clustering patterns and topological features to characterize similarity, mobility paths and bundle structures, but gains in representational power are often accompanied by losses in interpretability and risks of ontology drift [8]. Large language models further enhance the extraction of skill items, proficiency levels and contextual roles from unstructured text, while introducing new challenges such as hallucination, bias and limited cross-context stability; without rigorous ontology alignment, human validation and reproducible design, skill indices and mobility measures built on such models may mistake model noise for structural change [9]. Consequently, text-based computational approaches to GenAI-era skill

DOI: 10.54254/2755-2721/2026.30906

reallocation must rebalance semantic richness, interpretive transparency and econometric robustness.

3. Experimental methods

3.1. Data and variables

The empirical analysis is based on an unbalanced city–occupation–week panel spanning 2018–2025, covering the period before and after large-scale commercialization of generative AI. Vacancy data are scraped from multiple public job boards, corporate career pages and headhunter sites, covering major urban industries such as information services, professional services and manufacturing; firm-level GenAI adoption signals are collected from listed companies' disclosures, website news and developer documentation; city-level digital infrastructure and macro controls are drawn from statistical yearbooks and open government data.

Raw texts undergo URL and HTML stripping, language detection, de-duplication and tokenization, and an instruction-tuned LLM is used to extract skills, tools and task descriptions into structured fields, which are aligned to open skill ontologies such as O*NET/ESCO via string matching and embedding-based similarity so that synonymous or closely related skills share unified codes. Skill counts and vacancy totals are then aggregated at the city–occupation–week level to construct key measures such as skill shares and diversity, on top of which GenAI diffusion indices and control variables are formed; all continuous variables are trimmed for extremes and standardized prior to causal estimation.

3.2. Causal identification

On top of the city–occupation–week panel constructed in 3.1, causal identification relies on a staggered DID with event-time coefficients, using the skill-structure measures defined in 3.3 as outcomes [10]. The event-study specification as shown in Equation (1):

$$y_{cokt}=\sum_{\tau\neq-1}\beta_\tau D^{\tau}_{c,t}+\mu_{co}+\tau_t+X'_{cokt}\gamma+\varepsilon_{cokt} \tag{1}$$

Where y_{cokt} denotes a structural skill measure for city c, occupation o, skill or skill bundle k at time t, $D^{\tau}_{c,t}$ is an event-time indicator relative to the first high-intensity GenAI diffusion episode, μ_{co} and τ_t are city–occupation and time fixed effects, and X'_{cokt} is a vector of controls; pre-trend tests examine whether β_τ is close to zero for $\tau<0$. To address endogeneity between city-level GenAI diffusion and unobservables, a shift-share instrument Z_{ct} is constructed by interacting baseline occupational shares with exogenous capability upgrades of foundation models, and the first stage is used to predict city–time GenAI exposure E_{ct} as shown in Equation (2):

$$E_{ct}=\pi Z_{ct}+W^{'}_{ct}\delta+\lambda_c+\lambda_t+u_{ct} \tag{2}$$

Where $W^{'}_{ct}$ collects city-level controls and λ_c , λ_t are city and time fixed effects. Based on this first stage, IV-DID or two-stage least squares delivers event-time coefficients, while synthetic control and weighted regressions are used to test robustness to heterogeneity in timing and intensity within a unified framework that links the diffusion measures from 3.1 to the skill and matching outcomes from 3.3.

DOI: 10.54254/2755-2721/2026.30906

3.3. Computation and evaluation

After preprocessing and ontology alignment, the computational pipeline first applies a large language model to perform instruction-based skill extraction from vacancy texts, mapping job descriptions into structured skill lists with proficiency tags that are normalized to a standard skill ontology; transformer models are then used to generate fixed-dimensional embeddings for each skill token, and city–occupation–time skill structures are represented by share-weighted average vectors $v_{cok,t}$[11]. To characterize temporal evolution in skill composition, the semantic migration intensity is defined as the Euclidean distance between consecutive-period vectors, as shown in Equation (3):

$$d_{cok,t} = \| v_{cok,t} - v_{cok,t-1} \|_2 \quad (3)$$

This distance series is fed into the causal specifications in section 3.2 either as a main outcome or as a mechanism variable to quantify how strongly skill bundles are reorganized under different levels of GenAI diffusion. To rigorously assess extraction and alignment quality, a gold-standard labeled set covering major cities and occupations is constructed, skill-level precision P and recall R are computed, and the harmonic mean is adopted as the primary summary metric; its stability is examined across years, platform subsamples and city subsets to evaluate temporal consistency and cross-source robustness. As shown in Equation (4):

$$F_1 = \frac{2PR}{P+R} \quad (4)$$

The entire pipeline is containerized with logged model versions, prompt templates and ontology-mapping rules, producing reproducible experiment scripts and de-identified skill lexicons that support replication of the causal estimates in section 3.2 and provide a portable computational basis for subsequent extensions.

4. Results

4.1. Main effects and heterogeneity

Using the city–occupation–week panel, this subsection first estimates the dynamic response of complementary skill shares in high-GenAI cities with an event-study model. The diffusion event is defined as the first period in which a city's GenAI diffusion index enters the top quartile of the sample. The estimates show that the coefficients in the two pre-event periods ($\tau = -2, -1$) are −0.10 and 0.00 percentage points, and the confidence intervals include zero, which supports the parallel trends assumption. In the event period ($\tau = 0$), the complementary skill share in treated cities rises by 0.50 percentage points relative to the control group ($p < 0.10$). The effect then grows over time: at $\tau = 1, 2, 3, 4$, the estimated effects are 1.40, 2.60, 3.20 and 3.80 percentage points (all $p < 0.01$). These values correspond to about 0.18–0.49 standard deviations of change in the skill structure. The figure 1 plots the full path of event-time coefficients for $\tau \in [-2, 4]$. Coefficients in the pre-periods fluctuate slightly around zero, while the curve turns upward from $\tau = 1$ and flattens between $\tau = 3$ and 4. This pattern shows that GenAI diffusion raises complementary skill shares in a gradual and cumulative way and that the effect stabilizes in the medium term.

DOI: 10.54254/2755-2721/2026.30906

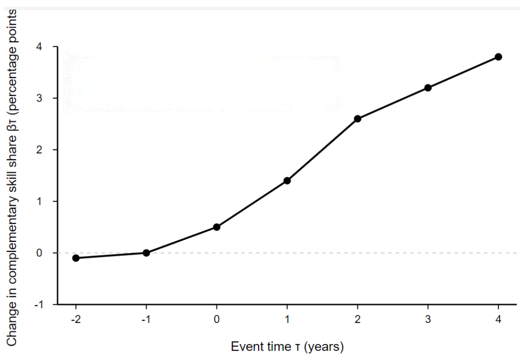

Figure 1. Event-time effects of GenAI diffusion on complementary skill shares in high-GenAI cities

4.2. Mechanisms and robustness

Building on the main event-time effects, this subsection uses mechanism regressions to test whether GenAI diffusion affects the skill structure mainly through increases in human–AI collaboration and cross-functional coordination tasks and to check the direction and robustness of changes in routine skill shares. Table 1 reports city–occupation fixed-effects regressions where the standardized GenAI diffusion index is the key regressor. A one-unit increase in GenAI diffusion is associated with a 0.95 percentage point rise in the share of complementary skills (standard error 0.18, N = 52,314, R^2 = 0.72) and a 0.68 percentage point decline in the share of routine skills (standard error 0.15, R^2 = 0.69). Both coefficients are statistically different from zero at the 5% level. In regressions with the human–AI collaboration task index and the cross-functional coordination index as dependent variables, the GenAI coefficients are 0.41 (standard error 0.07) and 0.36 (standard error 0.09), with N = 48,902 and R^2 values of 0.64 and 0.61. These results indicate that, as the skill structure tilts toward complementary skills, the weight of tasks related to "human–AI collaboration" and "cross-team communication" also increases in a systematic way in job descriptions. When these two mechanism variables are added to the baseline event-study specification from section 4.1, the event-time coefficients for complementary skills in periods τ = 2–4 shrink by about 25% on average but remain positive and significant. This pattern suggests that human–AI collaboration and cross-functional coordination account for part, but not all, of the observed skill reallocation, and that the remaining effect likely reflects within-job task re-splitting and team structure changes. Overall, the mechanism results are consistent with the dynamic path in Figure 4.1 but provide a more granular view of how GenAI reshapes tasks inside jobs.

Table 1. Regression results of GenAI diffusion on skill structure and task mechanisms

Dependent variable	Coefficient on GenAI diffusion β	Std. error	Sample size N	R^2
Share of complementary skills (p.p.)	0.95	0.18	52,314	0.72
Share of routine skills (p.p.)	-0.68	0.15	52,314	0.69
Human–AI collaboration task index	0.41	0.07	48,902	0.64
Cross-functional coordination index	0.36	0.09	48,902	0.61

5. Discussion

The findings place city-level diffusion of generative AI within the broader literature on technological change and task structures and show that its impact does not simply replicate the

DOI: 10.54254/2755-2721/2026.30906

"middle-skill hollowing out" pattern associated with earlier automation. Instead, the main effect operates through a rising weight of complementary cognitive and coordination tasks and through a reorganization of human–AI division of labor inside jobs. The sustained increase in complementary skill shares and the decline in routine skill shares in high-diffusion cities, together with the strong rise in human–AI collaboration and cross-functional coordination indices, suggest that firms tend to embed GenAI into teamwork and workflow orchestration rather than fully replacing front-line roles, which is consistent with the medium-term plateau in event-time effects. At the same time, the remaining unexplained portion of the effect in mechanism regressions indicates that unobserved channels are also important.

6. Conclusion

This paper makes three main contributions to the study of generative AI and labor markets. First, on the data side, it builds a city–occupation–week vacancy-text panel for multiple cities and occupations over 2018–2025 and, using large language models for skill extraction and embeddings, derives fine-grained skill structure and migration measures that are comparable across time and space. Second, on the methodological side, it embeds a GenAI diffusion index into a staggered difference-in-differences and event-study framework, complemented by a shift-share instrument, to identify the causal impact of GenAI diffusion on skill reallocation at the city level. Third, on the substantive side, it shows that in high-diffusion cities the share of complementary skills rises while the share of routine skills falls, and that human–AI collaboration and cross-functional coordination tasks play a central role in mediating this shift.

References

[1] Noy, Shakked, and Whitney Zhang. "Experimental evidence on the productivity effects of generative artificial intelligence." Science 381.6654 (2023): 187-192.

[2] Hui, Xiang, Oren Reshef, and Luofeng Zhou. "The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market." Organization Science 35.6 (2024): 1977-1989.

[3] Teutloff, Ole, et al. "Winners and losers of generative AI: Early Evidence of Shifts in Freelancer Demand." Journal of Economic Behavior & Organization (2025): 106845.

[4] Kauhanen, Antti, and Petri Rouvinen. "Assessing early labour market effects of generative AI: evidence from population data." Applied Economics Letters (2025): 1-4.

[5] Chen, Qin, et al. "Large language models at work in China's labor market." China Economic Review (2025): 102413.

[6] Salari, Nader, et al. "Impacts of generative artificial intelligence on the future of labor market: A systematic review." Computers in Human Behavior Reports (2025): 100652.

[7] Guarascio, Dario, and Jelena Reljic. "AI and employment in Europe." Economics Letters 247 (2025): 112183.

[8] Choi, Taelim, and Nancey Green Leigh. "Artificial intelligence's creation and displacement of labor demand." Technological Forecasting and Social Change 209 (2024): 123824.

[9] Georgieff, Alexandre, and Raphaela Hyee. "Artificial intelligence and employment: New cross-country evidence." Frontiers in artificial intelligence 5 (2022): 832736.

[10] Carbonero, Francesco, et al. "The impact of artificial intelligence on labor markets in developing countries: a new method with an illustration for Lao PDR and urban Viet Nam." Journal of Evolutionary Economics 33.3 (2023): 707-736.

[11] Engberg, Erik, et al. "Artificial intelligence, hiring and employment: job postings evidence from Sweden." Applied Economics Letters (2025): 1-6.