Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Uneven Generative AI Exposure and Skill Premiums in New Job Opportunities: Early Evidence from Beijing

Xiliu He
Tsinghua University
Haoxiang Zhao
Tsinghua University
Mingyi Ma
Tsinghua University
Edward Wenchuan Lai
Tsinghua University
Koei Enomoto
Tsinghua University
Anni Hu
Tsinghua University
Jiatong Li
Tsinghua University
Lingyun Chu
2077AI
Yuan Lai
yuanlai@mail.tsinghua.edu.cn

Tsinghua University

Article

Keywords:

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11606065/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Uneven Generative AI Exposure and Skill Premiums in New Job Opportunities: Early Evidence from Beijing

Xiliu He[1,†], Haoxiang Zhao[1,†], Mingyi Ma[1], Edward Wenchuan Lai[1], Koei Enomoto[1], Anni Hu[1], Jiatong Li[1], Lingyun Chu[2], Yuan Lai[1,3,*]

1. School of Architecture, Tsinghua University, Beijing, China.
2. 2077AI, Palo Alto, CA, USA.
3. Technology Innovation Center for Smart Human Settlements and Spatial Planning & Governance, Ministry of Natural Resources, Tsinghua University, Beijing, China.

† These authors contributed equally: Xiliu He, Haoxiang Zhao.

* Corresponding author: Yuan Lai. Email: yuanlai@mail.tsinghua.edu.cn

Abstract

Generative artificial intelligence (GenAI) is the first automation wave to reach high-cognitive tasks at scale, yet its effects on intra-urban labor markets remain unknown. Using 5.7 million job postings from Beijing, we build a neighborhood-level GenAI exposure index via five large language models and examine its spatial, structural, and causal mechanisms. GenAI exposure concentrates in core districts: the ChatGPT shock did not disperse high-skill job opportunities but unfolded within, and transmitted through, Beijing's existing core–periphery hierarchy. High-exposure neighborhoods attracted a rising share of high-skilled workers while advertised wages stagnated—a "high-skill trap" driven by task de-skilling and labor-market crowding. Difference-in-differences estimates support a causal interpretation: each 0.1-point higher neighborhood GenAI exposure corresponds to an additional 10.9% decline in advertised wages after the shock. These findings extend skill-biased technological change into the GenAI era and inform inclusive AI governance in global technology hubs.

Introduction

Generative Artificial Intelligence (GenAI) is reshaping the labor market unevenly[1-3]. The International Monetary Fund (IMF) estimates that AI will affect nearly 40% of global jobs, with the International Labor Organization (ILO) highlighting cognitive and clerical tasks as the most exposed[4,5]. Yet, it remains unknown how this macro-level shock unfolds within micro-level intra-urban labor markets, particularly regarding the new job opportunities. As the first automation wave to penetrate cognitive, creative, and communicative tasks that have long commanded an wage premium, GenAI raises a fundamental question. Will it act as a great leveler[6,7] that disperses new job opportunities to urban peripheries[8-10], fuel further agglomeration in urban cores and thereby deepen core–periphery divides and widen wage gaps[11-13], or drive both simultaneously? The answer has direct implications for intra-urban digital divides, labor protections in exposed occupations, and the governance of global tech hubs[3,14,15].

Previous waves of technology (Supplementary Text S1)—from computers to industrial robots—rewarded cognitive skills and reinforced agglomeration, creating a spatial polarization in which educated talent, high-value jobs, and rising wages clustered in urban cores[12,16-21]. GenAI raises the prospect of a departure from this pattern. Its capabilities span the full skill spectrum—from routine coding and writing to complex data analysis—and it may compress the very skill premium that high-skilled workers have historically commanded[22]. Three frameworks compete. The augmentation view holds that GenAI raises productivity across the skill distribution and creates new demand[23-26]. The substitution view warns of structural displacement and deepening inequality[1,27,28]. The task-reorganization view emphasizes that technology reshapes the content of work without wholesale replacement[3,21,29]. Our analysis points to a more counterintuitive scenario, which we call a “high-skill trap”: a surge of demand for high-skilled workers in exposed occupations coincides with advertised stagnant wages, as GenAI erodes the scarcity value of premium cognitive tasks[1,2,23,30] (Fig. 1).

Three critical research gaps persist. First, existing studies are predominantly national in scale, cross-sectional, and focused on high-income Western economies; they cannot reveal how GenAI exposure interacts with intra-urban core–periphery dynamics at fine spatial resolution[3,13,31]. Second, the literature lacks dynamic evidence linking exposure to synchronous changes in skill requirements and wages for new job openings, as well as causal evidence that exploits clear exogenous breaks—few studies leverage natural experiments such as the release of

ChatGPT[1,2]. Third, occupational exposure is typically measured using a single LLM or expert-elicitation framework[1,5,32,33], which embeds model-specific bias. These gaps are especially consequential for global tech-hub cities, which concentrate frontier AI capital alongside heterogeneous labor markets where the stakes for distribution are highest[34,35].

We address these gaps using 5,703,995 online job postings from Beijing (2018–2024), matched at the neighborhood level with multi-source geospatial data (Fig. 2). Beijing offers a compelling setting: it hosts a globally significant AI ecosystem (Baidu, ByteDance, leading research institutions) while exhibiting a pronounced dual structure of high-tech cores and low-skill peripheries[36-39]. To measure exposure, we designed a matching pipeline that used retrieval-augmented generation (RAG) to link job tasks to AI capabilities and then aggregated assessments from five leading LLMs to reduce model-specific bias. Our analysis yields three interconnected findings. First, GenAI exposure for new job openings is spatially locked into urban cores: the ChatGPT shock did not create a new digital divide but unfolded within, and transmitted through, Beijing's existing spatial hierarchy. Second, during the period, persistently hiring neighborhoods with high exposure experience both rising demand for high-skill workers and stagnant advertised wages—a high-skill trap. Third, this decoupling is associated with both de-skilling and labor-market crowding. To establish causality, we exploit the quasi-exogenous shock of ChatGPT's release, using pre-determined 2018 exposure as the treatment in a difference-in-differences (DID) design, with pre-trends validated through an event study.

This study makes four contributions. On measurement, we provide a fine-grained map of GenAI exposure at the neighborhood level, opening up an intra-urban perspective largely missing from existing work. On causal identification, we pair the quasi-exogenous shock of ChatGPT's release with pre-determined 2018 exposure to construct a rigorous DID design, supported by randomization inference and a Bartik shift-share instrumental variable (IV). On mechanisms, we jointly quantify the de-skilling and crowding channels that drive the wage penalty within a unified framework. On policy, we offer an empirical foundation for inclusive AI governance in global tech-hub cities.

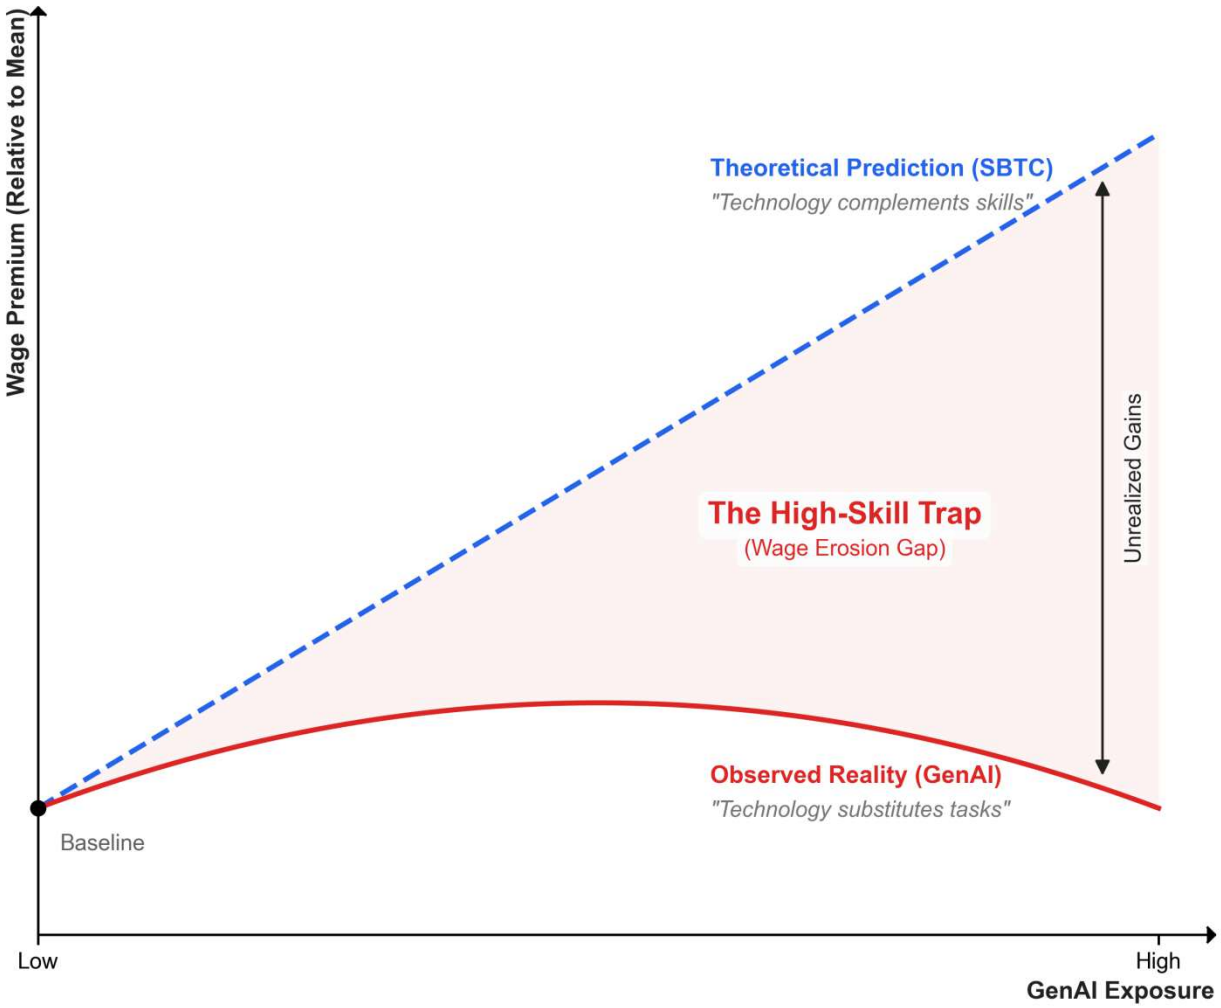

Fig. 1 | The "high-skill trap" under GenAI exposure.

The blue dashed line depicts the trajectory implied by skill-biased technological change (SBTC), extrapolated from earlier automation waves in which rising exposure to new technologies was accompanied by rising skill premia. The red solid line traces the pattern we observe in the GenAI era: after an initial rise, the skill premium flattens and ultimately stagnates. The shaded region marks the "high-skill trap"—the foregone wage gains that emerge because GenAI erodes the scarcity of high-skilled tasks even as exposure to them increases.

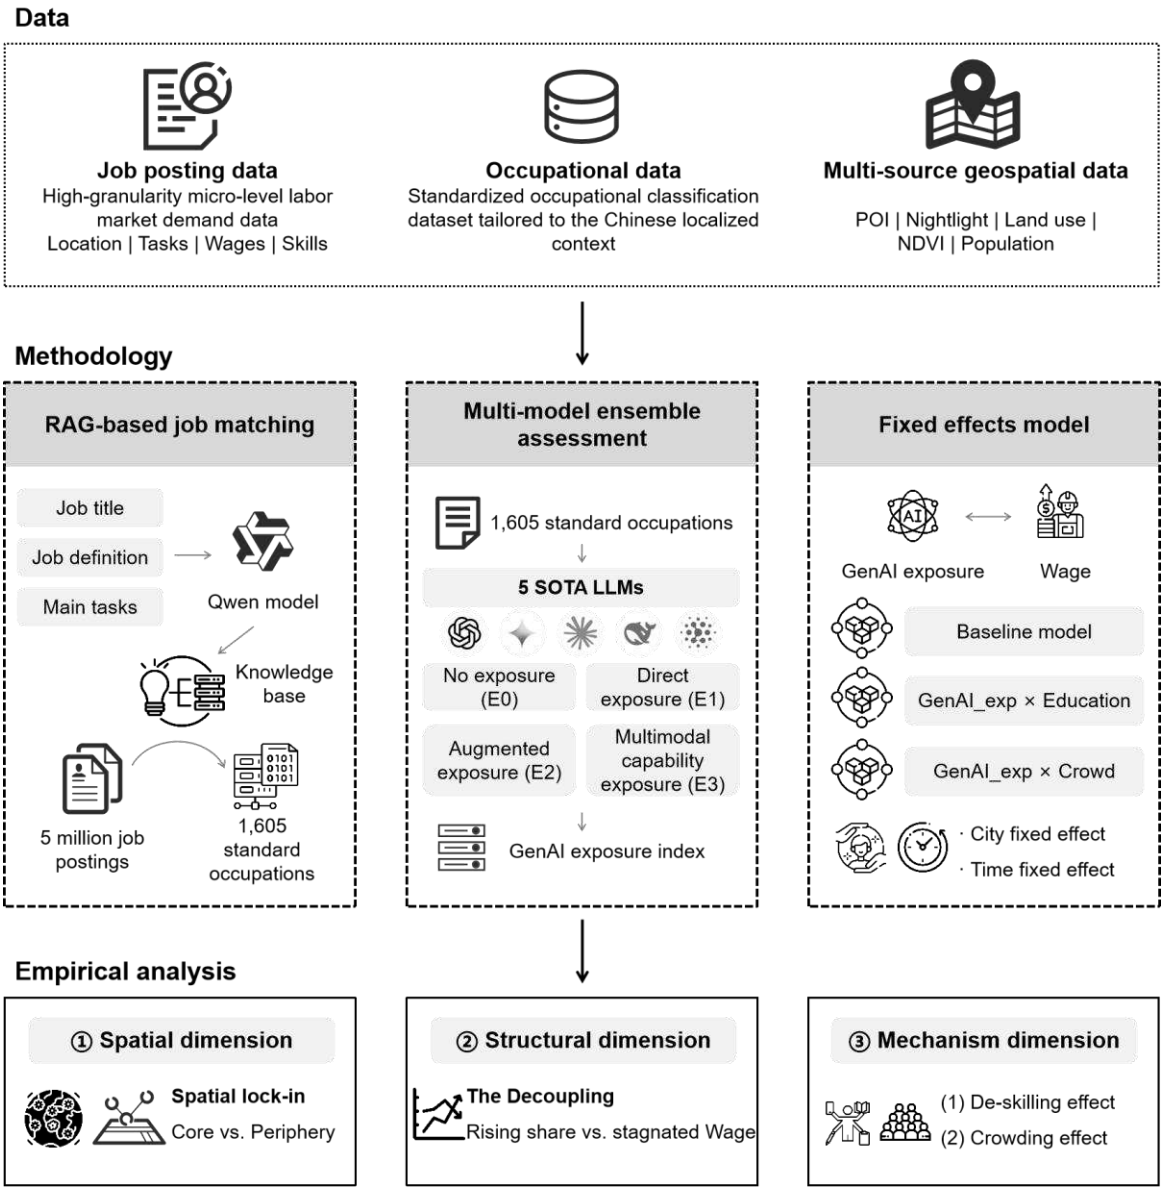

Fig. 2 | Measuring GenAI exposure and its spatial, structural, and mechanistic effects. The framework combines three steps: (1) millions of job postings are linked to standard occupational tasks using retrieval-augmented generation (RAG); (2) GenAI exposure scores are computed by aggregating assessments from five leading large language models; and (3) fixed-effects econometric models are used to isolate the distinct spatial, structural, and causal mechanisms of the GenAI shock.

Results

GenAI exposure is spatially locked into core districts

GenAI exposure does not spread evenly across the city. Contrary to the idea that new technologies dissolve geographic barriers[6,7], we find that for the new job opportunities reflected in our posting data, GenAI exposure reinforces existing spatial hierarchy in Beijing. Operating through agglomeration economies, the core–periphery gap in exposure—already substantial in 2018—persisted throughout the sample period. The ChatGPT shock thus did not create a new digital divide but unfolded within, and transmitted through, Beijing's existing spatial hierarchy of cognitive work.

The spatial pattern of GenAI exposure for new job opportunities (2018–2024) follows a monocentric, multi-ring structure with pronounced heterogeneity (Fig. 3a). Mean exposure declines monotonically from the inner rings outward—from roughly 0.33 inside the Fourth Ring Road to 0.21 beyond the Sixth Ring Road in 2024—and is strongly negatively correlated with distance from the urban center (core-to-periphery ratio ≈1.30, stable since 2021). Among new job openings, the highest exposure levels are locked into a golden triangle zone of knowledge-intensive services—Zhongguancun Science City (innovation), Financial Street (capital management) and the Guomao CBD (business services)—whose neighborhood-level means rank among the highest in the city. In contrast, low-exposure areas are scattered across the city's ecological-conservation zones and legacy industrial districts on the periphery, where low-skill services and traditional manufacturing dominate—tasks that AI currently struggles to replicate[40].

A temporal comparison confirms strong path dependence (Fig. 3b). First, core districts exhibit spatial lock-in and persistent relative advantage: although absolute exposure fell citywide after 2021, high-value clusters in Haidian and Chaoyang Districts retained their dominance in the exposure hierarchy, and the core – periphery gap actually widened as peripheral neighborhoods saw larger absolute declines. Second, within the market-based hiring captured by our data, we find no evidence of a post-ChatGPT spillover toward top-down planned sub-centers. The municipal sub-center (Tongzhou) recorded the highest neighborhood-mean exposure of any master-plan zone in 2018 (0.337, marginally above the core's 0.335), rose further by 2021 (0.347), and then declined steadily (0.292 in 2023, 0.283 in 2024), even as the core partially recovered—the post-shock period thus left the concentration of exposure in established cores intact rather than diffusing it toward planned growth poles.

A local spatial autocorrelation analysis (Local Indicators of Spatial Association, LISA) confirms this polarization (Fig. 3c). GenAI exposure is strongly and persistently spatially autocorrelated (global Moran's I = 0.19–0.23 across 2018–2024, all $p < 0.001$). High–high clusters remain locked into the core—79 of 118 significant high–high neighborhoods in 2024 lie within the Fourth Ring Road—while low–low clusters expanded outward, from 79 to 98 neighborhoods between 2018 and 2024 (90 of 93 located beyond the Fifth Ring Road), sprawling toward the southwest and northeast. The result is a spatially sticky low-exposure periphery where the local hiring market offers limited direct exposure and lacks a supportive environment for upgrading[41]. Transitional zones (high–low and low–high outliers) remain comparatively rare, accounting for roughly 5% of neighborhoods in each year and underscoring the sharpness of the core–periphery divide.

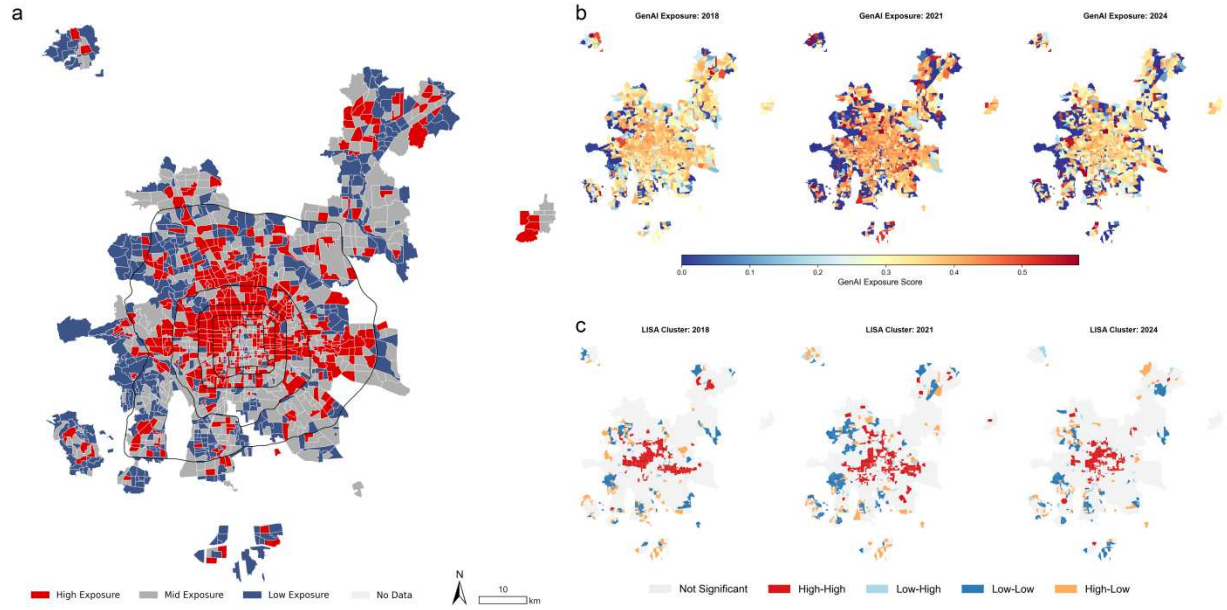

Fig. 3 | GenAI exposure across Beijing: spatial patterns, temporal change, and local clustering (2018–2024).

a, GenAI exposure intensity averaged over 2018-2024 across Beijing's neighborhoods. **b**, Exposure distribution in 2018, 2021 and 2024. **c**, Local spatial association (LISA) map: high–high clusters remain locked in the core with low–low clusters expanding across the periphery.

High-exposure neighborhoods attract talent but wages stagnate

Spatial polarization captures the geography of the shock; structural analysis illuminates how GenAI reshapes returns to human capital. Using micro-data from 2018–2024, we examine how GenAI exposure, skill demand, and wages evolve together, testing whether the shock has structurally reversed the skill premium. The results complement the conventional skill-biased technological change (SBTC) hypothesis[16] and reveal a "high-skill trap": although persistently hiring neighborhoods with high exposure continue to draw in highly educated workers, they have experienced marked advertised wage stagnation or even declines.

Fig. 4a shows the joint distribution of average years of schooling and standardized GenAI exposure across Beijing's neighborhoods in 2024. The upward-sloping fitted line confirms a significant positive correlation: exposure rises with education. A distinct high-exposure, high-education cluster emerges, where the workforce averages more than 14 years of schooling and GenAI exposure is at its highest—direct evidence that the GenAI shock is biased toward high-skill occupations[1]. At the other end, low-education neighborhoods exhibit limited GenAI integration: they are temporarily shielded from AI-driven displacement, but they also miss out on the potential productivity gains that AI augmentation could bring.

Yet the expected wage gains have not materialized. Fig. 4c tracks wage trajectories for groups with high, medium, and low exposure. Wages in the high-exposure group peaked in 2021 at ¥17,189 and subsequently declined. The descent accelerated after ChatGPT's release in late 2022, reaching ¥12,712 by 2024. Wages in low-exposure areas, while also peaking in 2021, exhibited substantially higher cumulative growth over the full period—consistent with a pattern in which GenAI compresses the wage distribution through both stagnation at the top and faster catch-up at the bottom[23]. Fig. 4b reinforces this picture: the relationship between initial 2018 exposure and subsequent six-year wage growth is strongly negative, indicating a clear wage penalty for high exposure. This finding revises the classical SBTC framework: in the age of GenAI, technological exposure no longer uniformly delivers a skill premium; instead, it can suppress wage growth through task-level reconfigurations[42].

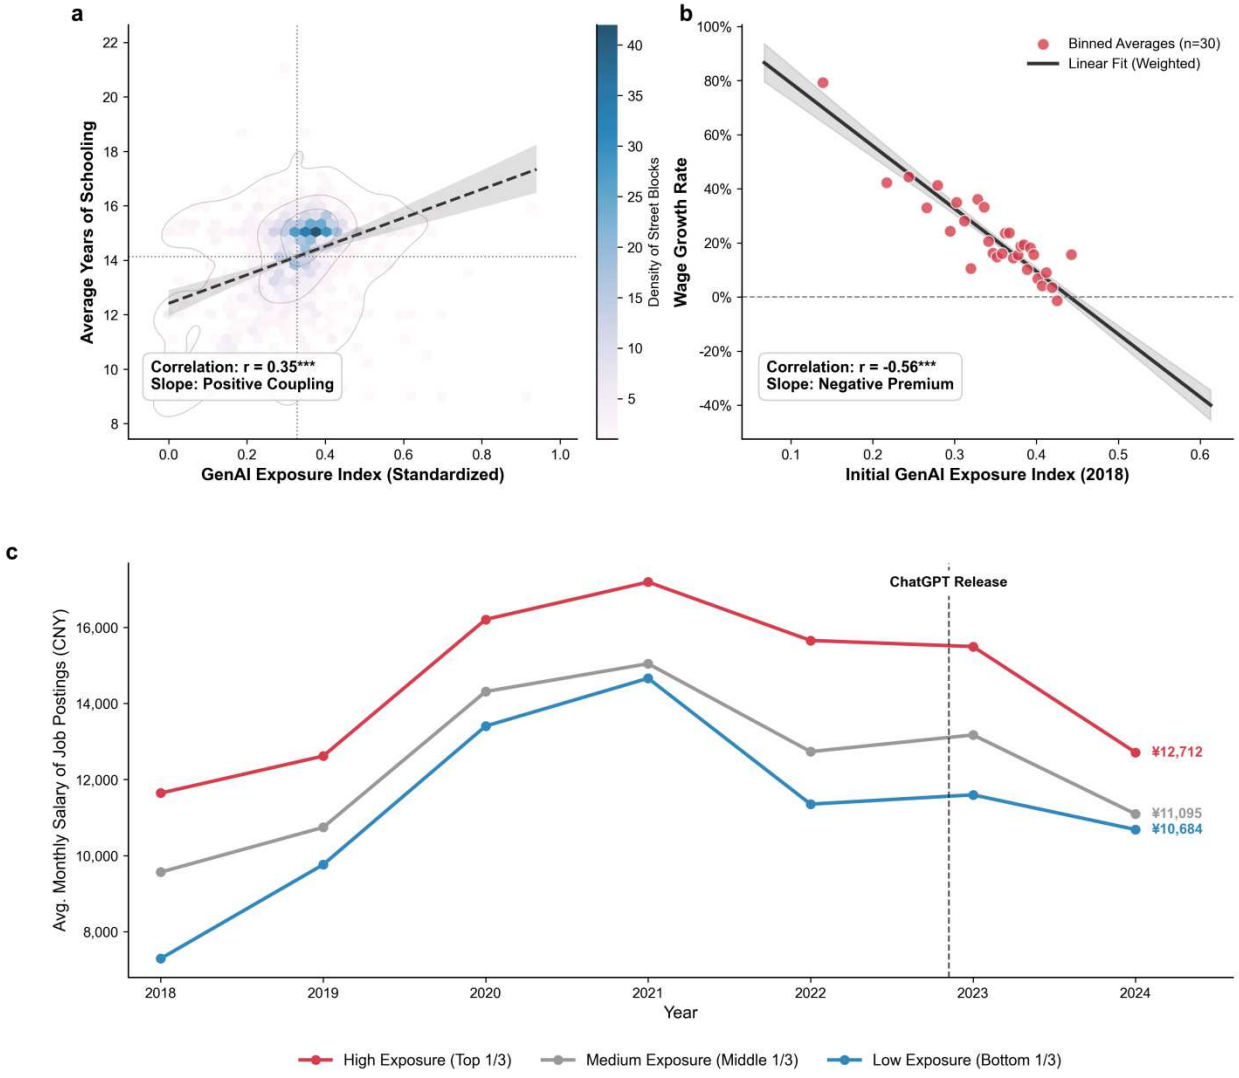

Fig. 4 | GenAI exposure, education, and the emerging wage penalty.
a, Hexagonal heatmap of standardized GenAI exposure versus average years of schooling in 2024. The dashed line is a linear fit (r = 0.35, $p < 0.001$). **b**, Binned scatter plot relating 2018 GenAI exposure to wage growth over 2018–2024. The negative slope (solid line) and 95% confidence interval (shaded) indicate a wage penalty for higher initial exposure. **c**, Average monthly wages (CNY) for high, medium, and low GenAI exposure groups. The dashed line marks ChatGPT's release; after that point, the high-exposure group (red) shows a marked downturn.

Why do neighborhoods that exhibit high demand for educated workers and concentrate high GenAI exposure still see wages stagnate? Fig. 5 further reveals a sharp contrast between low- and high-exposure areas. In low-exposure areas, the share of high-skill job opportunities and

advertised wage returns moved broadly together: the high-skill share of postings roughly doubled from 20% (2018) to 39% (2024), and the high-skill wage index climbed to 121 (2018 = 100)—the pattern that classical SBTC framework predicts. In high-exposure areas, however, the two decouple. The high-skill share rose from 50% in 2018 to a peak of 66% in 2021 and remained elevated at 61% in 2024, yet the high-skill wage index ended the period back at its 2018 baseline—about 22 points below the low-exposure counterpart. In level terms, the wage premium that high-exposure neighborhoods once commanded for high-skill postings narrowed from +26% in 2018 to roughly +3% in 2024. Employer demand for educated labor remained elevated; the returns did not follow.

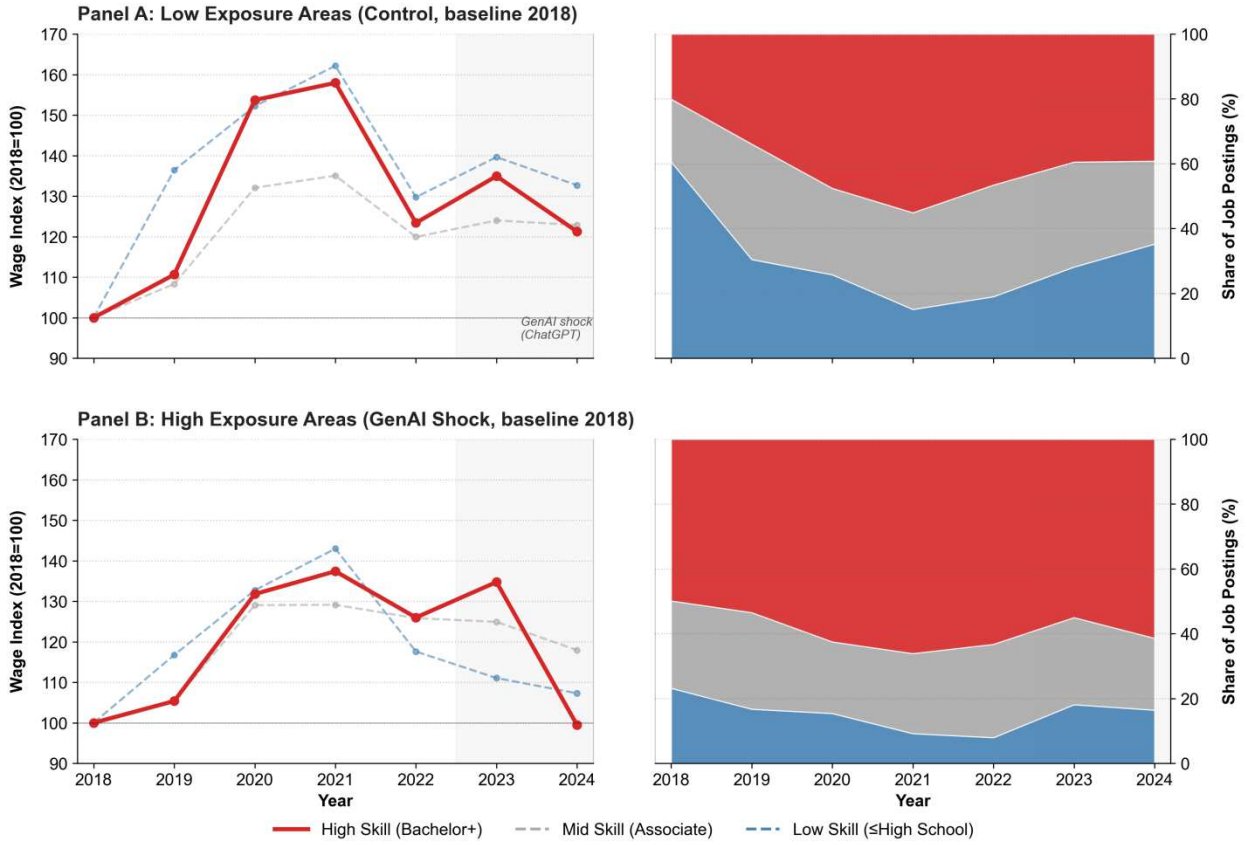

Fig. 5 | The decoupling of skill demand and wage returns in high-exposure neighborhoods.
Panel A (Low Exposure, control) and Panel B (High Exposure, treatment) are defined by 2018 baseline GenAI exposure terciles. The focal pattern is in Panel B: the high-skill share rises from roughly 50% to 61% while the high-skill wage index returns to its 2018 baseline, whereas in Panel A both measures rise together.

Causal identification of the GenAI wage penalty

The patterns shown so far are descriptive, but association is not causation. Between 2018 and 2024, Beijing's labor market faced overlapping shocks—antitrust tightening in the technology sector (from 2021), the COVID-19 pandemic and recovery (2020–2023), and a real-estate downturn (from 2022)—any of which could coincide spatially with GenAI exposure. To isolate the causal effect of GenAI, we use a DID design with a pre-determined treatment variable, validate the parallel-trends assumption with an event-study specification, and subject the estimate to an extensive battery of robustness checks—randomization inference, spatial fixed-effect specifications, alternative clustering structures, and weak-instrument-robust inference—summarized below and documented in full in Supplementary Text S5.

Two features of the research design support causal interpretation. First, the treatment is pre-determined: we use GenAI exposure in 2018, four years before ChatGPT's release (November 2022) and before our 2020–2024 analysis window, ruling out reverse causality. Second, the DID specification interacts this pre-determined exposure with an indicator for the post-treatment period ($\text{GenAI}_{2018} \times \text{Post}$, where Post = 1 in 2023–2024 and 2022 serves as the reference year), while controlling for pre-determined neighborhood exposure to the three concurrent shocks—technology-sector regulation, COVID recovery, and real-estate adjustment—each measured from the 2018 industry composition and likewise interacted with Post, so that their confounding influence is separated from the GenAI shock. The key identifying assumption is that, in the absence of the GenAI shock, high- and low-exposure neighborhoods would have followed parallel advertised wage trends.

Fig. 6 reports the event-study estimates with the concurrent-shock controls (2022 is the reference year; Equation 3). Pre-treatment coefficients (2020: $\beta = -0.164$, $P = 0.103$; 2021: $\beta = -0.074$, $P = 0.393$) are statistically insignificant, and a joint Wald test fails to reject the parallel trends assumption (Wald = 2.66, $P = 0.265$). A sharp structural break appears post-treatment: the coefficient for 2023 drops to −0.235 ($P = 0.008$), implying that a one-standard-deviation increase in pre-determined exposure is associated with an approximately 21.0% decline in wages during the first year of the shock. The 2024 coefficient is −0.242 ($P = 0.014$): rather than dissipating, the wage penalty persists through the second post-treatment year.

Table 1 reports the core DID estimates. The bare DID without additional regressors (Column 1) yields an interaction coefficient of −0.139 ($P = 0.020$). Adding the time-varying neighborhood controls (Column 2) leaves the estimate essentially unchanged at −0.140 ($P = 0.019$), indicating that the wage penalty is not an artefact of observable local characteristics. Adding the three concurrent-shock controls (Column 3) does not shrink the estimate either; it increases slightly to −0.160 ($P = 0.015$), equivalent to about 14.8% lower advertised monthly wages. Most simply, after the shock, advertised monthly wages are about 14.8% lower in neighborhoods with pre-shock GenAI exposure one standard deviation (0.14 GenAI exposure index points) above the average than in those one standard deviation below it—and every 0.1-point higher GenAI exposure index corresponds to roughly 10.9% lower advertised monthly wages. This stability is informative: if the concurrent shocks were driving the result, adding them as controls would have reduced the estimate; instead, the estimate remains robust, suggesting that any omitted-variable bias from those shocks is more likely to mask the GenAI effect than to create it spuriously[43].

The remaining columns of Table 1 probe the two threats to inference most frequently raised in neighborhood-level designs: unobserved spatially specific shocks and spatially correlated errors. Columns (4) and (5) re-estimate Column (3) with standard errors clustered at the administrative-district (G = 16) and ring-road-belt (G = 6) level; the point estimate is unchanged, and inference remains significant ($P < 0.001$ and, under conservative t(G−1) critical values given only six clusters, $P = 0.053$). Columns (6) and (7) replace the common year fixed effects with ring-belt-by-year and district-by-year fixed effects, which absorb all time-varying shocks common to a spatial unit—including unobserved post-2021 downturns in the golden triangle zone—so that identification relies solely on within-district-year variation in GenAI exposure across neighborhoods. The coefficient attenuates only modestly (−0.152, $P = 0.017$; −0.149, $P = 0.021$), indicating that the wage penalty is not an artefact of core-district labor-market trends.

Beyond Table 1, the estimate survives a broader battery of checks documented in Supplementary Text S5. Randomization inference—permuting the pre-determined exposure across neighborhoods 500 times—yields a placebo distribution tightly centered on zero, with the true estimate in the extreme left tail (permutation $P = 0.002$). Conley HAC standard errors with 10-km and 20-km cutoffs leave the estimate significant ($P = 0.038$ and $P = 0.002$), and residual

spatial autocorrelation is weak in all years (Moran's I ≤ 0.022). Finally, a Bartik shift-share instrument—the Beijing-wide, leave-one-out industry-level GenAI exposure weighted by each neighborhood's 2018 industrial composition, estimated in long differences—delivers a reduced-form estimate of −0.652 log points ($P < 0.001$). The first-stage F-statistic is 66.2, far exceeding the conventional weak-instrument threshold of 10 and confirming the relevance condition of the IV[44]. Weak-instrument-robust Anderson–Rubin inference decisively rejects a zero effect (AR F = 69.61, $P = 1.7 \times 10^{-16}$), providing corroborating evidence for the DID design.

Table 1 | Difference-in-Differences Estimates of the GenAI Wage Penalty

	(1) Bare DID	(2) + Controls	(3) + Confounder controls	(4) District cluster	(5) Ring cluster	(6) + Ring×Year FE	(7) + District×Year FE
Dependent variable: ln(Wage)							
$GenAI_{2018}$ × Post	−0.139** (0.060)	−0.140** (0.060)	−0.160** (0.065)	−0.160*** (0.030)	−0.160* (0.063)	−0.152** (0.063)	−0.149** (0.065)
TechReg × Post			0.070 (0.059)	0.070 (0.054)	0.070 (0.040)	0.068 (0.052)	0.076 (0.052)
COVID × Post			0.086 (0.085)	0.087 (0.070)	0.086 (0.060)	0.092 (0.076)	0.088 (0.077)
RealEstate × Post			−0.042 (0.106)	−0.042 (0.072)	−0.042 (0.021)	−0.045 (0.095)	−0.041 (0.094)
Controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Confounder controls	No	No	Yes	Yes	Yes	Yes	Yes
Entity FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	via Ring×Year	via District×Year
Ring×Year FE	No	No	No	No	No	Yes	No
District×Year FE	No	No	No	No	No	No	Yes
SE clustering	Neighborhood	Neighborhood	Neighborhood	District	Ring	Neighborhood	Neighborhood
Clusters (G)	1,379	1,379	1,379	16	6	1,379	1,379
Observations	6,895	6,895	6,895	6,895	6,895	6,895	6,895
Neighborhoods	1,379	1,379	1,379	1,379	1,379	1,379	1,379
R^2 (within)	0.0015	0.0293	0.0304	0.0304	0.0304	0.0520	0.0625

Notes. *The dependent variable is the natural logarithm of average monthly wages. $GenAI_{2018}$ denotes standardized pre-determined GenAI exposure measured in 2018. Post = 1 for 2023–2024. The analysis window spans 2020–2024 with 2022 as the base year. Columns (1)–(3) add control variables and confounder controls sequentially. Columns (4)–(5) replicate Column (3) with standard errors clustered at the administrative-district level (G = 16) and at the ring-road-belt level (G = 6, actual ring-road polygons), respectively; inference in these columns uses t(G−1) critical values given the small number of clusters. Columns (6)–(7) replicate Column (3) but replace the year fixed effects with ring-belt-by-year and district-by-year fixed effects, absorbing all time-varying shocks common to a spatial unit; identification relies on within-ring-year and within-district-year variation in GenAI exposure across neighborhoods. Confounder controls are industry-composition-based exposure measures for technology regulation, COVID recovery and real-estate adjustment, each interacted with Post. Control variables include ln(company count), ln(nightlight), NDVI, POI density and land-use ratio. Standard errors clustered at the neighborhood level are in parentheses unless otherwise noted. *** p < 0.01; ** P < 0.05; * P < 0.1.*

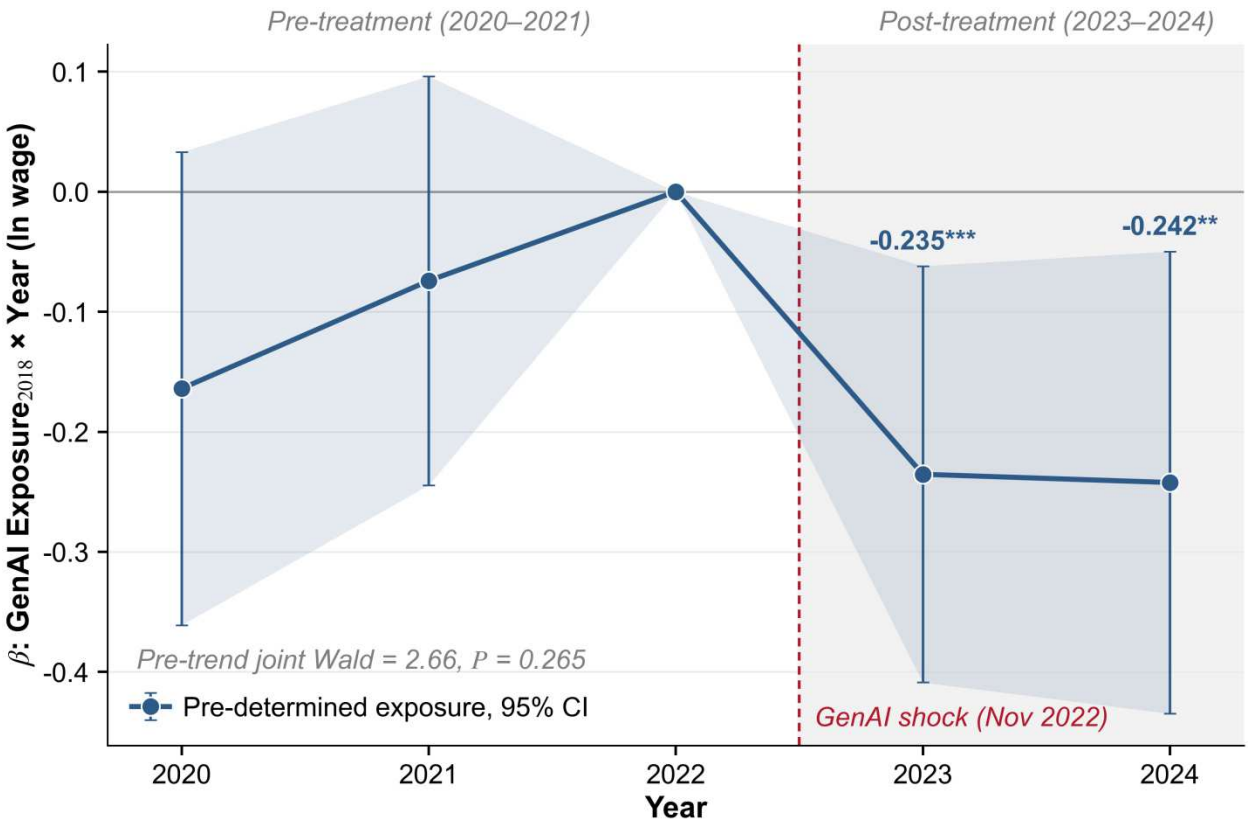

Fig. 6 | Causal identification of the GenAI wage penalty.

Event-study estimates of the dynamic effect of pre-determined (2018) GenAI exposure on log wages, controlling for concurrent shocks. The reference year is 2022. Green shading marks the pre-treatment period (2020 – 2021) and red shading marks the post-treatment period (2023 – 2024). Pre-treatment coefficients are statistically indistinguishable from zero (joint Wald: 2.66, $P = 0.265$), supporting the parallel-trends assumption. A sharp negative break emerges in 2023 and persists through 2024. Shaded bands around point estimates represent 95% confidence intervals based on neighborhood-clustered standard errors.

De-skilling and crowding erode the skill premium

Classical SBTC theory predicts that new technologies complement high-skilled workers and widen the skill premium[45,46]. Recent work on GenAI points to the opposite mechanism—de-skilling, or skill-levelling—in which large language models lower the cognitive barriers to tasks once reserved for educated workers (coding, copywriting, data analysis), allowing less-educated workers to substitute for high-skilled labor and compressing wages at the top[25,42,47,48]. A second, complementary channel operates through labor-market crowding: when the influx of talent in AI-exposed fields grows rapidly and demand for their tasks does not expand enough, productivity gains can paradoxically push wages down—a pattern often described as "involution"[49,50]. Viewed through the lens of labor demand, we capture the employer-side reflection of this dynamic using a measure of job-market heat, which reflects the intensity of competition for positions in each neighborhood. We measure education requirement as the mean required years of schooling across all job postings in a neighborhood-year, and job-market heat as the standardized recruitment intensity in that neighborhood-year.

Column 2 of Table 2 shows that, although the main effect of education on advertised wages remains significantly positive (1.455), the interaction between GenAI exposure and education is negative and highly significant (−1.287, $P < 0.001$). This means that the extra advertised wage gain associated with GenAI exposure shrinks as education requirements rise. Supplementary Figure S4 traces the implied marginal effect. For neighborhoods with below-average education requirements, GenAI exposure has a positive marginal effect—consistent with an assistive role in which AI helps workers clear skill-specific task barriers[23]. As education requirements increase, however, this effect declines and eventually turns negative; the marginal effect crosses zero at approximately 14.2 years of required schooling (0.54 standard deviations above the mean). For these high-education neighborhoods, GenAI's capabilities overlap substantially with the routine cognitive and textual tasks that define the jobs, and the technology acts more as a substitute than a complement, eroding relative bargaining power. These results provide strong support for the de-skilling hypothesis[2,25].

Skill substitution alone does not explain the full wage penalty. Column 3 of Table 2 reports a highly significant negative interaction between GenAI exposure and job-market heat (−

0.709, $P < 0.001$), indicating that greater competitive intensity strongly suppresses the GenAI-salary gradient. Supplementary Figure S5 traces the implied marginal effect. In low-competition niches (more than 0.45 standard deviations below mean heat), the marginal effect of GenAI exposure is positive, reflecting excess returns where competition remains limited. As recruitment intensity approaches and exceeds the mean, the premium shrinks and crosses zero, marking an equilibrium phase. In the most competitive segments, the marginal effect turns sharply negative, consistent with a crowding-driven wage penalty.

Together, these findings trace a technology-induced saturation mechanism: widespread GenAI adoption may lower barriers and homogenize the skill signals that once differentiated educated workers, intensifying competition for available positions in already hot recruitment markets, and compressing offered wages. The interaction models in Table 2 are conditional-correlation tests that help interpret the DID wage penalty; causal identification rests on the pre-determined DID design above. Supplementary Text S7 reports confounder and window robustness, channel-separation diagnostics with pre-determined-moderator checks, and placebo tests.

Table 2 | Mechanisms underlying the GenAI wage effect

	(1) Baseline	(2) De-skilling	(3) Crowding
Dependent variable: *ln*(Wage)			
GenAI Exposure (Std.)	0.108	0.690***	−0.318***
	(0.113)	(0.114)	(0.102)
Education (Std.)	3.196***	1.455***	2.909***
	(0.156)	(0.113)	(0.161)
Job Market Heat (Std.)	0.823***	0.791***	0.866***
	(0.120)	(0.125)	(0.168)
GenAI × Education		−1.287***	
		(0.079)	
GenAI × Job Market Heat			−0.709***
			(0.043)
ln(Company Count)	−0.072	−0.152***	−0.069
	(0.049)	(0.034)	(0.063)
ln(Nightlight)	0.209**	0.184**	0.176**
	(0.100)	(0.077)	(0.087)
NDVI (Std.)	0.052	−0.033	−0.049
	(0.097)	(0.065)	(0.143)

POI Density (Std.)	−0.015	−0.008	−0.028***
	(0.018)	(0.010)	(0.010)
Land Use Ratio (Std.)	0.178	0.095	0.176**
	(0.118)	(0.212)	(0.086)
Entity FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Observations	6,895	6,895	6,895
R^2 (within)	0.825	0.881	0.849
F-statistic	3,413.0	4,493.0	3,519.5

Notes. *The dependent variable is the natural logarithm of average monthly wages. All columns report two-way fixed-effects panel regressions estimated on the 2020–2024 window, aligned with the main DID analysis. Standard errors in parentheses are two-way clustered by neighborhood and year. *** P < 0.01; ** P < 0.05; * P < 0.1.*

Discussion

This study uses fine-grained data to reveal a triple polarization—spatial, structural and wage-related—that GenAI is producing in Beijing. Our first finding challenges the classic "death of distance" hypothesis, which predicts that technologies that slash the cost of moving information will disperse high-skilled work geographically[6]. In reality, GenAI-exposed jobs remain tightly clustered in the existing innovation zones, financial districts, the CBD, and this clustering persisted after ChatGPT's release rather than dispersing—the spatial hierarchy of AI exposure held its rank order even as the shock transmitted through it. We attribute this paradox to the changing nature of knowledge in the GenAI era. While codified information and routine analysis can now flow almost frictionlessly through large models, high-value creative work still relies on tacit knowledge—the kind of know-how that is exchanged through face-to-face interaction, informal social networks and localized innovation ecosystems[51-56]. GenAI, rather than dissolving the need for proximity, acts as a complementary asset that magnifies the returns to agglomeration, leaving peripheral neighborhoods that lack talents and supportive environment at persistent disadvantage[18,57,58].

More fundamentally, our results offer an empirical complement to SBTC in the GenAI era. The long-standing view that new technologies favor high-skilled workers and widen the skill premium[20,33] may not hold in the GenAI era. The magnitudes are substantial: each 0.1-point higher neighborhood exposure corresponds to an additional 10.9% decline in advertised wages

after the shock—14.8% per standard deviation (0.14 GenAI exposure index points), or roughly ¥1,050 per month at the sample mean advertised wage of ¥9,673. Our evidence identifies a de-skilling mechanism: GenAI exerts a negative effect on the advertised wages for roles requiring high education (in average-exposure neighborhoods, the education–wage gradient reaches zero at about 14.2 years of schooling), consistent with the levelling effect noted by others[23,59]. It simultaneously lowers the barriers to routine coding, writing, and analytical tasks, allowing employers to fill previously high-skilled roles with lower-credentialed candidates[60]. On the other hand, it directly encroaches on the core cognitive tasks traditionally embedded in high-skilled occupations, rapidly diluting their scarcity[23,61]. Wage stagnation, however, is not solely the product of substitution. It is amplified by labor-market crowding. In popular, AI-exposed fields, widespread adoption intensifies competition: as GenAI raises baseline output expectations, candidates face steeper competition and must invest more effort simply to secure a job offer[50]. External evidence points in the same direction: since ChatGPT's release, the job-finding rate of young workers aged 22-25 entering AI-exposed occupations has declined by roughly 14%[62]. When collective upskilling is not met by expanding demand for the newly accessible tasks, it fails to translate into productivity gains and instead weakens workers' bargaining power[42,63]. This demand-return decoupling echoes macro-level evidence: although AI-related skills may command a wage premium, employment in AI-vulnerable occupations is 3.6% lower five years on in regions with higher demand for AI skills[64].

The policy implications are clear. In responding to GenAI, policymakers should look beyond computing infrastructure or generic skills training toward inclusive AI governance. We highlight three directions. First, to loosen the monopoly of urban cores over talent, technology, and supportive environment, AI demonstration zones could be established in municipal sub-centers and suburban new towns, accompanied by lower access costs for digital infrastructure and support for building localized innovation ecosystems. Second, higher education should pivot away from cultivating codifiable standard cognitive skills toward value judgment, decision-making, complex collaboration, critical reasoning, and cross-cultural communication[65,66]—capabilities that are harder to automate—while embedding human-in-the-loop working modes that prepare graduates for human-AI collaborative problem-solving[67]. Third, policy should attend

to workers at both ends of the exposure distribution. For highly exposed occupations, re-employment transition funds, skill-transformation subsidies[68] and capability enhancement program[69] can cushion wage erosion and prevent displacement from cascading into talent outflows. For low-exposure workers, the risk is exclusion rather than displacement: basic AI and digital literacy programs[69], affordable access to AI tools[70], and sector-specific adoption support[71] can help these workers and industries use GenAI as a complement.

This study has limitations. The data come from Beijing alone; caution is needed when generalizing to other cities in China and worldwide. Online recruitment data capture new job opportunities—specifically, the flow of new hiring demand in the labor market—rather than AI-induced novel job roles or the existing stock of employed workers, and thus may over-represent volatility[72]. Our data also lack worker demographics such as age and tenure, so the youth-employment implications discussed above are inferred from neighborhood-level wage and hiring patterns rather than directly observed. Moreover, commercial online platforms primarily capture market-based hiring; public-sector recruitment, an important component of employment growth in policy-driven areas such as Beijing's municipal sub-center, proceeds through separate channels and is under-represented in our data. Finally, because GenAI has been widely deployed for only a short period, this study should be read as early evidence: our estimates capture the initial labor-market signals of the shock, while longer-run job-creation and productivity effects are not yet visible in the data. Tracking these evolutionary dynamics over longer horizons, and across a set of global technology hubs, is a priority for future work.

Methods

Study area selection

We selected Beijing as the study area for two reasons. First, it sits at the technological frontier. Beijing is a global hub for AI research and industry, concentrating China's leading AI talent, core research institutions (such as Tsinghua University and the Beijing Academy of Artificial Intelligence), and the headquarters of firms like Baidu and ByteDance. This high density of AI activity makes Beijing an early adopter of GenAI, and the resulting labor-market disruption offers an early signal of the dynamics that may soon appear in other leading AI cities—San Francisco, London, and Singapore among them[36,37,39]. Second, Beijing exhibits a

pronounced spatial duality. The city has a textbook core–periphery structure, combining world-class knowledge hubs such as Zhongguancun Science City with peripheral areas that retain traditional industrial or agricultural features[38]. This sharp spatial contrast provides an ideal setting to test whether GenAI functions as an "equalizer" that disperses opportunity to the periphery or as a force that deepens intra-urban divides[73,74].

Core recruitment indicators and control variables

We use online job postings from the RESSET Enterprise Big Data Platform, which aggregates vacancies from major Chinese recruitment platforms (51job, Zhaopin, BOSS Zhipin and Liepin). After cleaning—retaining only the first unique posting per company per month and dropping records with missing key fields—we occupation-matched the remaining postings and spatially assigned them to Beijing's neighborhoods (regulatory planning units). The final dataset contains 5,703,995 valid recruitment postings from 1 January 2018 to 31 December 2024. From these micro-data we constructed three core neighborhood-level measures. (1) Hiring activity. We compute the share of citywide postings located in each neighborhood (Job Share) to capture its relative position in the market, and a standardized log-transformed count (Job Index) to measure cross-neighborhood variation in recruitment intensity. (2) Average advertised monthly wage. All compensation components—annual salary, 13th-month pay, and bonuses—are converted to a monthly equivalent through a pre-built salary-standardization dictionary and averaged within each neighborhood; salary strings flagged as extreme outliers in the 0.1% tails of the raw salary distribution (monthly equivalents below ¥1,000 or above ¥67,500) are treated as missing before aggregation. (3) Average educational requirement. We map each posting's education requirement to years of schooling (unrestricted = 9, junior middle school or below = 9, high school/vocational/technical = 12, associate degree = 15, bachelor's = 16, master's = 19, doctorate = 23) and average these values within neighborhoods to capture the local skill threshold and degree of industrial sophistication.

To control for confounding influences, we include five neighborhood-level time-varying controls: firm presence, living convenience, development intensity, environmental quality, and economic activity. All controls are compiled as an annual panel for 2018–2024 at the

neighborhood scale. Full indicator definitions, calculation methods, and data sources appear in Supplementary Text S4.

RAG-based job matching method

Online job titles and descriptions are often unstructured and highly variable, so traditional keyword or fuzzy matching yields low accuracy. We therefore developed a standardized occupational matching framework based on RAG, achieving precise mapping between large-scale job postings and China's 2022 Occupational Classification Dictionary. First, we built a national-standard occupational knowledge base by extracting core fields—names, definitions and primary tasks—for all detailed occupations in the Dictionary and vectorizing this textual information with the Qwen3-Embedding-4B model, producing a high-dimensional semantic vector knowledge base of 1,605 standard occupations. Second, in the matching phase, the same embedding model converted job titles and responsibilities from online recruitment data into query vectors, which were then matched by semantic-similarity retrieval against the knowledge base, assigning each non-standardized posting to the closest standard occupational sub-category. To validate the method, we randomly sampled matched postings and had them checked by expert review; classification accuracy exceeded 93%, substantially outperforming traditional dictionary and fuzzy-matching approaches. Compared with end-to-end LLM-based generative classification, the RAG pipeline achieves similarly high accuracy at a fraction of the computational cost.

Multi-model ensemble assessment of GenAI exposure

Using the same 1,605 standard occupations defined above (excluding party/government organs and military personnel), we constructed an assessment matrix using five leading LLMs—ChatGPT-4o, Gemini-2.5-Pro, Claude-3.5-Sonnet, GLM-4 and DeepSeek-V3—to measure the GenAI shock and to mitigate the training-data biases and hallucinations inherent in any single model. Following the assessment criteria of Eloundou et al. (2024)[1] and Chen et al. (2025)[75], we designed standardized prompt engineering requiring each model to independently classify each occupation as exhibiting No Exposure (E0), Direct Exposure (E1), Application-Augmented Exposure (E2) or Multimodal Capability Exposure (E3). To minimize random error, we

conducted five independent assessment rounds (25 inferences per occupation) and used the mean as the final exposure score (Supplementary Text S3). The occupational-level GenAI exposure index was then obtained through aggregation, and neighborhood-level exposure was computed as the posting-weighted average of occupational exposure within each neighborhood. Assessments were completed in June 2025. Pairwise correlations among model assessments range from 0.66 to 0.85, indicating consistent ensemble behavior. A manual expert review of 100 randomly sampled occupations yielded 89% consistency with model assessments, confirming the robustness and reliability of the ensemble approach.

Pre-determined DID

To establish a causal link between GenAI exposure and wage changes, we adopt a pre-determined DID design. The core specification is

$$\ln(\text{wage})_{it} = \alpha + \delta\,(\text{GenAI}_i^{\text{pre}} \times \text{Post}_t) + \sum_k \gamma_k\,(\text{Confounder}_{k,i} \times \text{Post}_t) + \sum_j \beta_j X_{j,it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (1)$$

where $\ln(\text{wage})_{it}$ is the log average monthly salary of jobs posted in neighborhood i in year t; $\text{GenAI}_i^{\text{pre}}$ is the standardized pre-determined GenAI exposure measured in 2018; Post_t is the post-treatment dummy (= 1 for 2023–2024); $\text{Confounder}_{k,i}$ denotes the three concurrent-shock exposures based on the 2018 industrial structure; $X_{j,it}$ includes ln(company count), ln(nightlight), NDVI, POI density and land-use ratio; and μ_i and λ_t are neighborhood and year fixed effects, respectively. The analysis window spans 2020–2024, comprising three pre-treatment years (2020, 2021, 2022) and two post-treatment years (2023, 2024), totalling 6,895 neighborhood–year observations across 1,379 neighborhoods. Standard errors are clustered at the neighborhood level. The treatment variable's strict pre-determination is central: GenAI exposure is measured in 2018, four years before ChatGPT's release (November 2022) and prior to the analysis window (from 2020), ruling out reverse causality and anticipatory responses. The core parameter δ estimates the differential change in wages following the ChatGPT shock per standard-deviation increase in pre-determined GenAI exposure, conditional on concurrent-shock controls.

Confounder exposure construction

The three concurrent-shock exposures are constructed within a unified framework. For each shock (technology regulation, COVID-19 recovery, real-estate adjustment), we identify the set of directly affected industries and compute each neighborhood's 2018 posting share in that set:

$$\text{Confounder}_{k,i} = \frac{\sum_{j \in \mathcal{J}_k} N_{ij,2018}}{\sum_{j=1}^{18} N_{ij,2018}} \quad (2)$$

where $N_{ij,2018}$ is the number of recruitment postings for neighborhood i in industry j in 2018 and $\mathcal{J}_k$ is the set of industries directly affected by shock k. Technology-regulation exposure uses the combined posting share of Information Technology and Software (Industry Code I) and Education (P); COVID-19-recovery exposure uses the combined share of Accommodation and Catering (H), Transportation (G) and Wholesale/Retail (F); real-estate-adjustment exposure uses the combined share of Real Estate (K) and Construction (E). All three exposures are standardized and interacted with Post_t before inclusion.

Event-study specification

To verify parallel trends and characterize the dynamic effect, we estimate the event-study model

$$\ln(\text{wage})_{it} = \alpha + \sum_{k \neq 2022} \delta_k \left(\text{GenAI}_i^{\text{pre}} \times \mathbf{1}[t = k]\right) + \sum_m \gamma_m \left(\text{Confounder}_{m,i} \times \text{Post}_t\right) + \sum_j \beta_j X_{j,it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (3)$$

with 2022 as the base period and $\mathbf{1}[t = k]$ an indicator equal to one in year k. The identifying assumption requires that all pre-treatment coefficients $\delta_k (k = 2020,2021)$ satisfy $\delta_k = 0$, meaning that prior to the GenAI shock, neighborhoods with different exposure levels followed a common wage trend. We assess this assumption with a joint Wald test.

Randomization inference

To evaluate the robustness of statistical inference without relying on large-sample asymptotics or specific error-structure assumptions, we implement randomization inference grounded in Fisher's exact test. The procedure is: (1) randomly permute pre-determined GenAI exposure across the 1,379 neighborhoods, breaking the actual mapping between treatment and outcome; (2) re-estimate the DID coefficient in Equation (1) with concurrent-shock controls using the permuted exposures; (3) repeat 500 times to generate a placebo distribution; and (4) compute the two-sided permutation p-value. The procedure is exact in finite samples and robust to arbitrary spatial correlation and heteroskedasticity. The true estimate lies in the extreme left tail of the placebo distribution.

Bartik shift-share IV

As a complementary identification strategy, we construct a Bartik shift-share instrument. For each of the 18 major industry sectors, we compute a leave-one-out industry-average GenAI exposure that excludes the focal neighborhood's contribution:

$$\overline{\mathrm{GenAI}}_{j,-i} = \frac{\sum_{i' \neq i} N_{i'j,2018}\ \mathrm{GenAI}_{i',2018}}{\sum_{i' \neq i} N_{i'j,2018}} \quad (4)$$

where $N_{i'j,2018}$ is the number of postings in neighborhood i for industry j in 2018 and $\mathrm{GenAI}_{i',2018}$ is the pre-determined GenAI exposure of neighborhood i. Interacting neighborhood i's 2018 industry-posting share $s_{ij,2018} = N_{ij,2018}/N_{i,2018}$ with the industry-level exposure and summing over the 18 industries yields:

$$\mathrm{Bartik}_i = \sum_{j=1}^{18} s_{ij,2018}\ \overline{\mathrm{GenAI}}_{j,-i} \quad (5)$$

This instrument is driven solely by historical industrial structure and cross-neighborhood industry-level exposure, excluding endogenous adjustments by the focal neighborhood.

To maximize first-stage power, IV estimation employs a broader window than the event study. Using 2018–2019 averages as the pre-period and 2023–2024 averages as the post-period,

we compute long differences (Δ) for each variable and estimate cross-sectional two-stage least squares (2SLS) on all neighborhoods:

First stage:

$$\Delta \mathrm{GenAI}_i = \pi_0 + \pi_1 \mathrm{Bartik}_i + \pi_2 \Delta X'_i + u_i \quad (6)$$

Second stage:

$$\Delta \ln(\mathrm{wage})_i = \theta_0 + \beta_{\mathrm{IV}} \Delta \widehat{\mathrm{GenAI}}_i + \theta_2 \Delta X'_i + v_i \quad (7)$$

Reduced form:

$$\Delta \ln(\mathrm{wage})_i = \rho_0 + \beta_{\mathrm{RF}} \mathrm{Bartik}_i + \rho_2 \Delta X'_i + e_i \quad (8)$$

where $\Delta X'_i$ includes the change in log company count, the change in log nighttime-light intensity, the changes in NDVI, POI density and land-use intensity, and the three concurrent-shock exposure measures. Instrument strength is evaluated via the first-stage F-statistic (F = 66.2), far exceeding the critical value of 10 of Stock and Yogo (2002)[44]; weak-instrument-robust Anderson–Rubin inference further rejects a zero effect (AR F = 69.61, $P = 1.7 \times 10^{-16}$). The reduced form is negative and highly significant (−0.652 log points, s.e. = 0.078, $P < 0.001$), providing directionally consistent corroborating evidence for the DID design.

Fixed-effects model

To test the de-skilling and crowding mechanisms, we estimate two-way fixed-effects models in which the GenAI wage response is allowed to vary with a time-varying channel variable—mean required education for the de-skilling channel and job-market heat for the crowding channel.

The specifications are

$$\ln(\mathrm{wage})_{it} = \alpha + \delta_1 \mathrm{GenAI}_{it} + \delta_2 \mathrm{Educ}_{it} + \delta_3 (\mathrm{GenAI}_{it} \times \mathrm{Educ}_{it}) + \sum_j \beta_j X_{j,it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (9)$$

$$\ln(\text{wage})_{it} = \alpha + \delta_1 \text{GenAI}_{it} + \delta_2 \text{Crowd}_{it} + \delta_3 (\text{GenAI}_{it} \times \text{Crowd}_{it}) + \sum_j \beta_j X_{j,it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (10)$$

where GenAI_{it} is standardized time-varying GenAI exposure; Educ_{it} is standardized mean required education; Crowd_{it} is standardized recruitment heat (market crowding); δ_3 tests the de-skilling effect (Equation 9) and the crowding effect (Equation 10); $X_{j,it}$ is the vector of time-varying controls; and μ_i and λ_t are neighborhood and year fixed effects. Standard errors are two-way clustered by neighborhood and year. Because the channel variables may themselves respond to the GenAI shock, we additionally estimate a design in which the interaction is formed with pre-determined 2018 channel values ($\text{GenAI}_{it} \times \text{Educ}_{i,2018}$ and $\text{GenAI}_{it} \times \text{Crowd}_{i,2018}$); both channels remain significantly negative in this pre-determined-moderator specification (Supplementary Table S10, Panel F), confirming that the moderating patterns do not depend on endogenous post-treatment adjustment of the channel variables.

To assess the robustness of the interaction-term coefficients, we conducted a placebo test based on random permutation: the GenAI exposure variable was randomly shuffled across neighborhood–year observations, the shuffled variable was used to reconstruct the interaction terms, and the placebo interaction coefficients were re-estimated. Repeating this procedure 500 times yielded an empirical distribution of placebo coefficients centered on zero. The fact that the true interaction coefficients lie far in the tail of this distribution indicates that the observed de-skilling and crowding effects are not statistical noise (Supplementary Figures S6–S7).

Data availability

The processed, neighborhood-level datasets that support the findings of this study are publicly available in the GitHub repository accompanying this manuscript at https://github.com/tututu0331/genai-wage-beijing. The raw online job-posting data were obtained from third-party platforms under agreements that prohibit redistribution; they are therefore not released. All aggregated data used in the reported analyses are provided in full, and requests for further information should be directed to the corresponding author.

Code availability

All code used to process the data and to reproduce the analyses and figures of this study is publicly available in the GitHub repository at https://github.com/tututu0331/genai-wage-beijing.

Funding Declaration

This research was funded by the National Natural Science Foundation of China (#72274101).

References

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

2 Webb, M. The impact of artificial intelligence on the labor market. *Available at SSRN 3482150* (2019).

3 Frank, M. R. *et al.* Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences* **116**, 6531–6539 (2019).

4 Cazzaniga, M. *et al. Gen-AI: Artificial intelligence and the future of work*. (International Monetary Fund, 2024).

5 Gmyrek, P., Berg, J. & Bescond, D. Generative AI and jobs: A global analysis of potential effects on job quantity and quality. (2023).

6 Cairncross, F. The death of distance: How the communications revolution will change our lives. *(No Title)* (1997).

7 Friedman, T. L. *The world is flat: A brief history of the twenty-first century*. (Macmillan, 2005).

8 Autor, D. H. Why are there still so many jobs? The history and future of workplace automation. *Journal of economic perspectives* **29**, 3–30 (2015).

9 Barrero, J. M., Bloom, N. & Davis, S. J. Why working from home will stick. (National Bureau of Economic Research, 2021).

10 Ramani, A., Alcedo, J. & Bloom, N. How working from home reshapes cities. *Proceedings of the National Academy of Sciences* **121**, e2408930121 (2024).

11 Florida, R. *The new urban crisis: How our cities are increasing inequality, deepening segregation, and failing the middle class-and what we can do about it*. (Basic books, 2017).

12 Autor, D. H. in *AEA Papers and Proceedings.* 1–32 (American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203).

13 Moretti, E. The effect of high-tech clusters on the productivity of top inventors. *American Economic Review* **111**, 3328–3375 (2021).

14 Vinuesa, R. *et al.* The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature communications* **11**, 233 (2020).

15 Capraro, V. *et al.* The impact of generative artificial intelligence on socioeconomic inequalities and policy making. *PNAS nexus* **3**, pgae191 (2024).

16 Katz, L. F. & Murphy, K. M. Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics* **107**, 35–78 (1992).

17 Card, D. & DiNardo, J. E. Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of labor economics* **20**, 733–783 (2002).

18 Moretti, E. *The new geography of jobs*. (Houghton Mifflin Harcourt, 2012).

19 Bárány, Z. L. & Siegel, C. Job polarization and structural change. *American Economic Journal: Macroeconomics* **10**, 57–89 (2018).

20 Acemoglu, D. & Autor, D. in *Handbook of labor economics* Vol. 4 1043–1171 (Elsevier, 2011).

21 Autor, D. H., Levy, F. & Murnane, R. J. The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics* **118**, 1279–1333 (2003).

22 Bloom, D. E., Prettner, K., Saadaoui, J. & Veruete, M. Artificial intelligence and the skill premium. *Finance Research Letters* **81**, 107401 (2025).
23 Brynjolfsson, E., Li, D. & Raymond, L. Generative AI at work. *The Quarterly Journal of Economics* **140**, 889–942 (2025).
24 Noy, S. & Zhang, W. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **381**, 187–192 (2023).
25 Dell'Acqua, F. *et al.* Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science* **37**, 403–423 (2026).
26 Wu, Y., Lin, Z., Zhang, Q. & Wang, W. Artificial intelligence, wage dynamics, and inequality: Empirical evidence from Chinese listed firms. *International Review of Economics & Finance* **96**, 103739 (2024).
27 Acemoglu, D. The simple macroeconomics of AI. *Economic Policy* **40**, 13–58 (2025).
28 Cousins, J. J. Artificial intelligence's metabolic footprint and the political-industrial ecology of urban sustainability. *npj Urban Sustainability* (2026).
29 Kok, S. & Weel, B. t. Cities, tasks, and skills. *Journal of Regional Science* **54**, 856–892 (2014).
30 Hui, X., Reshef, O. & Zhou, L. The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science* **35**, 1977–1989 (2024).
31 Sun, Y., Zhu, H., Wang, L., Zhang, L. & Xiong, H. Large-scale online job search behaviors reveal labor market shifts amid COVID-19. *Nature Cities* **1**, 150–163 (2024).
32 Tolan, S. *et al.* Measuring the occupational impact of AI: tasks, cognitive abilities and AI benchmarks. *Journal of Artificial Intelligence Research* **71**, 191–236 (2021).
33 Acemoglu, D. & Restrepo, P. Tasks, automation, and the rise in US wage inequality. *Econometrica* **90**, 1973–2016 (2022).
34 Demombynes, G., Langbein, J. & Weber, M. *The exposure of workers to Artificial Intelligence in low-and middle-income countries*. (World Bank, 2025).
35 Castells, M. *The rise of the network society*. (John wiley & sons, 2011).
36 You, X. China intensifies push to become world leader in tech and AI. *Nature* **651**, 859–860 (2026).
37 Wu, F. *et al.* Towards a new generation of artificial intelligence in China. *Nature machine intelligence* **2**, 312–316 (2020).
38 Tian, G., Wu, J. & Yang, Z. Spatial pattern of urban functions in the Beijing metropolitan region. *Habitat International* **34**, 249–255 (2010).
39 Lai, Y. & Zhao, H. Comparative analysis of smart city scientific research trends in the USA and China. *Nature Cities* **2**, 875–883 (2025).
40 Acemoglu, D. & Restrepo, P. Robots and jobs: Evidence from US labor markets. *Journal of political economy* **128**, 2188–2244 (2020).
41 Kraay, A. & McKenzie, D. Do poverty traps exist? Assessing the evidence. *Journal of Economic Perspectives* **28**, 127–148 (2014).
42 Acemoglu, D. & Restrepo, P. Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives* **33**, 3–30 (2019).
43 Oster, E. Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics* **37**, 187–204 (2019).

44 Stock, J. H. & Yogo, M. (National Bureau of Economic Research Cambridge, Mass., USA, 2002).
45 Goldin, C. & Katz, L. F. (National Bureau of Economic Research Cambridge, Mass., USA, 2007).
46 Acemoglu, D. Why do new technologies complement skills? Directed technical change and wage inequality. *The quarterly journal of economics* **113**, 1055–1089 (1998).
47 Hartley, J., Jolevski, F., Melo, V. & Moore, B. The labor market effects of generative artificial intelligence. *Available at SSRN* (2024).
48 Ferdman, A. AI deskilling is a structural problem. *AI & SOCIETY* **41**, 3001–3013 (2026).
49 Hossain, T. & Li, K. K. Crowding out in the labor market: A prosocial setting is necessary. *Management Science* **60**, 1148–1160 (2014).
50 Cheng, S. & Fan, Q. From Interaction to Involution: Human-AI Interaction and Employee Work Involution. *International Journal of Human–Computer Interaction*, 1–18 (2026).
51 Storper, M. & Venables, A. J. Buzz: face-to-face contact and the urban economy. *Journal of economic geography* **4**, 351–370 (2004).
52 Glaeser, E. *Triumph of the city: How our greatest invention makes us richer, smarter, greener, healthier, and happier*. (Penguin, 2012).
53 Li, L. & Zhao, N. Explicit and tacit knowledge have diverging urban growth patterns. *NPJ Urban Sustainability* **3**, 34 (2023).
54 Glaeser, E. L. & Resseger, M. G. The complementarity between cities and skills. *Journal of regional science* **50**, 221–244 (2010).
55 Florida, R. L. *The rise of the creative class*. Vol. 512 (Basic books New York, 2019).
56 Shen, Y. & Zhang, X. The impact of artificial intelligence on employment: the role of virtual agglomeration. *Humanities and Social Sciences Communications* **11** (2024).
57 Goldfarb, A. & Tucker, C. Digital economics. *Journal of economic literature* **57**, 3–43 (2019).
58 Pan, Y., Gao, F. & Xing, Z. The impact of skill-biased technological change on urban-rural income inequality: Evidence from China. *Applied Geography* **187**, 103891 (2026).
59 Chuang, Y.-T., Chiang, H.-L. & Lin, A.-P. Insights from the Job Demands–Resources Model: AI's dual impact on employees' work and life well-being. *International Journal of Information Management* **83**, 102887 (2025).
60 Cui, K. Z. *et al.* The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. *Management Science* (2026).
61 Ozgul, P., Fregin, M.-C., Stops, M., Janssen, S. & Levels, M. High-skilled human workers in non-routine jobs are susceptible to AI automation but wage benefits differ between occupations. *arXiv preprint arXiv:2404.06472* (2024).
62 Massenkoff, M. & McCrory, P. Labor market impacts of AI: A new measure and early evidence. *Anthropic Research* **5** (2026).
63 Lazear, E. P. & Oyer, P. Internal and external labor markets: a personnel economics approach. *Labour economics* **11**, 527–554 (2004).
64 Jaumotte, M. F. *et al.* Bridging Skill Gaps for the Future: New Jobs Creation in the AI Age. (International Monetary Fund, 2026).
65 Deming, D. J. The growing importance of social skills in the labor market. *The quarterly journal of economics* **132**, 1593–1640 (2017).

66 Hou, C. *et al.* The effects of critical thinking intervention on reliance behaviors, problem-solving quality, and creativity during human-generative AI collaborative learning. *Computers & Education*, 105576 (2026).
67 Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., Bobes-Bascarán, J. & Fernández-Leal, Á. Human-in-the-loop machine learning: a state of the art: E. Mosqueira-Rey et al. *Artificial Intelligence Review* **56**, 3005–3054 (2023).
68 George, A. S. Artificial intelligence and the future of work: job shifting not job loss. *Partners Universal Innovative Research Publication* **2**, 17–37 (2024).
69 Yan, X., Zhu, Q., Yu, S. & Huang, X. AI technology and labor skill transformation: Theoretical and empirical evidence from China. *Technological Forecasting and Social Change* **227**, 124619 (2026).
70 Tambe, P. B. Reskilling the workforce for AI: Domain expertise and algorithmic literacy. *Management Science* **72**, 515–537 (2026).
71 Weber, M., Hein, A., Weking, J. & Krcmar, H. Orchestration logics for artificial intelligence platforms: From raw data to industry-specific applications. *Information Systems Journal* **35**, 1015–1043 (2025).
72 Hershbein, B. & Kahn, L. B. Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review* **108**, 1737–1772 (2018).
73 Graham, S. Bridging urban digital divides? Urban polarisation and information and communications technologies (ICTs). *Urban studies* **39**, 33–56 (2002).
74 Capello, R. & Lenzi, C. Spatial heterogeneity in knowledge, innovation, and economic growth nexus: conceptual reflections and empirical evidence. *Journal of Regional Science* **54**, 186–214 (2014).
75 Chen, Q., Ge, J., Xie, H., Xu, X. & Yang, Y. Large language models at work in China's labor market. *China Economic Review* **92**, 102413 (2025).

Ethical Approval

This article does not contain any studies with human participants performed by any of the authors.

Informed Consent

Not applicable.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- 20261009SupplementaryInformation.docx