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.

Generative AI impacts on intra-urban inequality and skill premium in Beijing

YUAN LAI

yuanlai@mail.tsinghua.edu.cn

Tsinghua University https://orcid.org/0000-0003-0664-5048

Xiliu He

https://orcid.org/0009-0007-8476-4831

Haoxiang Zhao

Tsinghua University https://orcid.org/0009-0004-4846-6392

Mingyi Ma

Tsinghua University

Edward Lai

Tsinghua University https://orcid.org/0009-0005-0263-700X

Koei Enomoto

Tsinghua University

Anni Hu

Tsinghua University

Jiatong Li

Tsinghua University

Lingyun Chu

ZODA LAB

Article

Keywords:

Posted Date: July 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9955657/v1

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

Additional Declarations: There is **NO** Competing Interest.

Generative AI impacts on intra-urban inequality and skill premium in Beijing

Abstract

Generative artificial intelligence (GenAI) is the first automation wave to reach high-cognitive tasks at scale, yet its effects on intra-urban inequality remain largely unknown. Using 5 million job postings from Beijing (2018–2024), we construct a neighborhood-level GenAI Exposure Index by aggregating task-level assessments from five leading large language models. We examine the spatial, structural and causal mechanisms of this shock. We find that GenAI exposure is highly concentrated in the city's core districts, deepening the intra-urban AI divide. Since 2023, high-exposure neighborhoods have experienced wage stagnation even as they continue to attract high-skilled workers—a "high-skill trap." This wage penalty is driven by task de-skilling and intensified labor-market crowding. A difference-in-differences design centered on ChatGPT's release supports a causal interpretation. These findings challenge the prevailing theory of skill-biased technological change and provide a basis for inclusive AI governance in global technology hubs.

Introduction

Cities concentrate the workers and industries most exposed to Generative Artificial Intelligence (GenAI)[1-3]. As the first automation wave to penetrate cognitive, creative, and communicative tasks—capabilities that have long commanded an urban wage premium—GenAI raises a fundamental question: will it act as a "great leveler"[4,5] that disperses opportunity to urban peripheries[6-8], or will it deepen core–periphery divides and widen wage gaps[9-11]? The answer has direct implications for intra-urban digital divides, labor protections in exposed occupations, and the governance of global tech hubs[3,12,13].

Previous waves of technology (Supplementary Text S1)—from computerization 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 city cores[10,14-19]. GenAI breaks 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. Three frameworks compete. The augmentation view holds that GenAI raises productivity across the skill distribution and creates new demand[20,21]. The substitution view warns of structural displacement and deepening inequality[1,22]. The task-reorganization view emphasizes that technology reshapes the content of work without wholesale replacement[3,19,23]. Recent evidence points to a more counterintuitive scenario, which we call a "high-skill trap": a surge of skilled workers in exposed occupations coincides with stagnant wages, as GenAI erodes the scarcity value of premium cognitive tasks[1,2,20,24] (Fig. 1).

Three critical gaps persist. First, existing studies are predominantly national in scale, cross-sectional, and focused on high-income Western economies; they cannot reveal how GenAI reshapes intra-urban core–periphery dynamics at fine spatial resolution[3,11,25]. Second, the literature lacks dynamic evidence linking exposure to synchronous changes in skill supply and wages, 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[1,26-28], 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[29,30].

We address these gaps using 4,995,615 online job postings from Beijing (2018–2024), matched at the neighborhood scale 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[31-34]. To measure exposure, we designed a matching pipeline that uses retrieval-augmented generation (RAG) to link job tasks to AI capabilities and then aggregated assessments from five state-of-the-art (SOTA) LLMs to reduce model-specific bias. Our analysis yields three interconnected findings. First, GenAI exposure is spatially locked into urban cores, deepening the intra-urban AI divide. Second, after ChatGPT's release, neighborhoods with high exposure experience both rising skilled-labor supply and stagnant 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 the first 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.

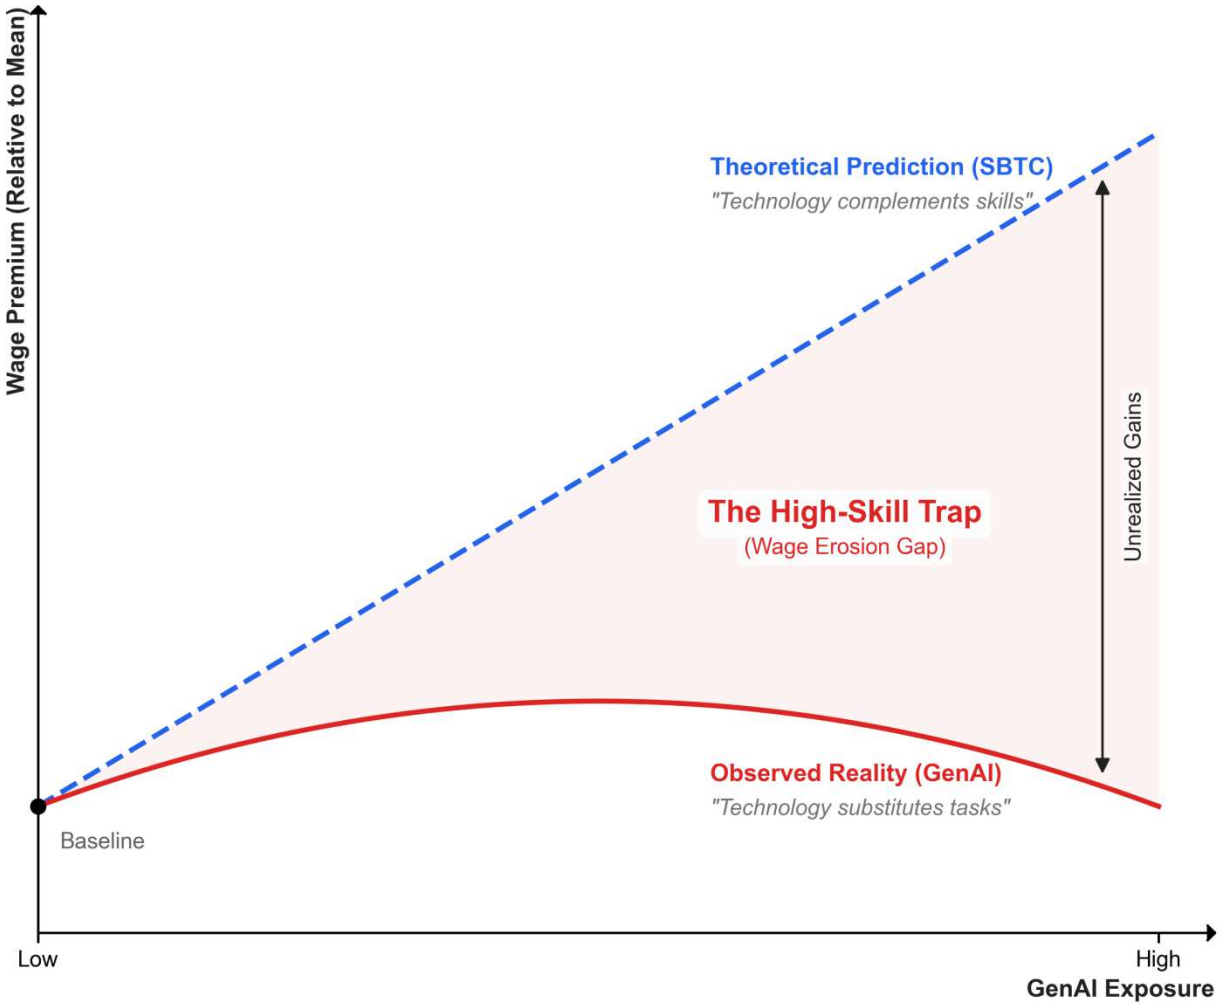


Fig. 1 | The "high-skill trap" under GenAI exposure. The blue dashed line shows the skill premium predicted by traditional skill-biased technological change (SBTC) theory: as exposure to technology rises, high-skilled labor and the skill premium grow linearly. The red solid line traces the pattern we observe empirically: after an initial rise, the skill premium flattens and ultimately stagnates, revealing diminishing marginal returns. 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 increases.

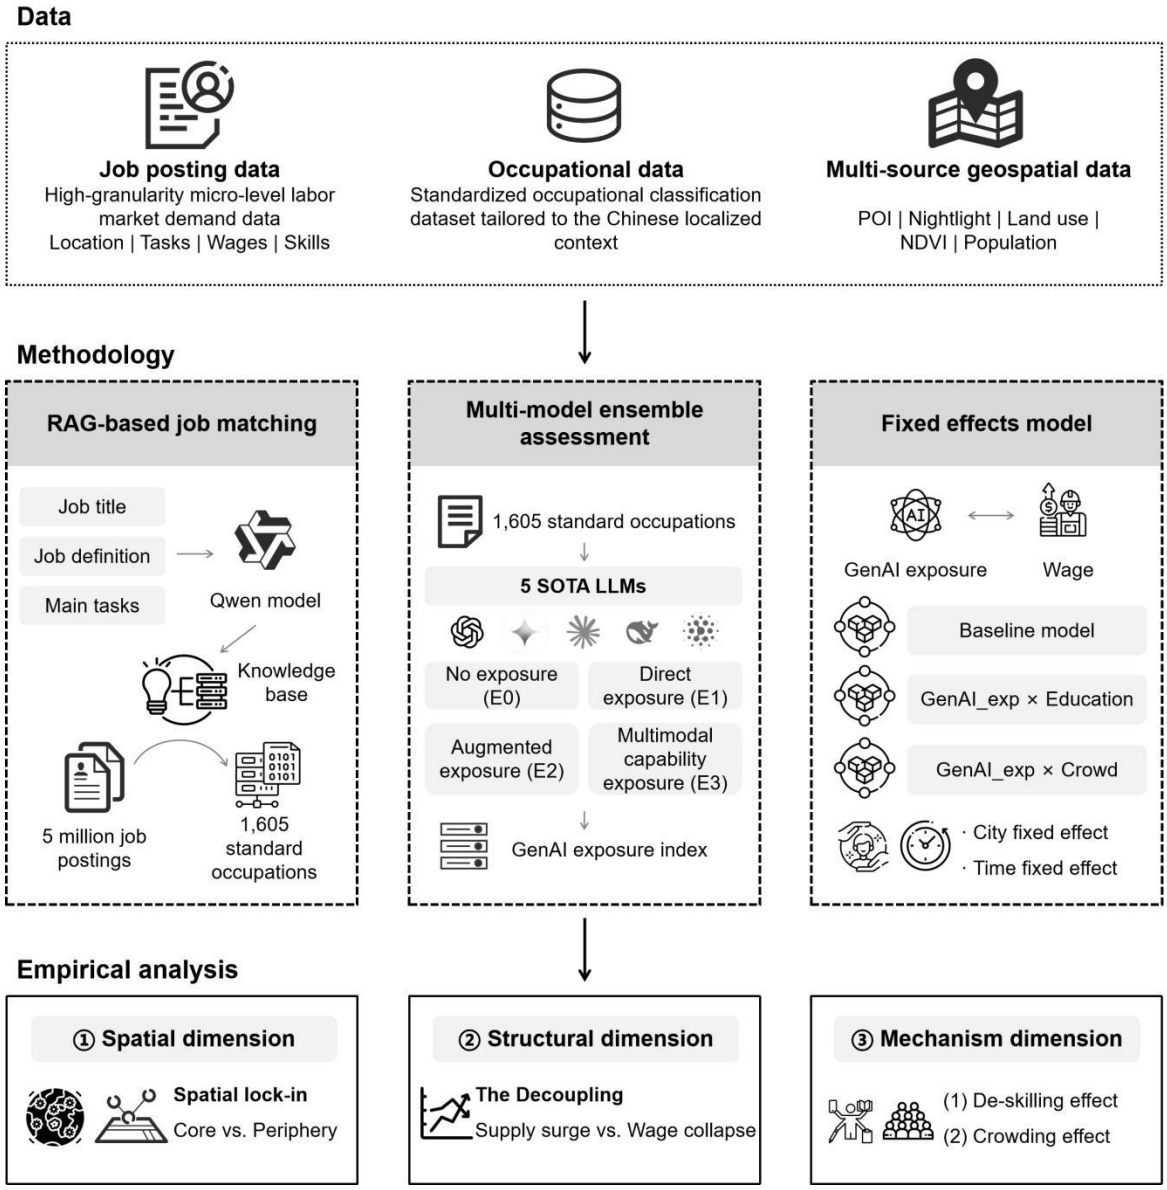


Fig. 2 | Measuring GenAI exposure and its spatial, structural, and mechanistic effects. The framework combines three steps: (1) millions of job posting are linked to standard occupational tasks using retrieval-

augmented generation (RAG); (2) GenAI exposure scores are computed by aggregating assessments from five SOTA LLMs; 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 does not spread evenly across the city. Contrary to the idea that new technologies dissolve geographic barriers[4,5], we find that GenAI exposure reinforces Beijing's existing spatial hierarchy. Operating through agglomeration economies, the shock deepens the core–periphery divide and projects a new digital divide onto the urban landscape. The spatial pattern of GenAI exposure across Beijing (2018–2024) follows a monocentric, multi-ring structure with pronounced heterogeneity (Fig. 3a). High-exposure areas are locked into a "golden triangle" of knowledge-intensive services—Zhongguancun Science City (innovation), Financial Street (capital management) and the Guomao CBD (business services)—where exposure scores consistently exceed 0.3. This concentration indicates that high-cognitive, creative occupations most susceptible to AI (such as algorithm engineers and financial analysts) still depend on the face-to-face interaction and knowledge spillovers that dense urban cores provide[35,36]. 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[37].

A temporal comparison (Fig. 3b) reveals two trajectories that confirm strong path dependence. First, core districts exhibit spatial lock-in and cumulative advantage: despite the pandemic and industry shocks, high-value clusters in Haidian and Chaoyang intensified rather than dispersed, with AI-related employment continuing to concentrate in talent-rich, information-dense areas[16]. Second, the post-ChatGPT GenAI boom triggered a selective spillover. Exposure rose markedly in Beijing's municipal sub-center between 2023 and 2024, showing that emerging clusters with sufficient absorptive capacity begin to capture these spillovers. Yet this diffusion is highly selective—outer suburban counties remain low-exposure, suggesting that the spread of AI exposure is constrained by gaps in digital infrastructure and complementary assets[38] and fails to reach the periphery.

A local spatial autocorrelation analysis (Local Indicators of Spatial Association, LISA) confirms this polarization (Fig. 3c). High–high clusters (neighborhoods with high exposure surrounded by similarly high values) are persistently locked into the core, and their spatial connectivity strengthened between 2018 and 2024; spillovers circulate mainly within this cluster, forming a closed network of knowledge exchange. Low–low clusters (low-exposure neighborhoods surrounded by low-exposure areas) sprawl contiguously toward the southwest and northeast, creating a "spatial poverty trap" where workers lack both direct exposure and a supportive environment for upgrading[39]. Transitional zones (high–low and low–high outliers) are largely absent, underscoring a spatial fragmentation so severe that the digital divide becomes especially difficult to bridge.

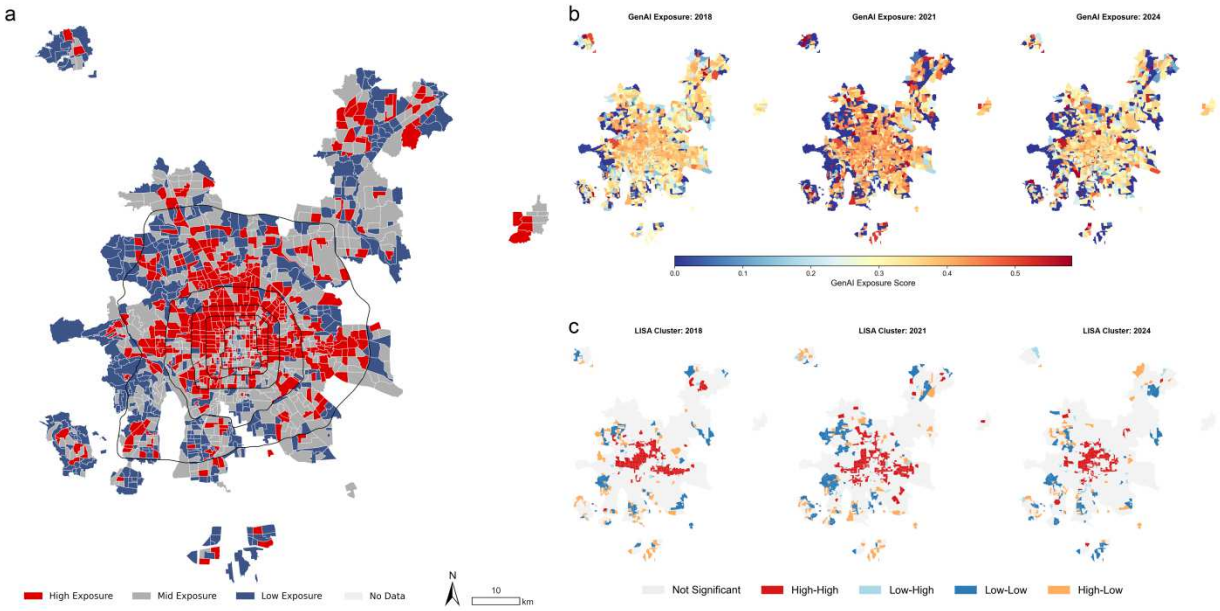


Fig. 3 | GenAI exposure across Beijing: spatial patterns, temporal change, and local clustering (2018–2024). a, GenAI exposure intensity averaged over 2018-2024 across Beijing's neighborhoods. **b**, Exposure distribution in 2018, 2021 and 2024, showing persistent core concentration and selective post-2023 spillover. **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 supply, and wages evolve together, testing whether the shock has structurally reversed the skill premium. The results challenge the conventional skill-biased technological change (SBTC) hypothesis[14] and reveal a "high-skill trap": although high-exposure neighborhoods continue to draw in highly educated workers, they have experienced marked wage stagnation and 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 strong positive correlation: exposure rises with education. A pronounced "elite lock-in" zone emerges, where the workforce averages more than 12 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 fall into a low-exposure, low-skill trap: 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 declined sharply after peaking in 2021. Following ChatGPT's release in late 2022, they fell further to 13,673 CNY per month, narrowing the gap with the medium- and low-exposure groups. Wage in low-exposure areas, by contrast, held relatively steady—consistent with a pattern in which GenAI compresses the wage distribution primarily by pulling down high-end wages[20]. 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 substitution[40].

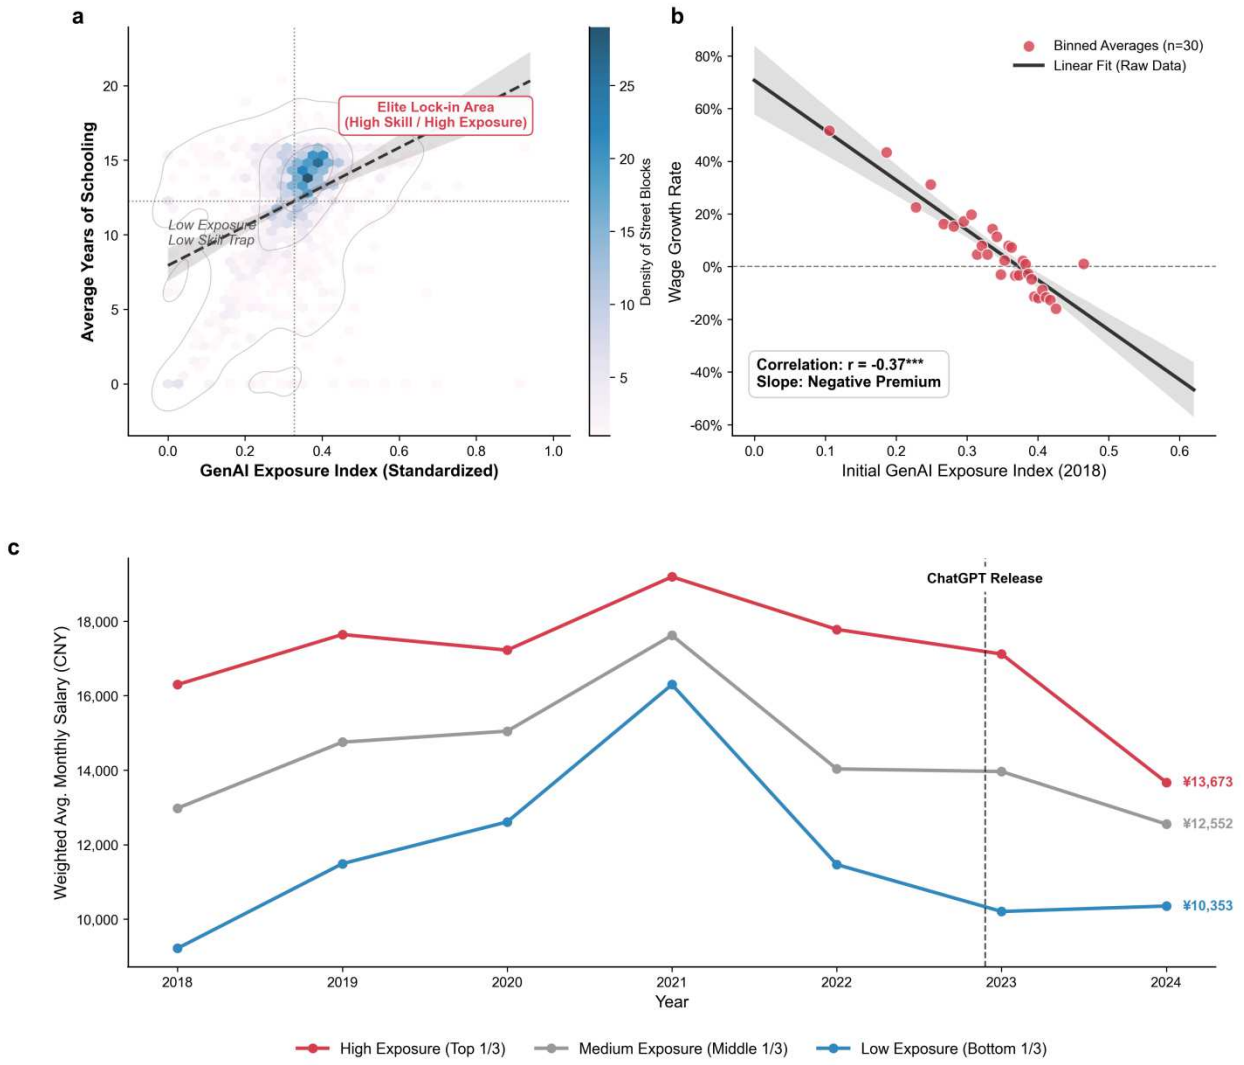


Fig. 4 | GenAI exposure, education, and the emerging wage penalty. a, Hexagonal heatmap of standardized GenAI exposure versus average years of schooling in 2024. The dashed line is a linear fit; the region where high education and high exposure overlap forms an "elite lock-in" zone. **b**, Binned scatter 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 attract highly educated workers and concentrate high GenAI exposure still see wages stagnate? Fig. 5 reveals a sharp contrast between low- and high-exposure areas. In low-exposure areas, skill composition and wages move broadly together. In high-

exposure areas, however, the two decouple: the share of high-skill jobs surges after 2021 as the urban core continues to draw in educated talent, yet the wage index for those same high-skill jobs declines substantially over the same period. This decoupling points to two mutually reinforcing channels. First, task substitution and de-skilling: GenAI lowers the cognitive threshold for tasks once concentrated in high-skill occupations—coding, copywriting, data analysis—making them more routine and eroding the bargaining power of skilled workers[41]. Second, crowding and hyper-competition: the steady inflow of educated talent into the core collides with a task space compressed by AI, generating an oversupply of workers relative to the remaining high-value tasks[42].

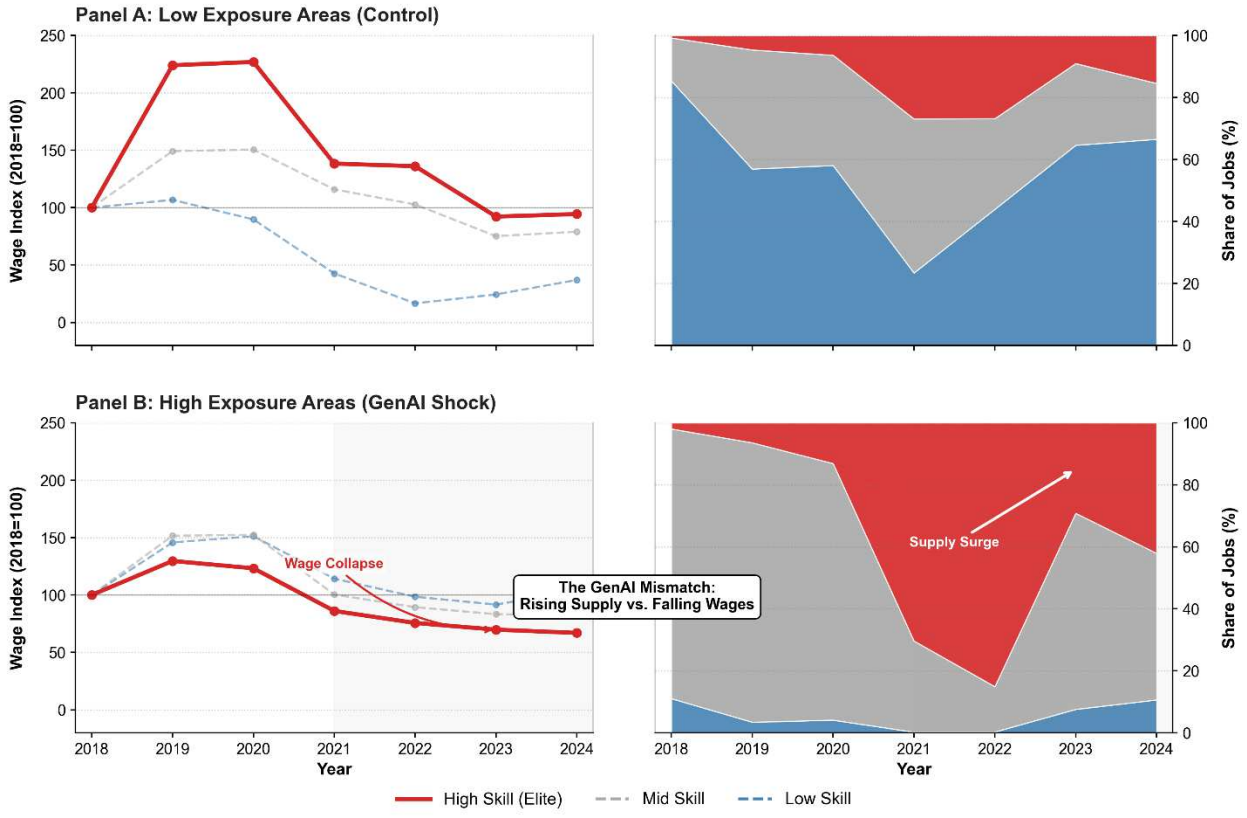


Fig. 5 | The decoupling: wage stagnation versus up-skilling in high-exposure areas. Panel A shows low-exposure (control) areas; Panel B shows high-exposure (treatment) areas. Lines depict the wage index and stacked areas depict the employment structure. The divergence in Panel B—a labor-supply surge alongside a wage collapse—is the focal pattern.

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 test robustness through randomization inference (Methods).

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 model interacts this pre-determined exposure with a post-2022 dummy indicator ($Post_t = \mathbf{1}[t \geq 2023]$) while controlling for the three concurrent shocks—each also interacted with post-period—to separate their confounding effects 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 wage trends.

Fig. 6 reports the event-study estimates with the concurrent-shock controls (2022 is the reference year; Equation 3). Pre-treatment coefficients ($\hat{\beta}_{2020} = 0.025$, $p = 0.791$; $\hat{\beta}_{2021} = -0.014$, $p = 0.874$) are statistically insignificant. A joint Wald test fails to reject the parallel trends assumption ($\chi^2 = 0.176$, $p = 0.916$). A sharp structural break appears post-treatment. The coefficient for 2023 drops to −0.193 ($p = 0.033$), implying that a one-standard-deviation increase in pre-determined exposure is associated with an approximately 17.5% decline in wages during the first year of the shock. The 2024 coefficient is −0.103 ($p = 0.284$), still negative but with weakened significance, consistent with partial labor-market adjustment through occupational transitions, skill restructuring, or the gradual emergence of new AI-complementary roles.

Table 1 reports the core DID estimates. In the baseline (Column 1), the interaction coefficient is −0.140 ($p < 0.05$), equivalent to a wage decline of about 13.1% post-shock. When the three concurrent-shock controls are added (Column 2), the coefficient does not shrink; it increases slightly to −0.151 ($p < 0.05$). 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]. None of the interactions with the concurrent shocks are statistically significant, ruling out the possibility that those shocks independently account for the findings.

We further test the robustness of the DID results with randomization inference. Shuffling the pre-determined exposure measure randomly across neighborhoods 500 times produces a distribution of placebo estimates tightly centered on zero. The true estimate (−0.151) lies in the extreme left tail of this distribution (permutation $p = 0.004$; only 2 of 500 simulations yielded a more extreme coefficient). This confirms that the wage penalty is neither noise nor an artefact of data structure.

As complementary evidence, we construct a Bartik shift-share IV that predicts local GenAI exposure using national AI-related job growth weighted by each neighborhood's 2018 industrial composition (Equations 4–8). The reduced-form estimate from this instrument implies that a one-standard-deviation increase in the Bartik predictor reduces wage growth by 0.436 log points ($p < 0.001$), which aligns closely with the DID estimate. The first-stage F-statistic of 53.5 far exceeds the weak-instrument threshold[44]. However, because the Bartik event-study fails the parallel-trends test (joint Wald $p = 0.004$), the structural parameter should be interpreted with caution. Full diagnostics results appear in Supplementary Text S6.

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

	(1) Baseline DID	(2) + Confounder controls
Dependent variable: ln(Wage)		
$GenAI_{2018}$ × Post	−0.140**	−0.151**
	(0.061)	(0.068)
TechReg × Post		0.035
		(0.061)
COVID × Post		0.092
		(0.087)
RealEstate × Post		−0.039
		(0.109)
Controls	Yes	Yes
Entity FE	Yes	Yes

	(1) Baseline DID	(2) + Confounder controls
Time FE	Yes	Yes
Observations	6,895	6,895
Neighborhoods	1,383	1,383
Pre-trend joint test (p)	0.758	0.916
Permutation p-value	0.004	0.004

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. 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(population), nightlight intensity, NDVI, POI density and land-use ratio. Standard errors clustered at the neighborhood level are in parentheses. ** $p < 0.05$; * $p < 0.1$.

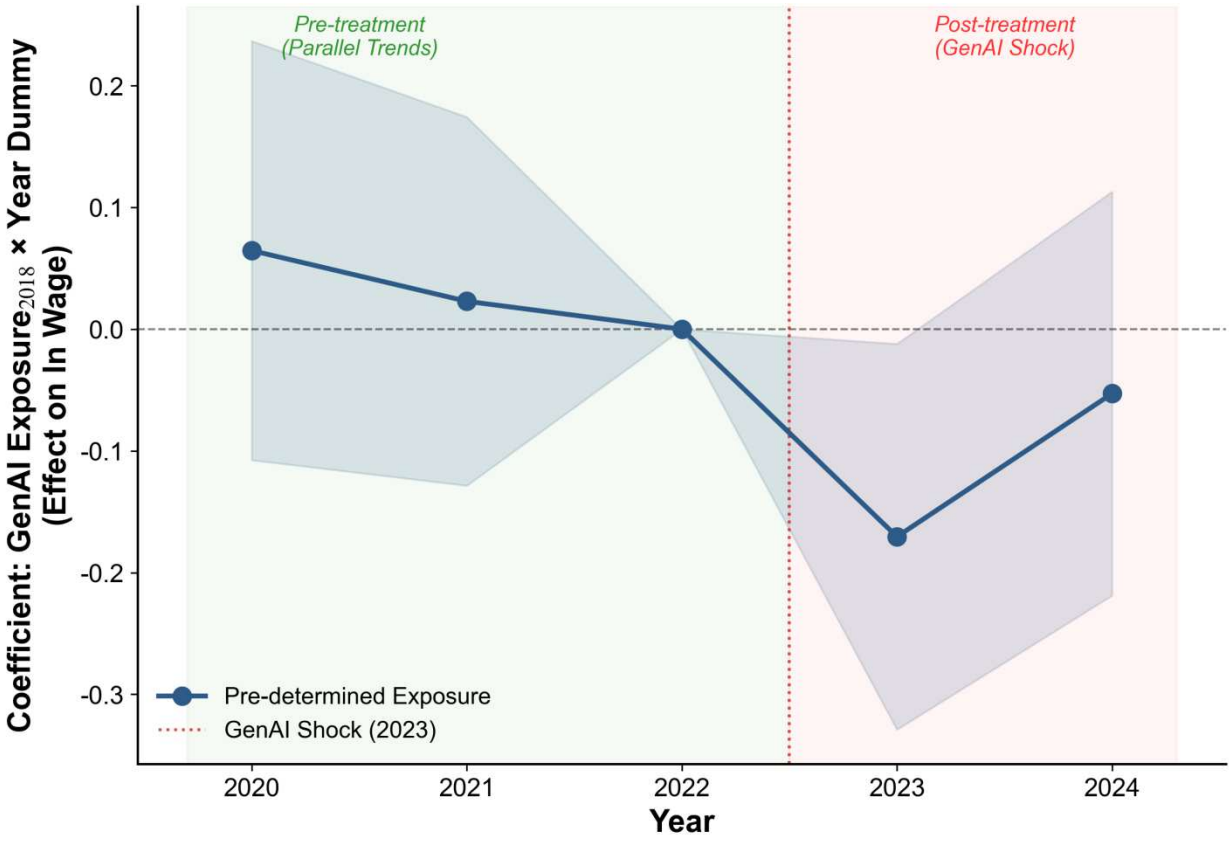


Fig. 6 | Causal identification of the GenAI wage penalty. Event-study estimates of the dynamic effect of pre-determined (2018) GenAI exposure on log wages, controlling for concurrent shocks. The reference year is 2022. Green shading marks the pre-treatment period (2020-2021) and red shading marks the post-treatment period (2023-2024). Pre-treatment coefficients are statistically indistinguishable from zero (joint Wald: $\chi^2 = 0.176$,

$p = 0.916$), supporting the parallel-trends assumption. A sharp negative break emerges in 2023 ($\beta = -0.193$, $p = 0.033$), followed by partial recovery in 2024 ($\beta = -0.103$, $p = 0.284$). 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, widening the skill premium[45]. Recent work on GenAI points to the opposite mechanism—de-skilling, or skill-levelling. The idea is that GenAI lowers the cognitive barriers to tasks once reserved for high-skilled workers (such as coding, copywriting, and data analysis), allowing lower-skilled workers to substitute for high-skilled labor and compressing wages at the top[40,41,46]. We test this directly by examining how education moderates the wage effect of GenAI exposure.

Column 2 of Table 2 shows that, although the main effect of education on wages remains significantly positive (1.150), the interaction between GenAI and education is negative and highly significant (−1.286). This means that the extra wage gain associated with GenAI exposure shrinks as education rises. Supplementary Fig. S1a reveals two regimes. For low-education workers, GenAI exposure has a positive marginal effect—consistent with an assistive role in which AI helps workers clear skill-specific task barriers[20]. As education increases, however, this effect declines and eventually turns negative for the most highly educated workers, including those with master's degrees and above. For this elite cohort, GenAI's capabilities overlap substantially with their own human capital (programming, translation, writing), and the technology acts more as a substitute than a complement, eroding their bargaining power. These results provide strong support for the de-skilling hypothesis[2,46].

Skill substitution alone does not explain the full wage penalty. A second, complementary channel operates through labor-market crowding. According to the crowding hypothesis[42], when the supply of workers in AI-exposed fields grows rapidly and the demand for their tasks does not expand enough, productivity gains can paradoxically push wages down—a pattern often called "involution." We capture this dynamic using a measure of job-market heat, which reflects the intensity of competition for positions in each occupation. Column 3 of Table 2 reports a highly significant negative interaction between GenAI exposure and job-market heat (−1.619), indicating that greater competitive intensity strongly suppresses the GenAI skill premium. Supplementary

Fig. S1b traces three distinct phases. In low-competition niches, the marginal effect of GenAI exposure is positive, generating a "blue ocean" of excess returns through scarcity. As adoption diffuses and competition approaches the mean, the premium shrinks and crosses zero, marking an equilibrium phase. In the most competitive segments—"red ocean" phase—the marginal effect turns sharply negative, confirming a crowding-driven wage penalty.

Together, these findings reveal a technology-induced saturation mechanism: widespread GenAI adoption lowers barriers and homogenizes the skill signals that once differentiated high-skilled workers, intensifying competition and compressing wages. To move from panel associations to causal evidence, we estimate a triple-DID specification that interacts pre-determined exposure, the post-2022 period, and either education or job-market heat. The triple interaction coefficients are -0.135 ($p = 0.177$) for the de-skilling channel and -0.063 ($p = 0.439$) for the crowding channel. Both signs are consistent with the panel results. Given the short post-treatment window, neither reaches conventional significance, but their directional consistency offers preliminary causal support (Supplementary Text S7).

Table 2 | Mechanisms underlying the GenAI wage effect

Dependent Variable: ln(Wage)	(1) Baseline	(2) De-skilling	(3) Crowding
GenAI Exposure (Std.)	0.835*** (0.068)	0.800*** (0.087)	-0.233* (0.138)
Education (Std.)	1.699*** (0.120)	1.150*** (0.099)	1.369*** (0.074)
Job Market Heat (Std.)	1.736*** (0.168)	1.194*** (0.176)	1.633*** (0.260)
Interactions			
GenAI × Education		-1.286*** (0.053)	
GenAI × Job Market Heat			-1.619*** (0.088)
Controls			
Population (ln)	-101.250*** (31.637)	-71.956** (29.076)	-71.868*** (27.459)
Nightlight	0.147 (0.143)	0.164 (0.116)	0.113 (0.115)
NDVI (Std.)	0.219 (0.208)	0.126 (0.111)	0.024 (0.147)
POI Density (Std.)	0.034	-0.006	-0.087

Dependent Variable: ln(Wage)	(1) Baseline	(2) De-skilling	(3) Crowding
	(0.100)	(0.049)	(0.061)
Land Use Ratio (Std.)	0.433 (0.385)	0.174 (0.284)	0.208 (0.249)
Model Statistics			
Entity FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Observations	9,653	9,653	9,653
R-squared	0.613	0.747	0.723
F-statistic	1636.2	2706.8	2397.1

Notes. The dependent variable is the natural logarithm of monthly wages. Column 1 reports the baseline two-way fixed-effects specification. Column 2 introduces the interaction between GenAI exposure and education to test the de-skilling channel. Column 3 introduces the interaction between GenAI exposure and job-market heat to test the crowding channel. Standard errors (in parentheses) are clustered at the neighborhood level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Variance inflation factors are below 5 for all variables, ruling out serious multicollinearity. While the DID analysis is restricted to the 2020–2024 window to allow parallel-trends verification, these regressions use the full 2018–2024 panel (N = 9,653 neighborhood–year observations) to maximize statistical power for estimating conditional correlations and their heterogeneity.

Discussion

This study uses fine-grained data from Beijing (2018–2024) to reveal a triple polarization—spatial, structural and wage-related—that GenAI is producing in a global megacity. 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[4]. In reality, GenAI-exposed jobs remain tightly clustered in the "golden triangle" of Zhongguancun, Financial Street, and the CBD, and this clustering did not weaken after ChatGPT's release—if anything, it intensified. 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[35,47,48]. GenAI, rather than dissolving the need for proximity, acts as a complementary asset that magnifies the returns to agglomeration, leaving peripheral neighborhoods that lack this cognitive infrastructure at deepening disadvantage[16,38].

More fundamentally, our results document the breakdown of SBTC theory in the context of GenAI. The long-standing view that new technologies inherently favor high-skilled workers and therefore widen the skill premium[18,27,45,49] does not hold up against the evidence. Instead, we identify a new de-skilling mechanism: GenAI exerts a unique "top-squeeze" effect on the most highly educated workers, consistent with the levelling effect noted by others[20]. It simultaneously lowers the barriers for routine coding, writing, and analytical tasks, allowing lower-skilled workers to move into previously high-skilled roles. On the other hand, it directly encroaches on the core cognitive tasks traditionally embedded in high-skilled occupations, rapidly diluting their scarcity[50]. This explains why the supply of educated workers and their wages have decoupled in high-exposure neighborhoods. Accumulating academic credentials no longer guarantees rising returns; the logic of returns to education itself is being restructured[2,20,51]. 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, workers must invest more effort simply to hold their position[46]. 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[40,52].

Our findings resonate with recent global reports. A 2026 study from Anthropic concludes that actual AI adoption lags behind its theoretical potential but notes a 14% drop in hiring for the 22-25-year-old cohort[53]. Our spatial micro-data reveal that this decline is concentrated in the high-exposure core neighborhoods, pointing to strong spatial heterogeneity in AI's youth-employment effects. Similarly, the IMF(2026)[54] observes that AI skills command a wage premium yet simultaneously reduces entry-level hiring by 3.6%. Our identification of the "high-skill trap" provides an intra-urban mechanism for that macro pattern: AI does not eliminate jobs wholesale; rather, it lowers the barriers to cognitive tasks, triggering excessive competition in spatially clustered cores.

The policy implications are clear. In responding to GenAI, policymakers should look beyond computing infrastructure or generic skills training. We highlight three directions. First, to loosen the technological monopoly of urban cores, AI demonstration zones could be established in municipal sub-centers and suburban new towns, accompanied by lower access cost for digital infrastructure. Second, higher education should pivot away from cultivating "codifiable" standard

cognitive skills—which AI replicates readily—toward complex collaboration, critical reasoning, and cross-cultural communication, capabilities that are harder to automate[55]. Third, labor protection policies such as re-employment transition funds and skill-transformation subsidies targeted at highly exposed occupations can prevent the wage erosion of the "high-skill trap" from cascading into talent outflows.

This study has limitations. The data come from Beijing alone; caution is needed when generalizing to other global cities or smaller Chinese cities. Online recruitment data capture hiring flows rather than the stock of employed workers and may over-represent volatility[56]. Finally, because GenAI has been widely deployed for only a short period, longer-run productivity effects—including the possible emergence of entirely new occupational categories—may not yet be visible. Tracking these evolutionary dynamics over longer horizons is a priority for future work.

Methods

Study area

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 disruptions offers an early signal of the dynamics that may soon appear in other leading AI cities—San Francisco, London, and Singapore among them[31,32,34]. 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[33]. 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[36,57].

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 spatially matched the remaining postings to Beijing's neighborhoods (regulatory planning units). The final dataset contains 4,995,615 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 normalized log-transformed count (Job Index) to measure cross-neighborhood variation in employment intensity. (2) Average monthly wage. All compensation components—annual salary, 13th-month pay, and bonuses— are converted to a monthly equivalent, winsorized at the 0.1% tails (excluding values below 1,000 or above 280,000 RMB), and averaged within each neighborhood. (3) Average educational requirement. We map each posting's education requirement to years of schooling (unrestricted = 0, 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 talent threshold and degree of industrial sophistication.

To control for confounding influences, we include five categories of neighborhood-level characteristics: living convenience, development intensity, environmental quality, economic activity, and population distribution. 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 Table S2.

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 using 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 500 matches and had them manually checked by experts. 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 SOTA LLMs—ChatGPT-4o, Gemini-2.5-pro, Claude-3.5-Sonnet, GLM-4 and Deepseek-R1—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)[58], 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 total) and used the mean as the final exposure score (Supplementary Text S4). The occupational-level GenAI exposure index was then obtained through aggregation. Assessments were completed in June 2025. Pairwise correlations among model assessments range from 0.66 to 0.86, 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:

$$\ln(Wage)_{it} = \alpha + \delta(GenAI_{i,2018} \times Post_t) + \sum_k \gamma_k \, (Confounder_{k,i} \times Post_t) + \sum_j \beta_j \, X_{j,it} + \mu_i + \lambda_t + \varepsilon_{it}$$

(1)

where $GenAI_{i,2018}$ is the standardized pre-determined GenAI exposure measured in 2018; $Post_t = \mathbf{1}[t \geq 2023]$ is the post-treatment dummy; $Confounder_{k,i}$ denotes the three concurrent-shock exposures based on the 2018 industrial structure; $X_{j,it}$ includes ln(population), nightlight intensity, 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 1,383 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 employment share in that set:

$$Confounder_{k,i} = \frac{\sum_{j \in \mathcal{J}_k} N_{ij,2018}}{\sum_{j=1}^{18} N_{ij,2018}}$$

(2)

where $N_{ij,2018}$ is the number of recruitment postings for neighborhood i in industry j in 2018. Technology-regulation exposure uses the employment share of the Information Technology and Software industry (Industry Code I); COVID-19-recovery exposure uses the combined share of Accommodation and Catering (H) and Transportation (G); real-estate-adjustment exposure uses the combined share of Real Estate (K) and Construction (E). All three exposures are standardized and interacted with $\mathrm{Post_t}$ before inclusion.

Event-study specification

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

$$\ln(Wage)_{it} = \alpha + \sum_{k \neq 2022} \beta_k \left(GenAI_{i,2018} \times \mathbf{1}[t = k]\right) + \sum_{m} \gamma_m \left(Confounder_{m,i} \times Post_t\right) + \sum_{j} \phi_j X_{j,it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (3)$$

with 2022 as the base period. The identifying assumption requires that all pre-treatment coefficients satisfy $\beta_{2020} = \beta_{2021} = 0$, 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 1,383 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.

Bartik shift-share IV

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

$$\overline{\mathrm{GenAI}}_{\mathrm{j},-\mathrm{i}} = \frac{\sum_{\mathrm{k}\neq\mathrm{i}} \mathrm{N}_{\mathrm{kj}} \cdot \mathrm{GenAI}_{\mathrm{k}}}{\sum_{\mathrm{k}\neq\mathrm{i}} \mathrm{N}_{\mathrm{kj}}} \quad (4)$$

where N_{kj} is the number of postings in neighborhood k for industry j and $GenAI_k$ is the GenAI exposure of neighborhood k. Interacting neighborhood i's 2018 industry-employment share $s_{ij,2018}$ with the industry-level exposure and summing over industries yields:

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

This IV is driven solely by historical industrial structure and national-level industry-specific 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 ($\Delta = \overline{Post} - \overline{Pre}$) for each variable and estimate cross-sectional two-stage least squares (2SLS):

First stage:

$$\Delta GenAI_i = \pi_0 + \pi_1 Bartik_i + \gamma' Z_i + u_i \quad (6)$$

Second stage:

$$\Delta \ln(Wage)_i = \alpha + \beta_{IV} \Delta \widehat{GenAI}_i + \gamma' Z_i + \varepsilon_i \quad (7)$$

Reduced form:

$$\Delta \ln(Wage)_i = \alpha + \beta_{RF} Bartik_i + \gamma' Z_i + \varepsilon_i \quad (8)$$

where Z_i includes population change, nightlight change and the three concurrent-shock exposures. Instrument strength is evaluated via the first-stage F-statistic, benchmarked against the critical value of 10 of Stock and Yogo (2002)[44].

Triple-DID mechanism tests

To test the de-skilling and crowding mechanisms within a causal framework, we extend the DID model into a triple-DID design. Taking the de-skilling channel as an example:

$$\ln(Wage)_{it} = \alpha + \delta_1(GenAI_{i,2018} \times Post_t) + \delta_2(M_{i,2018} \times Post_t) + \delta_3(GenAI_{i,2018} \times M_{i,2018} \times Post_t) + \Gamma + \varepsilon_{it} \quad (9)$$

where $M_{i,2018}$ is the 2018 moderator (education level or recruitment heat), and Γ collects concurrent-shock controls, baseline controls and neighborhood and year fixed effects. The key parameter δ_3 measures heterogeneity in the GenAI wage penalty—specifically, whether it is exacerbated by higher education (or greater market crowding). All moderators use 2018 pre-determined values to preserve exogeneity.

Fixed-effects model

To examine the impact of GenAI exposure on labor-market compensation and employment structure, we estimate a two-way fixed-effects model that controls for time-invariant individual heterogeneity and individual-invariant time trends, mitigating omitted-variable bias. To uncover the de-skilling and crowding mechanisms, we add interaction terms between GenAI exposure and education level, and between GenAI exposure and job-market heat, respectively:

$$\ln(Wage)_{it} = \alpha_0 + \alpha_1 GenAI_{it} + \alpha_2 Educ_{it} + \alpha_3(GenAI_{it} \times Educ_{it}) + \sum_k \gamma_k X_{k,it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (10)$$

$$\ln(\mathrm{Wage})_{it} = \alpha_0 + \alpha_1 \mathrm{GenAI}_{it} + \alpha_2 \mathrm{Crowd}_{it} + \alpha_3(\mathrm{GenAI}_{it} \times \mathrm{Crowd}_{it}) + \sum_k \gamma_k X_{kit} + \mu_i + \lambda_t + \varepsilon_{it} \quad (11)$$

where $\ln(\mathrm{Wage})_{it}$ is the natural logarithm of the wage level; GenAI_{it} is standardized GenAI exposure; Educ_{it} is the standardized education level; Crowd_{it} is standardized recruitment heat (market crowding); $\mathrm{GenAI}_{it} \times \mathrm{Educ}_{it}$ tests the de-skilling effect; $\mathrm{GenAI}_{it} \times \mathrm{Crowd}_{it}$ tests the crowding effect; $X_{k,it}$ is a vector of controls covering living convenience, development intensity, environmental quality, economic activity and population distribution; and μ_i and λ_t are entity and time fixed effects.

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 coefficients were re-estimated. Repeating this procedure 500 times yielded an empirical distribution of placebo coefficients. The fact that the true interaction coefficient lies far in the tail of this distribution indicates that the observed de-skilling and crowding effects are not statistical noise (Supplementary Fig. S6 and Supplementary Text S8).

Data availability

The datasets generated and analyzed during the current study are not publicly available because they contain location-identifiable spatial information that raises privacy concerns, but they are available from the corresponding author on reasonable request.

Code availability

Code underlying this study is not publicly available but may be made available to qualified researchers on reasonable request from the corresponding author.

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

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

6 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).

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

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

9 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).
10 Autor, D. H. in *AEA Papers and Proceedings*. 1–32 (American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203).
11 Moretti, E. The effect of high-tech clusters on the productivity of top inventors. *American Economic Review* **111**, 3328–3375 (2021).
12 Vinuesa, R. *et al.* The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature communications* **11**, 233 (2020).
13 Capraro, V. *et al.* The impact of generative artificial intelligence on socioeconomic inequalities and policy making. *PNAS nexus* **3**, pgae191 (2024).
14 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).
15 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).
16 Moretti, E. *The new geography of jobs*. (Houghton Mifflin Harcourt, 2012).
17 Bárány, Z. L. & Siegel, C. Job polarization and structural change. *American Economic Journal: Macroeconomics* **10**, 57–89 (2018).
18 Acemoglu, D. & Autor, D. in *Handbook of labor economics* Vol. 4 1043–1171 (Elsevier, 2011).
19 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).
20 Brynjolfsson, E., Li, D. & Raymond, L. Generative AI at work. *The Quarterly Journal of Economics* **140**, 889–942 (2025).
21 Noy, S. & Zhang, W. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **381**, 187–192 (2023).
22 Acemoglu, D. The simple macroeconomics of AI. *Economic Policy* **40**, 13–58 (2025).
23 Kok, S. & Weel, B. t. Cities, tasks, and skills. *Journal of Regional Science* **54**, 856–892 (2014).
24 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).
25 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).
26 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).
27 Acemoglu, D. & Restrepo, P. Tasks, automation, and the rise in US wage inequality. *Econometrica* **90**, 1973–2016 (2022).
28 Gmyrek, P., Berg, J. & Bescond, D. Generative AI and jobs: A global analysis of potential effects on job quantity and quality. (2023).
29 Demombynes, G., Langbein, J. & Weber, M. *The exposure of workers to Artificial Intelligence in low-and middle-income countries*. (World Bank, 2025).
30 Castells, M. *The rise of the network society*. (John wiley & sons, 2011).
31 You, X. China intensifies push to become world leader in tech and AI. *Nature* **651**, 859–860 (2026).

32 Wu, F. *et al.* Towards a new generation of artificial intelligence in China. *Nature machine intelligence* **2**, 312–316 (2020).

33 Tian, G., Wu, J. & Yang, Z. Spatial pattern of urban functions in the Beijing metropolitan region. *Habitat International* **34**, 249–255 (2010).

34 Lai, Y. & Zhao, H. Comparative analysis of smart city scientific research trends in the USA and China. *Nature Cities* **2**, 875–883 (2025).

35 Storper, M. & Venables, A. J. Buzz: face-to-face contact and the urban economy. *Journal of economic geography* **4**, 351–370 (2004).

36 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).

37 Acemoglu, D. & Restrepo, P. Robots and jobs: Evidence from US labor markets. *Journal of political economy* **128**, 2188–2244 (2020).

38 Goldfarb, A. & Tucker, C. Digital economics. *Journal of economic literature* **57**, 3–43 (2019).

39 Kraay, A. & McKenzie, D. Do poverty traps exist? Assessing the evidence. *Journal of Economic Perspectives* **28**, 127–148 (2014).

40 Acemoglu, D. & Restrepo, P. Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives* **33**, 3–30 (2019).

41 Hartley, J., Jolevski, F., Melo, V. & Moore, B. The labor market effects of generative artificial intelligence. *Available at SSRN* (2024).

42 Bergmann, B. R. Occupational segregation, wages and profits when employers discriminate by race or sex. *Eastern economic journal* **1**, 103–110 (1974).

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. in *Inequality in the 21st Century* 49–54 (Routledge, 2018).

46 Dell'Acqua, F. *et al.* Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard business school technology & operations mgt. Unit working paper* (2023).

47 Glaeser, E. *Triumph of the city: How our greatest invention makes us richer, smarter, greener, healthier, and happier*. (Penguin, 2012).

48 Polanyi, M. in *Knowledge in organisations* 135–146 (Routledge, 2009).

49 Autor, D., Chin, C., Salomons, A. & Seegmiller, B. New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics* **139**, 1399–1465 (2024).

50 Acemoglu, D. & Restrepo, P. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review* **108**, 1488–1542 (2018).

51 Deming, D. J. The growing importance of decision-making on the job. (National Bureau of Economic Research, 2021).

52 Lazear, E. P. & Oyer, P. Internal and external labor markets: a personnel economics approach. *Labour economics* **11**, 527–554 (2004).

53 Massenkoff, M. & McCrory, P. Labor market impacts of AI: A new measure and early evidence. *Anthropic Research* **5** (2026).
54 Jaumotte, M. F. *et al.* Bridging Skill Gaps for the Future: New Jobs Creation in the AI Age. (International Monetary Fund, 2026).
55 Deming, D. J. The growing importance of social skills in the labor market. *The quarterly journal of economics* **132**, 1593–1640 (2017).
56 Hershbein, B. & Kahn, L. B. Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review* **108**, 1737–1772 (2018).
57 Graham, S. Bridging urban digital divides? Urban polarisation and information and communications technologies (ICTs). *Urban studies* **39**, 33–56 (2002).
58 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).

Supplementary Files

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

- SupplementaryInformation.docx