The AI Shock in Multidimensional Skill Space

By Anonymous*

*Generative AI does not advance along a one-dimensional high-skill or routine-task margin. It expands unevenly across bundles of text, communication, clerical, coding, analytic, and managerial skills. This paper develops a multidimensional skill-space model in which jobs are skill bundles, AI is a time-varying capability set, and exposure is the proximity between a job's skill vector and the AI penetration frontier. Using 6.73 million online job postings by Chinese listed firms from 2016 to 2025, we construct a time-varying AI technology exposure index (AITE) from O*NET detailed work activities, multilingual semantic matching, LLM task scores, model Elo ratings, and market-share weights. Double machine learning and Bartik shift-share instruments show that AI exposure sharply reduces labor demand, with realized hires falling much more than posted vacancies. Wage effects are shaped by selection: exposed jobs that continue to post wages are positively selected, and correction turns the apparent wage premium negative. Quantile estimates reveal a monotone gradient from positive effects at the bottom to negative effects at the top, consistent with skill-bundle rather than one-dimensional skill sorting. Finally, graph attention networks show that the estimated direct effect largely disappears after controlling for neighboring exposure in the skill network. AI's labor-market impact is therefore best understood as network propagation across adjacent skill bundles.*
JEL: J23, J24, O33, C14, C45
Keywords: Artificial intelligence, labor demand, skill space, double machine learning, graph neural networks

Generative artificial intelligence has moved automation from the factory floor into the office. Large language and multimodal models now perform tasks that were previously treated as hard-to-automate human cognitive work: drafting legal documents, summarizing financial reports, writing code, producing marketing copy, screening resumes, and responding to customers. In China, the rapid diffusion of models such as Wenxin Yiyan, Tongyi Qianwen, Doubao, and DeepSeek created a short and unusually sharp window in which labor demand could react to a new class of general-purpose cognitive technology.

The standard economics vocabulary is useful but incomplete for this episode. Skill-biased technical change emphasizes complementarity with high-skill labor (Goldin and Katz, 2008). Routine-biased technical change emphasizes substitution of codifiable middle-skill tasks (Autor, Dorn and Hanson, 2015). Generative

* Prepared for an NBER-style working paper submission. Author information is omitted in this draft.

AI cuts across both categories. It does not simply help high-skill workers or replace routine workers. Its frontier expands first through particular bundles of language, communication, document processing, sales, programming, and data-analysis tasks. The relevant unit of exposure is therefore not a worker's location on a single skill ladder. It is the location of a job's skill bundle in a high-dimensional skill space.

This paper builds and estimates that idea. The theoretical object is a cognitive penetration frontier in multidimensional skill space. Each job is represented by a vector of detailed work activities. AI is a time-varying capability set rather than a scalar productivity shock. A job becomes exposed when its skill bundle lies inside, or close to, the region in which AI can perform the bundle at lower effective cost than human labor. This formulation generates four empirical predictions. First, demand falls most in exposed skill clusters, and realized hires should fall more than posted vacancies because firms can maintain search or branding postings while cutting actual hiring. Second, wage effects should vary by the skill composition of wage quantiles, not by a one-dimensional high-low skill ranking. Third, observed wages are selected because exposed jobs that keep reporting explicit wages are not representative of the full exposed population. Fourth, if skill bundles form local neighborhoods, exposure should propagate across the skill network; omitting neighboring exposure overstates the direct effect of a job's own exposure.

We test these predictions using 6,734,619 job postings issued by Chinese listed firms from 2016 through 2025. The empirical exposure measure, AITE, is built from the bottom up. We first score 2,087 O*NET detailed work activities (DWAs) for AI substitutability using large-language-model assessments. We then match Chinese job descriptions to DWAs with multilingual semantic embeddings and retain each posting's top ten matched activities. Finally, we weight the task score by time-varying AI capability, combining model Elo ratings and market-share weights. This creates a job-by-year measure of exposure that changes both with the task content of jobs and with the AI capability set available in each year.

The empirical results are consistent with the skill-space model. Baseline double machine learning (DML) estimates imply that AITE reduces log posted vacancies by 1.330 and log realized hires by 4.767. The effect on hires is roughly 3.6 times the effect on postings, matching the recruiting-funnel prediction. Bartik estimates using global compute shocks interacted with local baseline shares of compute-intensive jobs confirm the negative effect on realized hiring. Wage results initially look positive or weakly negative, but this is a selection artifact: exposed jobs with explicit wage postings are positively selected. Heckman correction and Lee bounds move the wage effect negative.

The distributional results sharpen the mechanism. Quantile estimates range from positive at the 10th percentile to strongly negative at the 90th percentile. This is not the U-shaped polarization of routine-biased technical change. It is a monotone gradient: low-wage jobs temporarily absorb displaced workers, while

upper and upper-middle wage jobs are closer to AI-exposed standardized cognitive skill clusters. Finally, graph attention networks (GATs) trained on the DWA skill-similarity network show that direct exposure estimates collapse once neighboring exposure is included. For realized hires, the direct coefficient falls from -5.266 to -0.971 and becomes statistically insignificant. This pattern is exactly what the model predicts if labor-market adjustment is a network propagation process rather than a collection of isolated job-level substitution events.

I. Literature Review

The task framework argues that automation affects labor demand through three forces. First, substitution changes the task content of production by allowing capital to take over tasks previously performed by labor. Second, automation can make task allocation across factors more flexible, raising productivity and increasing labor demand in nonautomated tasks. Third, automation may create new tasks in which labor has comparative advantage, generating a reinstatement effect. The net effect of automation on labor demand therefore depends on the balance among substitution, productivity, and reinstatement effects (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018). Acemoglu and Restrepo (2022) further shows that changes in group wages are directly linked to task displacement; 50 to 70 percent of the change in the US wage structure over the past four decades can be attributed to falling relative wages of workers specializing in routine tasks in rapidly automating industries. Autor and Salomons (2018) documents that the labor-share displacement effect of productivity growth was almost absent in the 1970s in eighteen OECD countries, but became increasingly salient and peaked after the 2000s, consistent with automation gradually shifting from labor-augmenting to labor-substituting.

The task framework motivates a systematic measurement of how occupations are exposed to AI. Brynjolfsson, Mitchell and Rock (2018) decomposes work activities in O*NET and scores the suitability of individual activities for machine learning, constructing a suitability-for-machine-learning index. Webb (2020) measures the semantic overlap between AI patent titles and occupation texts. Felten, Raj and Seamans (2021) maps AI application areas to underlying abilities through a large-scale crowd-evaluation system and constructs an occupational AI exposure index. Eloundou et al. (2023) combines expert ratings and LLM-assisted scoring to show that a broad range of white-collar jobs face productivity thresholds induced by large models, including complex communication and execution tasks. These shocks may be especially concentrated among traditional cognitive workers near the top of the income distribution.

Exposure indices, however, do not by themselves imply net employment or welfare effects. Under the task framework, AI can displace labor, but it can also expand demand through reinstatement and complementarity. Worker adjustment also depends on nontechnical factors. Pizzinelli et al. (2023) extends the AI occupational exposure index by incorporating potential complementarity and shows

that more educated, higher-income workers are more likely to be in highly exposed occupations, while those occupations also tend to have greater complementarity. Gathmann and Schönberg (2007) introduces task-specific human capital and measures skill distance using cosine distances between occupational task vectors. Manning and Aguirre (2026) combines transferability, liquid assets, geographic density, and age to construct an adaptability index for 356 occupations, finding that AI exposure is positively correlated with adaptability.

Recent Chinese evidence also studies AI and the labor market. Wang and Dong (2020) uses industry robot adoption and listed-firm data to show that a one-percent increase in robot penetration reduces firm labor demand by 0.18 percent. Chao and Zhou (2021) finds that AI lowers the labor income share in provincial panels. Yao et al. (2024) constructs firm-level AI measures from annual reports and patents and shows that AI raises productivity by reducing routine low-skill labor demand and increasing nonroutine high-skill labor demand. Zhang et al. (2025) uses online job postings and job-task descriptions to construct large-language-model exposure and estimates its effects on labor demand, educational requirements, and wages.

Methodologically, double machine learning (DML) eliminates regularization bias through orthogonalized objectives and cross-fitting (Chernozhukov et al., 2018). XGBoost is a natural first-stage learner for tabular data because it captures nonlinear relationships effectively (Chen and Guestrin, 2016; Grinsztajn, Oyallon and Varoquaux, 2022). For network causal inference, Leung and Loupos (2025) develops a theoretical foundation for graph neural networks under approximate neighborhood interference. When interference decays with network distance, a shallow GNN can approximate low-dimensional nuisance functions. This provides the basis for using graph attention networks to control network confounding in a skill network.

A. Evaluation of the Literature and Theoretical Need

The existing literature leaves room for extension on three margins. First, measurement remains largely static. Existing AI exposure indices, including SML (Brynjolfsson, Mitchell and Rock, 2018), AIOE (Felten, Raj and Seamans, 2021), and GPT exposure (Eloundou et al., 2023), typically score tasks once and treat AI capability as fixed. Since large-model capability has evolved continuously since 2022, incorporating time variation can improve measurement.

Second, identification can be strengthened. Many empirical studies rely on OLS or difference-in-differences designs. In a high-dimensional confounding environment, nonparametric causal methods such as DML improve robustness.

Third, the network dimension is underdeveloped. Existing studies generally impose the stable unit treatment value assumption (SUTVA), implicitly assuming that AI exposure in one job does not affect other jobs. Relaxing SUTVA and introducing skill-network spillovers adds an important dimension to the causal identification problem (Manski, 2013).

These three gaps motivate the next section. Building on the general-equilibrium task logic in Acemoglu and Restrepo (2018, 2019), we develop a multidimensional skill-space cognitive penetration frontier for generative AI and derive testable hypotheses for aggregate substitution, skill-composition distributional effects, wage selection, and network propagation.

II. A Skill-Space Model of AI Penetration

A. Jobs, AI Capabilities, and the Penetration Region

Let each job i be represented by a skill vector

$$\mathbf{x}_i = \sum_{k \in \mathcal{T}_i} \omega_{ik} \mathbf{e}_k \in \mathbb{R}^K, \tag{1}$$

where $\mathcal{T}_i$ is the set of detailed work activities matched to job i, $\mathbf{e}_k$ is the embedding of activity k, and ω_{ik} is the semantic weight of that activity in the posting. The dimensions of $\mathbf{x}_i$ include cognitive processing, communication, coordination, creativity, programming, management, and physical or on-site operation.

AI at time t is a capability set $\mathcal{A}_t$, not a scalar. Let $a_t(\mathbf{x})$ denote AI's effective productivity on skill bundle $\mathbf{x}$, and let $c_H(\mathbf{x})$ denote the human effective cost of performing that bundle. The AI penetration region is

$$\mathcal{P}_t = \{\mathbf{x} \in \mathcal{X} : a_t(\mathbf{x}) \geq c_H(\mathbf{x})\}. \tag{2}$$

Technological progress is an expansion of $\mathcal{P}_t$ through skill space. This expansion is nonuniform: document summarization, sales scripting, customer response, and structured data work may enter the region quickly, while on-site physical work or high-liability judgment may remain outside.

Empirically, exposure is the match between a job's skill bundle and the AI capability set:

$$D_{it} = g(\mathbf{x}_i, \mathcal{A}_t). \tag{3}$$

The AITE index described below is the empirical counterpart of D_{it}.

B. Labor Demand and the Recruiting Funnel

Let L_t^C be human labor demand for cognitive skill bundles outside the AI penetration region:

$$L_t^C = \int_{\mathcal{X} \setminus \mathcal{P}_t} l^c(\mathbf{x}) dF(\mathbf{x}). \tag{4}$$

If the penetration region expands from $\mathcal{P}_t$ to $\mathcal{P}_{t+1}$, with $\Delta\mathcal{P}_{t+1} = \mathcal{P}_{t+1} \setminus \mathcal{P}_t$, the displaced labor demand is

$$\Delta L^C = -\int_{\Delta\mathcal{P}_{t+1}} l^c(\mathbf{x}) dF(\mathbf{x}) < 0. \tag{5}$$

PROPOSITION 1 (Skill-bundle displacement): *If AI capability expansion creates a positive-measure addition to $\mathcal{P}_t$, and new task creation does not offset the displaced skill bundles, labor demand falls in exposed skill clusters. If cognitive and on-site tasks are complements, aggregate labor demand also falls.*

Observed vacancies need not equal realized demand. Let V_{ct} denote posted vacancies and H_{ct} realized hires in city c and year t. Suppose firms post

$$V_{ct} = H^*_{ct} + B_{ct}, \tag{6}$$

where H^*_{ct} is desired hiring and B_{ct} is a posting component related to branding, option value, or search. Realized hiring satisfies

$$H_{ct} = q(D_{ct}) H^*_{ct}, \qquad q'(D_{ct}) < 0. \tag{7}$$

Firms may keep vacancies visible while reducing actual hiring. Therefore the hiring response should exceed the posting response in absolute value.

C. Selection, Quantiles, and Skill-Neighborhood Propagation

Wage observations are selected. Let $S_i = 1$ indicate that job i posts an explicit wage rather than a negotiable or missing wage. We write

$$S_i = \mathbf{1}\{m(\mathbf{x}_i, D_{it}, Z_i) \geq u_i\}. \tag{8}$$

If high-exposure jobs that continue to report wages are positively selected, then

$$E[w_i \mid D_{it}, S_i = 1] > E[w_i \mid D_{it}]. \tag{9}$$

The observed wage effect can therefore be positive even when the true effect on the full exposed population is negative.

PROPOSITION 2 (Wage selection): *If explicit wage reporting among exposed jobs is positively selected on skill rents or bargaining strength, the observed mean wage effect is upward biased. Correcting selection can turn an apparent wage premium into a negative wage effect.*

Different wage quantiles correspond to different skill-bundle distributions. Let $F_\tau(\mathbf{x})$ denote the distribution of skill bundles behind wage quantile τ. The quan-

tile effect is

$$\beta(\tau) = \int_{\mathcal{X}} b(\mathbf{x}) dF_{\tau}(\mathbf{x}), \tag{10}$$

where $b(\mathbf{x})$ combines substitution, reallocation, and selection. If upper quantiles contain more standardized cognitive, text, sales, clerical, and computer-related bundles, then upper quantiles are closer to $\mathcal{P}_t$ and should experience stronger negative effects.

PROPOSITION 3 (Skill-composition quantile effects): *AI affects wage quantiles through their underlying skill bundles, not through a scalar skill rank. If exposure intensity rises with a quantile's concentration in standardized cognitive bundles, the treatment effect declines monotonically across quantiles.*

Finally, skill bundles form neighborhoods. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a skill graph with edge weights

$$w_{ij} = \kappa(d(\mathbf{x}_i, \mathbf{x}_j)), \tag{11}$$

where $d(\cdot, \cdot)$ is semantic distance. Labor-demand shocks satisfy

$$\Delta L_i = \theta D_i + \delta \sum_{k \in \mathcal{N}(i)} w_{ik} D_k + \epsilon_i. \tag{12}$$

If $\delta < 0$ and a job's exposure is positively correlated with neighboring exposure, omitting the second term makes the direct coefficient too negative.

PROPOSITION 4 (Network omitted-variable bias): *In the true model* (12)*, a regression that omits neighboring exposure has probability limit*

$$\operatorname{plim} \hat{\theta} = \theta + \delta \frac{\operatorname{Cov}(D_i, N_i)}{\operatorname{Var}(D_i)}, \qquad N_i = \sum_{k \in \mathcal{N}(i)} w_{ik} D_k. \tag{13}$$

If $\delta < 0$ *and* $\operatorname{Cov}(D_i, N_i) > 0$*, the isolated-job estimate is more negative than the true direct effect. Adding GNN embeddings or neighboring exposure should shrink the direct effect toward zero.*

III. Data and Measurement

The analysis combines seven sources: online job postings by listed firms, O*NET detailed work activities, LLM task exposure scores, LMSYS Chatbot Arena Elo ratings, large-model market shares, global AI compute shocks, and city- and firm-level controls. The main sample contains 6,734,619 postings from 2016 to 2025. The main panel is at the city-by-year level, with 404 cities and 3,232 observations.

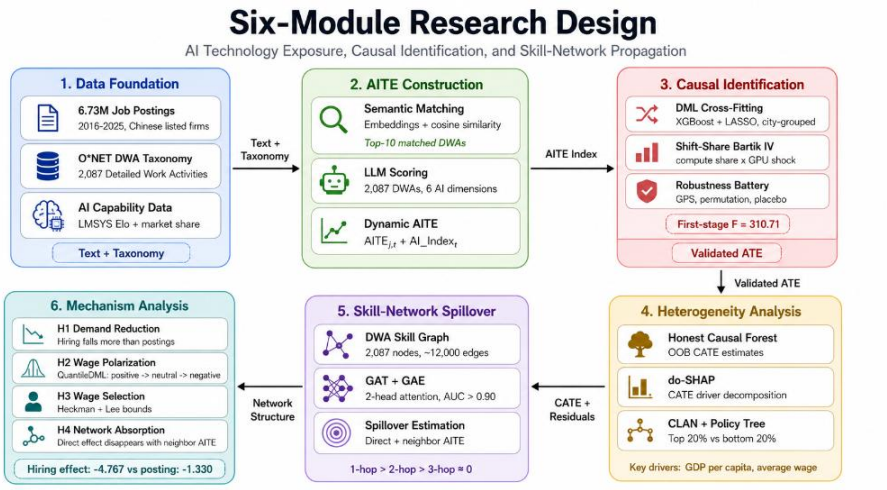


FIGURE 1. RESEARCH DESIGN

Note: The design links job text to O*NET detailed work activities, constructs dynamic AI exposure, estimates causal effects with DML and Bartik instruments, and uses GNN embeddings to model propagation in skill space.

We also use city-by-occupation-by-year, province-by-year, and firm-by-year panels for robustness and mechanism tests.

A. Constructing AITE

The AITE construction has three steps. First, each of 2,087 O*NET detailed work activities receives an LLM exposure score $S_k^{LLM} \in [0, 5]$. Second, Chinese job descriptions are matched to O*NET activities using multilingual semantic embeddings; each posting keeps the top ten matched DWAs and their normalized similarities. Third, task exposure is scaled by the time-varying AI capability index

$$AI_Index_t = \sum_m \omega_{m,t} \frac{ELO_{m,t}}{1000}, \tag{14}$$

where $\omega_{m,t}$ is model m's market share and $ELO_{m,t}$ is its Chatbot Arena rating. The posting-level exposure measure is

$$AITE_{j,t} = \sum_{k \in Top10_j} sim_{j,k} \cdot (S_k^{LLM}/5) \cdot AI_Index_t. \tag{15}$$

For 2016–2022, generative-model market shares are set to zero, so AITE is zero before the large-model diffusion period.

Table 1—Core Variables

Block	Variable	Definition
Outcomes	$\ln(Postings)_{ct}$	Log number of posted vacancies in city c and year t.
	$\ln(Hires)_{ct}$	Log realized recruitment headcount.
	Wage CV_{ct}	Coefficient of variation of posted wages.
Treatment	AITE_{ct}	City-year average AI technology exposure, weighted from posting-level AITE.
Instrument	Bartik IV_{ct}	Global AI compute shock interacted with local baseline share of compute-intensive jobs.
Controls	X_{ct}	City macro controls and listed-firm financial controls.

Note: The treatment maps the theoretical object $D_{it} = g(\mathbf{x}_i, \mathcal{A}_t)$ to data. The skill vector $\mathbf{x}_i$ comes from job-text-to-DWA matching; the capability set $\mathcal{A}_t$ is approximated by model Elo and market-share weights.

IV. Empirical Strategy

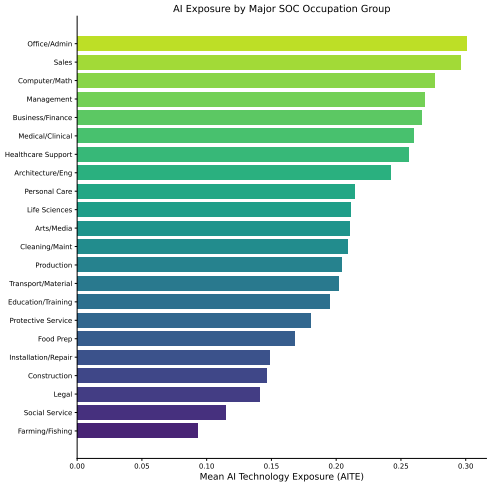


Figure 2. AI Exposure by Major SOC Group

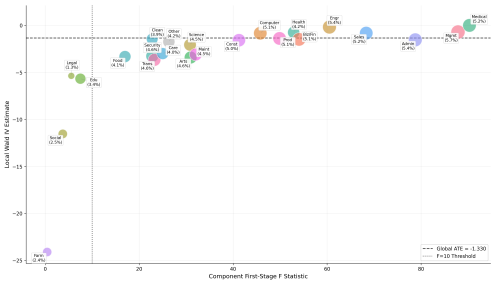


Figure 3. Rotemberg Weight Bubble Chart for IV Strategy

The main estimating equation is

$$Y_{ct} = \beta AITE_{ct} + X'_{ct}\Gamma + \lambda_c + \tau_t + \varepsilon_{ct}. \tag{16}$$

We report OLS, double machine learning, and Bartik IV estimates. The DML specification uses a partially linear model with orthogonalization and grouped cross-fitting by city. The first-stage nuisance functions are estimated with stacked XGBoost and LASSO learners. Standard errors are clustered at the city level.

The Bartik instrument is

$$Z_{ct} = s_{c,2016}^{compute} \times Shock_t^{GPU}, \tag{17}$$

where $s_{c,2016}^{compute}$ is the baseline city share of compute-intensive jobs and $Shock_t^{GPU}$ is the global AI compute shock proxied by NVIDIA data-center revenue. The identifying logic is that global compute supply shifts change the capability frontier, while pre-period local compute-job shares govern exposure to that shift. To enhance the exogenous variation of the instrument and decouple it from cryptocurrency mining cycles, we perform a residual orthogonalization on the Bartik-shift instrument. Specifically, we regress the NVIDIA data center revenue growth rate on the Bitcoin price and extract the residuals to construct $Shock_t^{GPU}$. This isolates the pure AI-driven compute shock, significantly strengthening the validity of our instrument.

V. Baseline Effects

The baseline pattern is the recruiting funnel in the data. Posted vacancies fall, but realized hiring falls much more. The DML coefficient on realized hires is 3.6 times as large as the coefficient on postings. Firms appear to maintain part of their vacancy posting behavior even as they sharply reduce actual hiring. This distinction is important for macroeconomic monitoring: vacancy postings alone understate the contraction in actual employment demand.

A. Robustness and Placebo Tests

Figures 4 and B1 summarize the multidimensional DML baseline estimates and robustness checks. The robustness exercises include GPS trimming, excluding pure AI algorithm-development jobs, a one-year lag of exposure, and two-way clustering. The estimates remain negative and statistically significant across specifications.

The pre-trend placebo test uses 2022 as the base year and the post-shock cross-sectional AITE distribution as fixed exposure intensity. The six pre-period coefficients from 2016 to 2021 are jointly insignificant ($F = 0.646$, $p = 0.693$), while the post-shock coefficient in 2024 is sharply negative (-7.158, $p < 0.01$). This rules out the alternative explanation that high-AITE cities already had differential employment trends before the AI shock.

Table 2—Baseline Effects of AI Exposure

	(1) OLS	(2) DML	(3) 2SLS
Panel A: Outcome = log posted vacancies			
AITE	-1.529^{***}	-1.330^{***}	-0.874
	(0.073)	(0.375)	(1.311)
Panel B: Outcome = log realized hires			
AITE	-5.559^{***}	-4.767^{***}	-6.335^{***}
	(0.101)	(0.751)	(0.874)
Panel C: Outcome = wage coefficient of variation			
AITE	-0.828^{***}	-0.615^{***}	–
	(0.027)	(0.230)	–
Controls	No	Yes	Yes
City and year fixed effects	Yes	Yes	Yes
Observations	3,232	3,232	3,232

Note: Standard errors clustered by city are in parentheses. The DML estimates use grouped cross-fitting by city. The 2SLS specification uses the Bartik compute-shock instrument. The first-stage F statistic in the SOC-level panel is 310.71. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

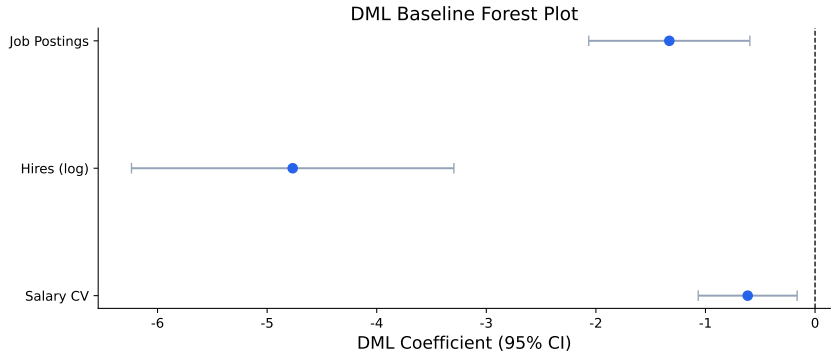


Figure 4. DML Baseline Forest Plot

Note: Each row reports a DML point estimate and 95 percent confidence interval for one outcome or specification.

Table 3—Hypothesis H1: Main Regressions and Robustness Checks

Specification	Estimate	95 percent CI	Observations
Panel A: Main and alternative outcome specifications			
DML baseline, log postings	-1.330^{***} (0.375)	$[-2.065, -0.595]$	3,232
DML baseline, log hires	-4.767^{***} (0.751)	$[-6.238, -3.296]$	3,232
Alternative hires imputation	-4.373^{***} (0.744)	$[-5.831, -2.914]$	3,232
Panel B: Sample and identification robustness			
GPS trimming, 3 percent	-0.316^{**} (0.148)	$[-0.605, -0.026]$	3,135
Two-way clustering, city and occupation	-1.279^{***} (0.445)	$[-2.151, -0.407]$	49,240
Excluding pure AI algorithm jobs	-0.552^{***} (0.073)	$[-0.695, -0.409]$	3,228
Lagged AITE, $AITE_{t-1}$	-1.735^{***} (0.393)	$[-2.506, -0.965]$	2,894
Panel C: Placebo permutation test			
5,000 permutations, empirical p-value	0.0006^{***}	–	–

Note: Standard errors are in parentheses. All specifications control for firm financial covariates and city and year fixed effects where applicable. $^{***}p < 0.01$, $^{**}p < 0.05$.

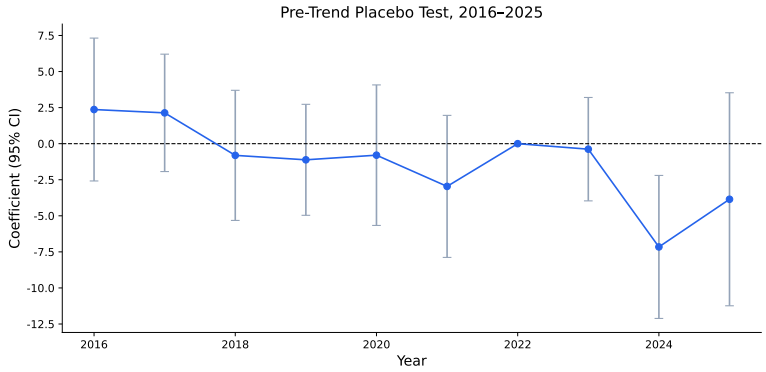


Figure 5. Pre-Trend Placebo Test, 2016–2025

VI. Mechanisms: Quantiles, Firm Buffering, and Wage Selection

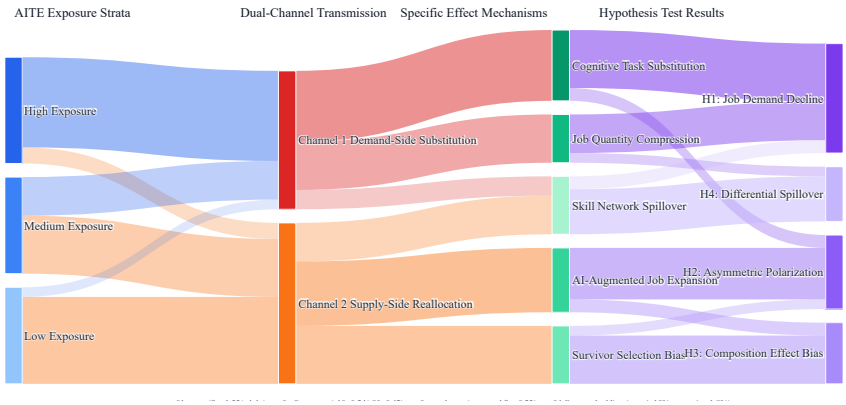


Figure 6. Two-Channel Mechanism in Multidimensional Skill Space

A. *Quantile Gradient*

The skill-space model predicts a monotone quantile gradient rather than classic U-shaped routine-task polarization. The estimates in Table 4 support this prediction. The effect is positive at the bottom, close to zero at the median, and negative at the top. This pattern is consistent with reallocation pressure at the bottom and substitution pressure in upper-middle and upper wage skill bundles.

B. *Wage Selection*

Chinese online recruiting platforms frequently allow firms to post wages as negotiable or undisclosed. This creates a selection problem. The Heckman first-stage Probit shows that high-AITE jobs are less likely to use negotiable wage postings and more likely to post explicit wages. Thus the observed wage sample among exposed jobs is positively selected. After correcting for selection with inverse Mills ratios, and after applying Lee bounds, the wage effect turns negative. The apparent wage premium is therefore not evidence that AI raises exposed workers' compensation. It is evidence that low-wage exposed jobs exit the observed wage sample first.

VII. Heterogeneity

The causal forest results show that the AI shock is not evenly distributed. The omnibus heterogeneity test rejects homogeneous treatment effects with $F = 254.286$ and $p < 0.001$. The largest negative conditional effects are concentrated

TABLE 4—MECHANISM TESTS: QUANTILES AND FIRM BUFFERING

Panel A: Wage quantile regressions			
Quantile τ	Estimate	p-value	Interpretation
0.10	$+0.236^{***}$ (0.023)	$< 10^{-23}$	Bottom absorption
0.25	$+0.140^{***}$ (0.019)	$< 10^{-13}$	Positive but declining
0.50	$+0.014$ (0.020)	0.506	Turning point
0.75	-0.137^{***} (0.029)	$< 10^{-6}$	Upper displacement
0.90	-0.424^{***} (0.045)	$< 10^{-21}$	Strong substitution
Panel B: Triple-difference test, AITE × R&D × Post			
Outcome	Estimate	p-value	R^2 (N)
Posted vacancies	$+0.675^{**}$ (0.331)	0.042	0.270 (43,336)
Average wage	-0.261^{**} (0.108)	0.016	0.081 (42,242)

Note: Panel A reports marginal effects at wage quantiles. Panel B tests whether R&D-intensive firms buffer employment losses after the AI shock. $^{***}p < 0.01$, $^{**}p < 0.05$.

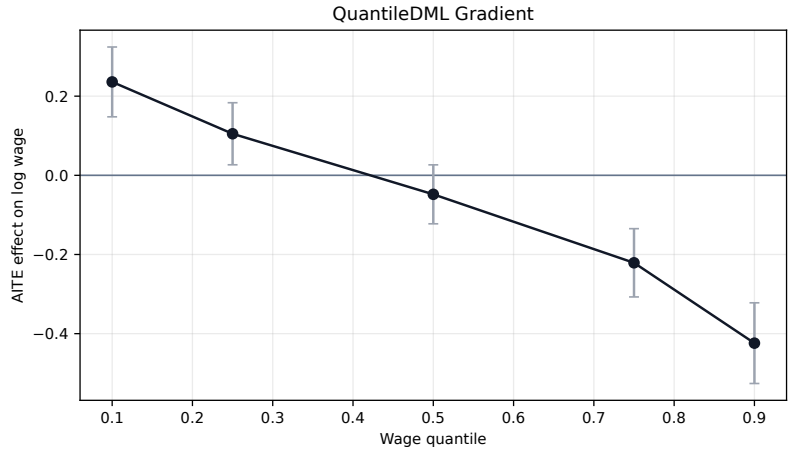


FIGURE 7. QUANTILE GRADIENT OF AI EXPOSURE EFFECTS

Note: The gradient is monotone from the lower tail to the upper tail. The pattern supports a skill-bundle interpretation of distributional effects.

in richer and higher-wage city-firm environments, where jobs are more likely to contain the standardized cognitive bundles that generative AI can perform.

Table 5—Causal Forest Heterogeneity

Panel A: Omnibus heterogeneity test			
$F = 254.286$, $p < 0.001$, $R^2 = 0.184$			
Panel B: Best linear projection of CATE			
Variable	Estimate	Variable	Estimate
Secondary-industry share	0.022^{***}	CapEx	-1.467^{***}
Tertiary-industry share	0.014^{***}	Cash-flow ratio	-0.458^{***}
GDP per capita	$2.7 \times 10^{-7***}$	Board size	-0.323^{***}
Population	0.0001^{***}	ROA	0.143^{***}
Firm size	-0.027^{***}		
Panel C: CLAN comparison, top versus bottom 20 percent CATE groups			
Variable	Gap	Cohen's d	Interpretation
Average wage	$36,775^{***}$ RMB	0.899	High-wage exposed
GDP per capita	$39,120^{***}$ RMB	0.816	Rich-city exposed
CapEx	-0.12^{***} sd	–	Consistent
Board size	-0.85^{***} sd	–	Consistent
Cash-flow ratio	$+0.08^{***}$ sd	–	Consistent
Hierarchy	-0.45^{***} sd	–	Consistent
Secondary-industry share	$+0.05^{***}$ sd	–	Consistent

Note: The forest uses a Bartik-IV doubly robust estimator. Panel C compares characteristics of the most and least negatively affected CATE groups. $^{***}p < 0.01$.

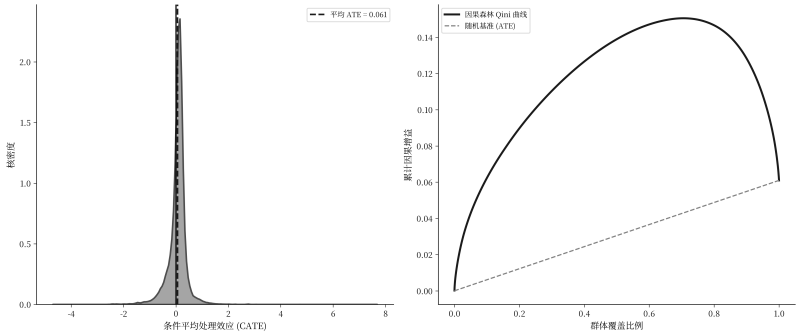


Figure 8. CATE Distribution and Qini Curve

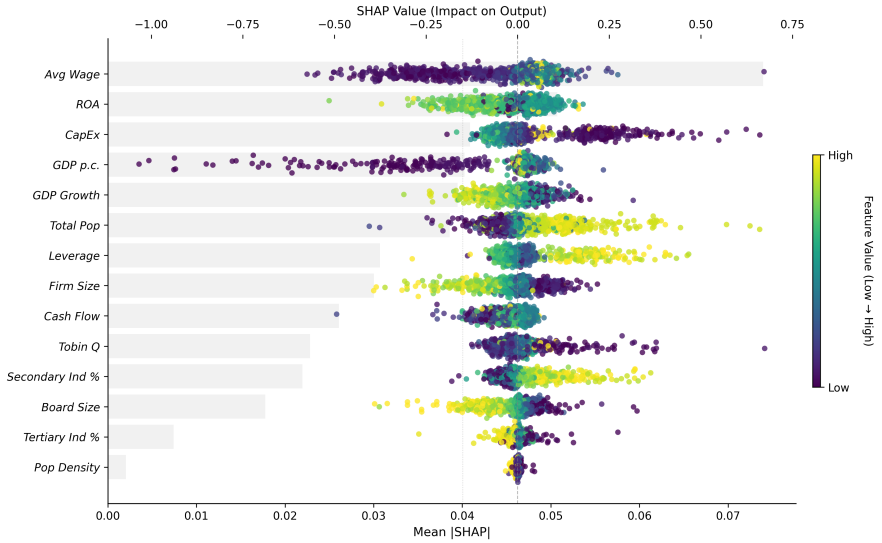


FIGURE 9. DO-SHAP DECOMPOSITION OF TREATMENT-EFFECT HETEROGENEITY

Note: The SHAP decomposition highlights GDP per capita and average wage as central drivers of heterogeneous effects.

VIII. Skill-Network Propagation

The final prediction is that the AI shock propagates through skill-space neighborhoods. We estimate a graph attention autoencoder on the DWA skill graph. Nodes are DWA skill bundles and edges are semantic similarities. The GAT attention weights estimate the model's w_{ij} rather than imposing an external spatial weight matrix.

The second-stage panel specification is

$$Y_{ct} = \theta_0 AITE_{ct} + \gamma \bar{z}_{ct} + \delta \overline{AITE}_{ct}^{neighbor} + X'_{ct}\beta + \lambda_c + \tau_t + \epsilon_{ct}, \tag{18}$$

where $\bar{z}_{ct}$ aggregates GAT embeddings to the city-year level and $\overline{AITE}_{ct}^{neighbor}$ measures neighboring exposure. Under approximate neighborhood interference, spillovers should decay with graph distance (Leung and Loupos, 2025).

IX. Conclusion

This paper argues that generative AI should be modeled as a multidimensional expansion in skill space. The core empirical object is not a one-dimensional task complexity score but the relationship between a job's skill bundle and the evolving AI capability set. This perspective unifies four facts: realized hiring falls much more than vacancy postings; wage effects are selected; distributional effects

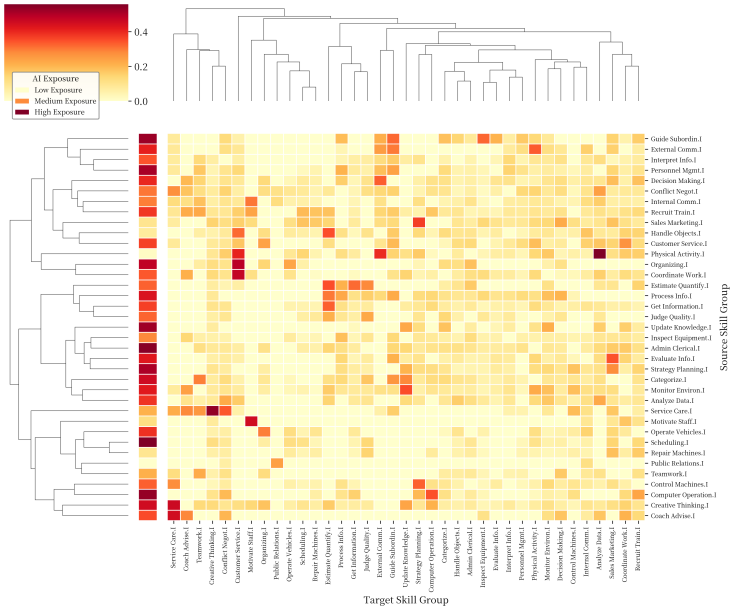


FIGURE 10. GAT ATTENTION WEIGHTS IN THE SKILL NETWORK

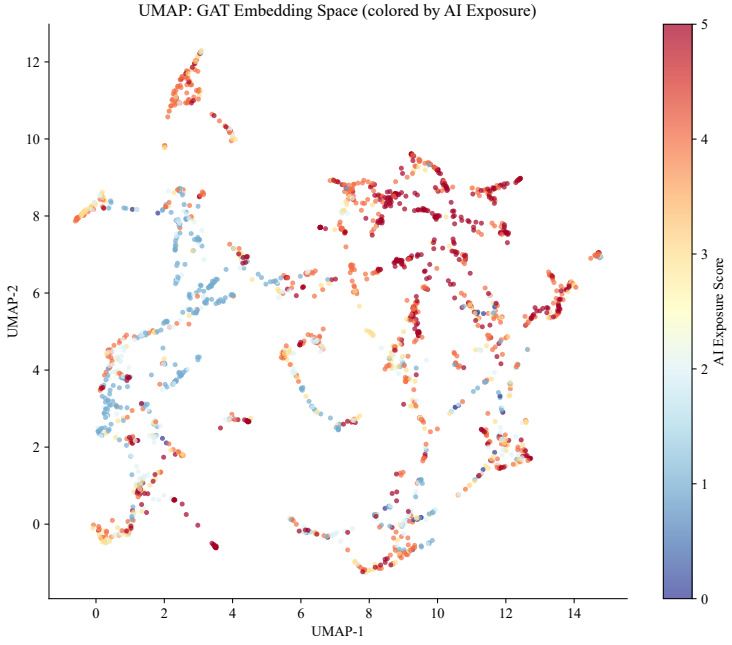


FIGURE 11. UMAP PROJECTION OF THE DWA SKILL-SPACE EMBEDDING

Table 6—Skill-Network Spillovers and Threshold Sensitivity

	(1) DML	(2) DML+GAT	(3) DML spillover	(4) OLS spillover
Panel A: Outcome = log posted vacancies				
AITE	-1.177^{***}	-1.263^{***}	-1.536	-0.736^{**}
	(0.297)	(0.295)	(1.296)	(0.341)
1-hop neighbor AITE	–	–	Yes	-0.736^{**}
2-hop neighbor AITE	–	–	Yes	-0.685^{**}
Panel B: Outcome = log realized hires				
AITE	-5.266^{***}	-4.338^{***}	-0.971	0.029
	(0.608)	(0.571)	(2.391)	(0.552)
Panel C: Outcome = average wage				
AITE	-0.069	-0.067	-0.267	-0.095
	(0.200)	(0.201)	(0.797)	(0.150)
Panel D: Outcome = wage coefficient of variation				
AITE	-0.717^{***}	-0.720^{***}	-2.441^{***}	-0.041
	(0.218)	(0.217)	(0.845)	(0.123)
GAT embedding controls	No	Yes	Yes	No
City and year fixed effects	Yes	Yes	Yes	Yes
Observations	3,186	3,186	3,186	3,186

Note: Standard errors are clustered by city. Column 3 partials out 1-hop and 2-hop neighbor AITE in the high-dimensional first-stage model. The direct effect on realized hires shrinks from -5.266 to -0.971 and becomes statistically insignificant, consistent with network propagation rather than isolated direct substitution.

decline monotonically across wage quantiles; and the direct effect of exposure largely disappears once skill-network propagation is modeled.

The policy implications differ from a standard automation narrative. Training policy should not simply move workers up a skill ladder. It should move workers across skill-space neighborhoods toward bundles that are distant from the AI penetration region and complementary with AI. Labor-market monitoring should distinguish postings from realized hiring. Finally, empirical work on AI and labor should relax SUTVA when skills are densely connected. In such settings, traditional estimates can misattribute network transmission to direct displacement.

REFERENCES

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics*, 4B: 1043–1171.

Acemoglu, Daron, and Pascual Restrepo. 2018. "Artificial Intelligence, Automation and Work." NBER Working Paper 24196.

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33(2): 3–30.

Acemoglu, Daron, and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6): 2188–2244.

Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica*, 90(5): 1973–2016.

Athey, Susan, Julie Tibshirani, and Stefan Wager. 2019. "Generalized Random Forests." *Annals of Statistics*, 47(2): 1148–1178.

Autor, David H., and Anna Salomons. 2018. "Is Automation Labor Share-Displacing?" *Brookings Papers on Economic Activity*, 2018(1): 1–87.

Autor, David H., David Dorn, and Gordon H. Hanson. 2015. "Untangling Trade and Technology: Evidence from Local Labour Markets." *Economic Journal*, 125(584): 621–646.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change." *Quarterly Journal of Economics*, 118(4): 1279–1333.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings*, 108: 43–47.

Chao, Xiaojing, and Wenhui Zhou. 2021. "The Impact of Artificial Intelligence on the Labor Income Share: A Skill-Biased Perspective." *Research on Economics and Management (Jingji yu Guanli Yanjiu)*, 42(2): 82–94.

Chen, Tianqi, and Carlos Guestrin. 2016. "XGBoost: A Scalable Tree Boosting System." 785–794.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. "Double/Debiased Machine Learning for Treatment and Structural Parameters." *Econometrics Journal*, 21(1): C1–C68.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." OpenAI.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence." *Strategic Management Journal*, 42(12): 2195–2217.

Gathmann, Christina, and Uta Schönberg. 2007. "How General Is Human Capital? A Task-Based Approach." IZA Discussion Paper No. 3067.

Goldin, Claudia, and Lawrence F. Katz. 2008. *The Race between Education and Technology.* Harvard University Press.

Grinsztajn, Léo, Edouard Oyallon, and Gaël Varoquaux. 2022. "Why Do Tree-Based Models Still Outperform Deep Learning on Tabular Data?"

Leung, Michael, and Panagiotis Loupos. 2025. "GNN for Causal Inference Under Network Confounding." *AEA Papers & Proceedings*, 115: 1–6.

Manning, Sam J., and Tomás Aguirre. 2026. "How Adaptable Are American Workers to AI-Induced Job Displacement?" NBER 34705.

Manski, Charles F. 2013. "Identification of Treatment Response with Social Interactions." *Econometrics Journal*, 16(1): S1–S23.

Pizzinelli, Carlo, et al. 2023. "Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications." IMF WP/23/216.

Wang, Yongqin, and Wen Dong. 2020. "How Does the Rise of Robots Affect the Chinese Labor Market? Evidence from Listed Manufacturing Firms." *Economic Research Journal (Jingji Yanjiu)*, 55(10): 159–175.

Webb, Michael. 2020. "The Impact of Artificial Intelligence on the Labor Market." Stanford University.

Yao, Jiaquan, Kunpeng Zhang, Lipeng Guo, et al. 2024. "How Does Artificial Intelligence Improve Firm Productivity? A Labor Skill Structure Perspective." *Management World (Guanli Shijie)*, 40(2): 101–116.

Zhang, Dandan, Hang Yu, Lixing Li, et al. 2025. "Measuring AI Technology Exposure and Its Impact on Labor Demand: New Evidence from Large Language Models." *Management World (Guanli Shijie)*, 41(7): 59–75.

Descriptive Statistics and Robustness Diagnostics

Table A1 reports descriptive statistics for the main city-by-year panel. The sample contains 3,232 city-year observations. Mean AITE is 0.265, with a standard deviation of 0.412, reflecting the fact that exposure is zero in 2016–2022 and rises quickly after 2023.

Table A1—Descriptive Statistics, City-by-Year Panel

Variable	*N*	Mean	SD	Min	P25	Median	Max
Log postings	3,232	5.314	2.083	0.000	4.094	5.425	12.062
Log average wage	3,201	9.082	0.491	6.898	8.833	9.058	12.874
Log hires	3,232	3.562	3.383	0.000	0.000	3.761	12.723
Wage coefficient of variation	3,201	0.994	0.885	0.000	0.468	0.620	16.010
AITE	3,232	0.265	0.412	0.000	0.000	0.000	1.264
Bartik instrument	3,232	0.107	0.110	−0.059	0.034	0.085	1.438
Firm size	3,232	23.587	1.146	0.000	23.206	23.615	28.502
Leverage	3,232	0.508	0.090	0.000	0.467	0.511	2.286
ROA	3,232	0.034	0.060	−1.852	0.025	0.037	1.135
Capital expenditure	3,232	0.047	0.020	−0.089	0.036	0.045	0.252
Tobin's Q	3,232	1.844	0.658	0.000	1.541	1.749	18.375

Note: The sample covers 404 Chinese prefecture-level and above cities from 2016 to 2025. Firm variables are city-level averages across listed firms.

A1. Nonzero Pre-Period AITE Sensitivity

The baseline sets AITE to zero in 2016–2022 because generative large models had not yet diffused. Table A2 raises the pre-period AITE value from 0 to 0.30. All estimates remain negative and highly significant.

A2. DML Nuisance Diagnostics

Mechanism and Heterogeneity Figures

B1. Additional Robustness Checks

To evaluate the sensitivity of the baseline Double Machine Learning (DML) estimates, Figure B1 presents a coefficient forest plot across alternative econometric specifications. The specifications include standard OLS, two-way fixed

Table A2—Sensitivity to Nonzero Pre-Period AITE

Pre-period AITE value	DML estimate	Standard error	p-value
0.00, baseline	-1.330^{***}	0.375	0.0004
0.02	-1.351^{***}	0.382	0.0004
0.05	-1.386^{***}	0.392	0.0004
0.08	-1.422^{***}	0.402	0.0004
0.10	-1.444^{***}	0.409	0.0004
0.15	-1.508^{***}	0.428	0.0004
0.20	-1.573^{***}	0.449	0.0005
0.30	-1.709^{***}	0.494	0.0005

Note: The outcome is log postings. Raising the pre-period baseline compresses treatment variation and mechanically raises the absolute coefficient, but the sign and significance are stable. $^{***}p < 0.01$.

Table A3—Out-of-Fold Diagnostics for DML Nuisance Models

Outcome	$ml_l(Y)$ RMSE	$ml_l(Y)$ R^2	$ml_m(D)$ RMSE	$ml_m(D)$ R^2
Postings	0.580	0.036	0.041	−0.004
Hires	1.092	0.030	0.041	−0.004
Wage CV	0.613	0.000	0.040	−0.006

Note: $ml_l(Y)$ is the outcome nuisance model and $ml_m(D)$ is the treatment nuisance model. GroupKFold cross-fitting is grouped by city. The near-zero treatment-model R^2 reflects conservative cross-fitting after two-way demeaning.

effects (FE), and varying combinations of covariates. The point estimates remain consistently negative and statistically significant across all models, confirming that the negative impact of AI exposure on labor demand is not driven by the functional form assumptions of the baseline DML estimator.

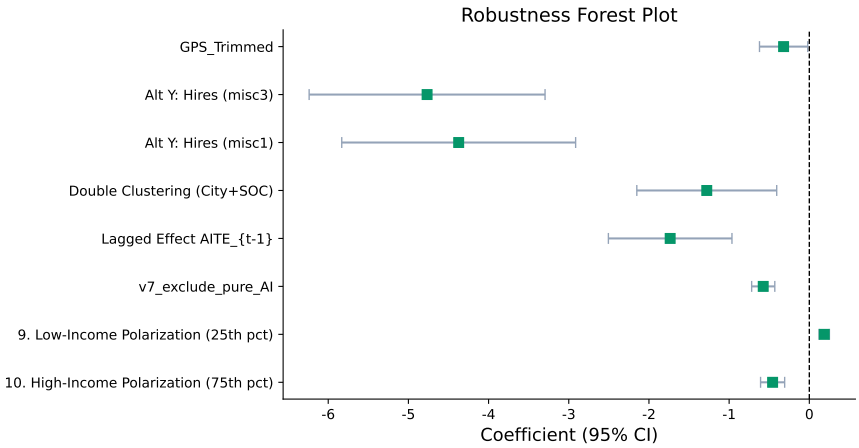


FIGURE B1. ROBUSTNESS FOREST PLOT ACROSS ALTERNATIVE SPECIFICATIONS

To rule out the possibility that the estimated effects are driven by unobserved random noise or spurious correlations, Figure B2 illustrates the distribution of placebo estimates from 5,000 random permutations of the AI exposure treatment variable. The actual DML estimate (indicated by the dashed line) lies at the extreme tail of the placebo distribution ($p < 0.001$), strongly rejecting the null hypothesis of zero effect under randomized treatment assignment.

B2. Extended Heterogeneity and SHAP Analysis

Figure B3 shows the QuantileDML gradient across five wage quantiles. The gradient moves from positive at the bottom to negative at the top, with a crossing near the median.

The subsequent figures provide an interpretative framework for the heterogeneous treatment effects (HTE) generated by the Causal Forest model. Figures B4 and B5 plot the one-way Partial Dependence and SHAP value curves for key regional covariates, such as Tertiary Industry GDP share and Population Density. These curves reveal significant nonlinearities: the marginal impact of AI exposure on labor demand shifts abruptly beyond certain thresholds of regional economic development.

Figures B6 and B7 extend this framework to two-dimensional feature interactions, demonstrating how firm-level financial characteristics moderate the regional

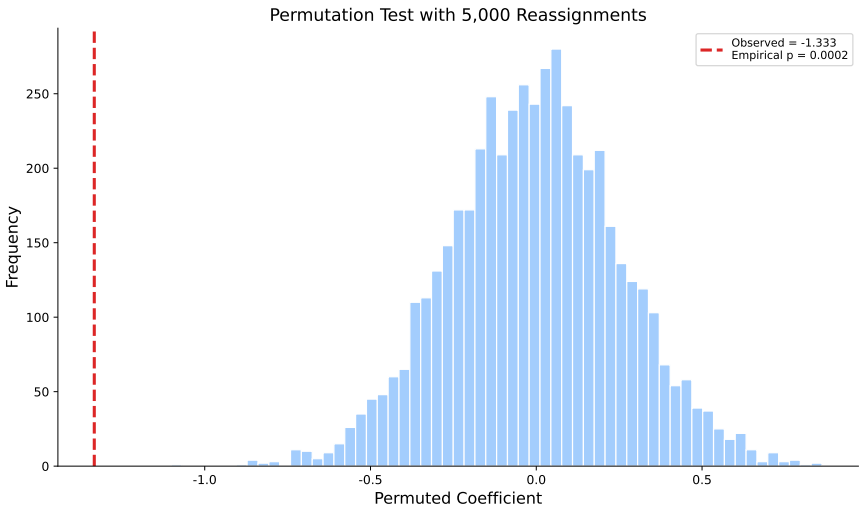


FIGURE B2. PLACEBO TEST: DISTRIBUTION OF 5,000 RANDOM PERMUTATIONS

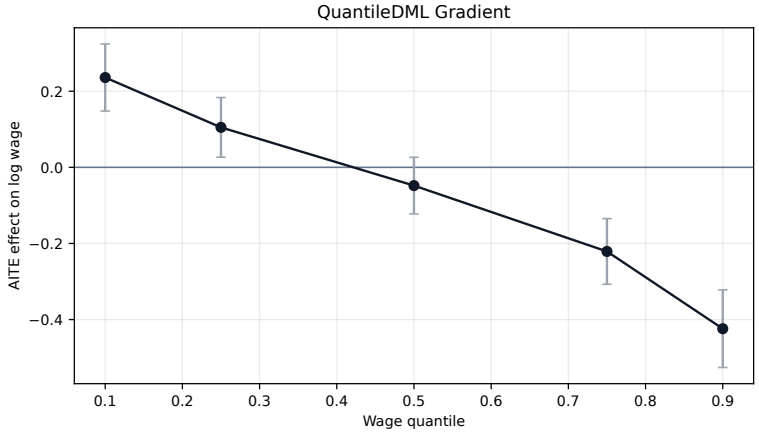


FIGURE B3. QUANTILEDML GRADIENT WITH 95 PERCENT CONFIDENCE INTERVALS

baseline effects. Finally, Figure B8 visualizes a shallow decision tree trained directly on the estimated CATEs. This policy tree discretizes the continuous heterogeneity space into interpretable terminal nodes, delineating the specific subpopulations most vulnerable to AI-induced labor substitution.

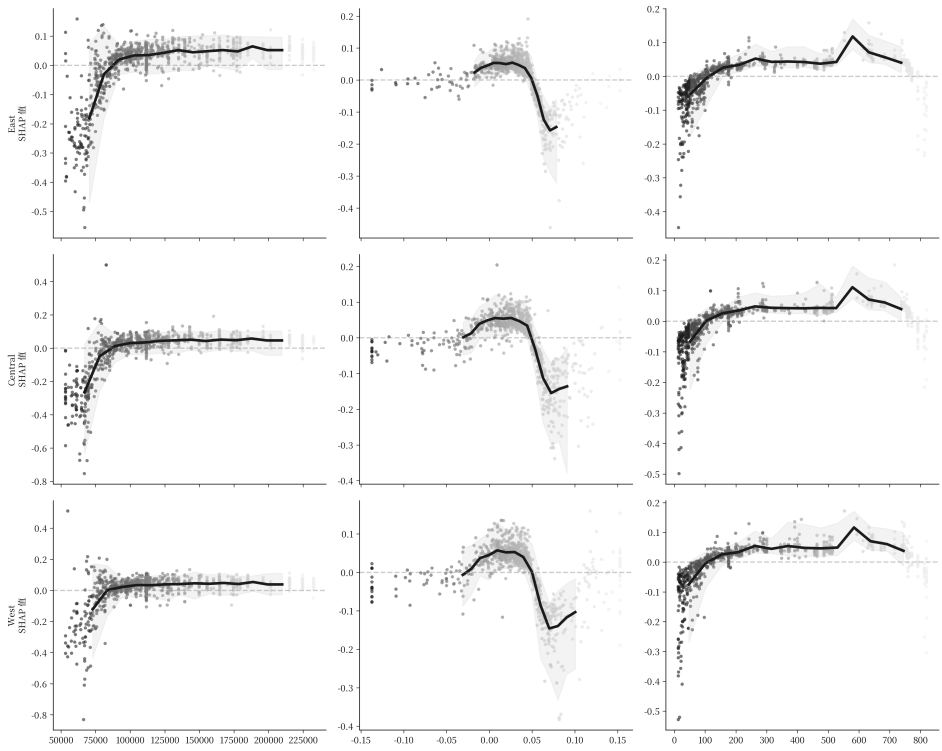


Figure B4. Regional One-Way SHAP Dependence

GNN Spillover Appendix

In the context of the spatial skill network, it is critical to validate the structural parameters of the Graph Autoencoder (GAE) and the mechanism of network spillovers. Figure C2 presents the training loss convergence and latent space clustering diagnostics for the GAE embedding process, verifying the stability of the topological representation.

Figure C4 evaluates the robustness of the network construction by testing alternative edge-weight threshold values for the adjacency matrix. The fundamental network topology and subsequent spillover estimates remain invariant across reasonable threshold specifications. Finally, Figure **??** quantifies the spatial decay of AI-induced spillovers across the skill network. The magnitude of the spillover ef-

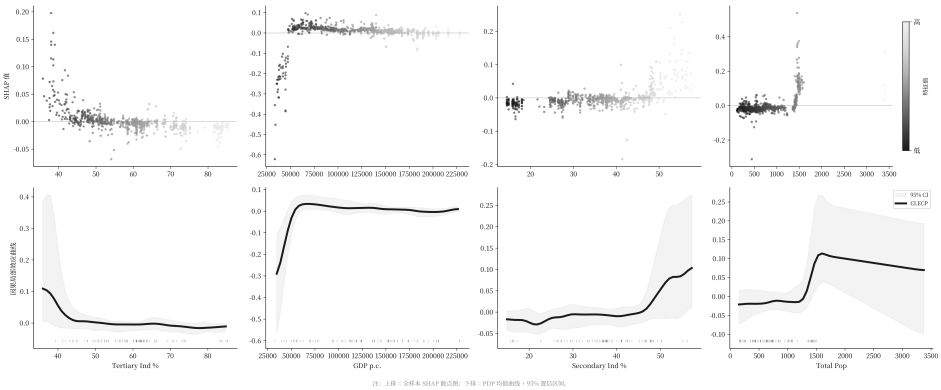


FIGURE B5. ONE-WAY SHAP DEPENDENCE AND PDP CURVES

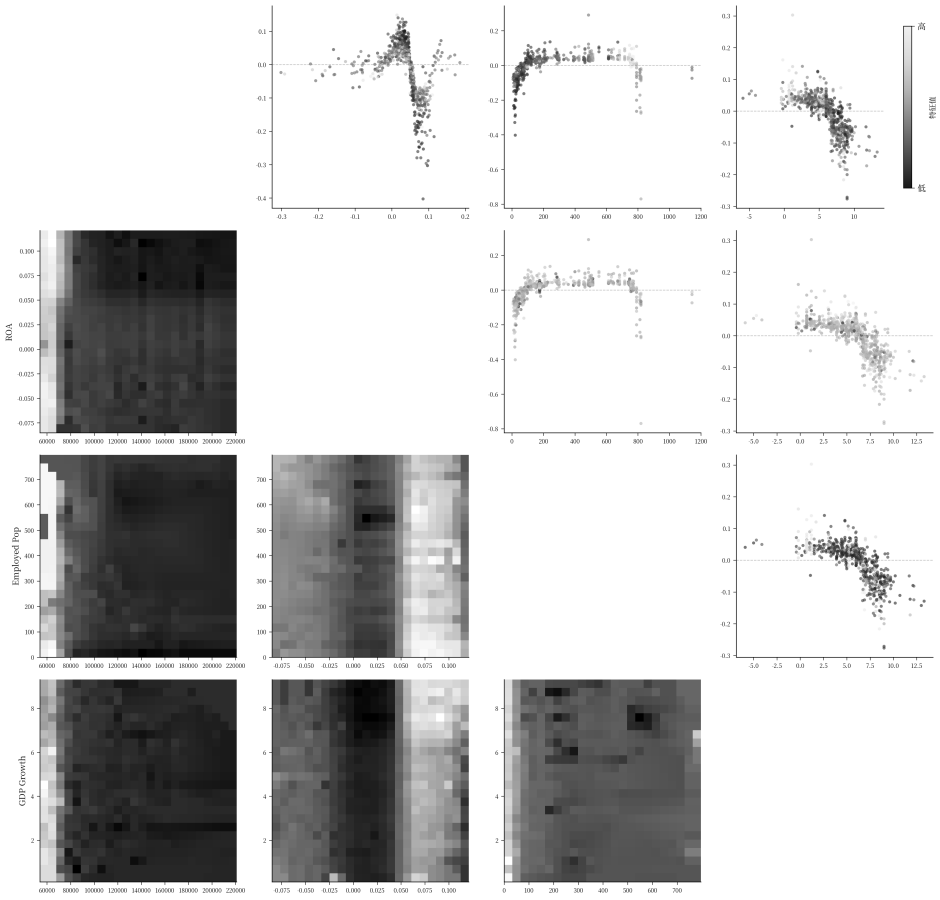


FIGURE B6. BIVARIATE SHAP INTERACTION HEATMAP

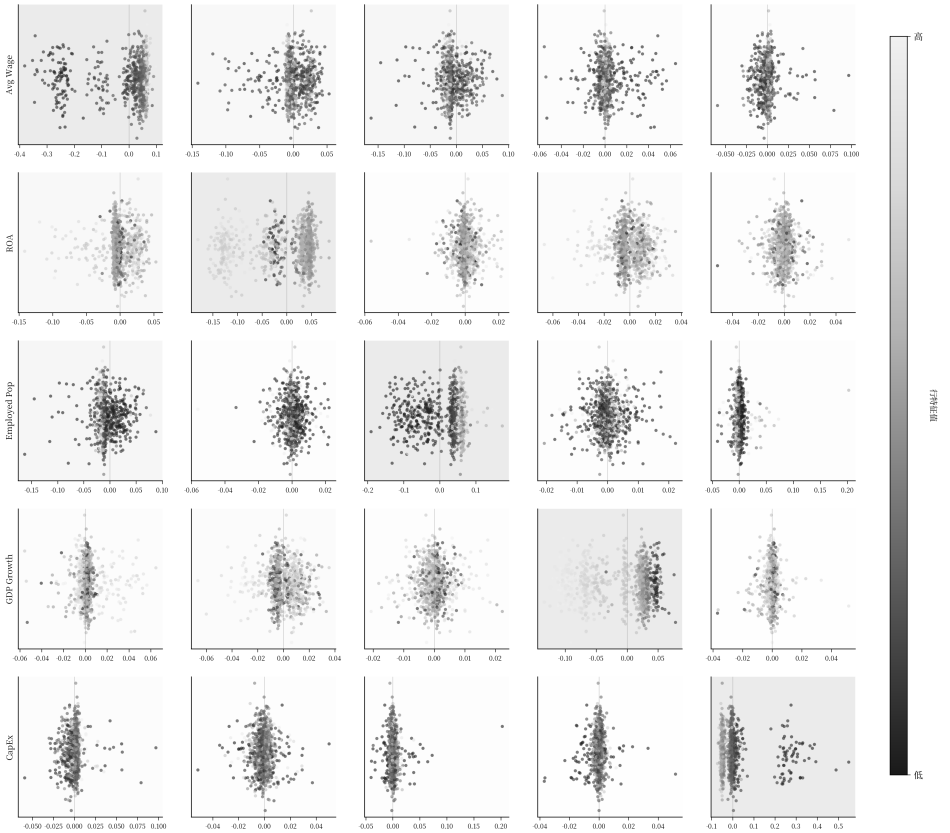


Figure B7. SHAP Interaction Matrix, Five-by-Five Grid

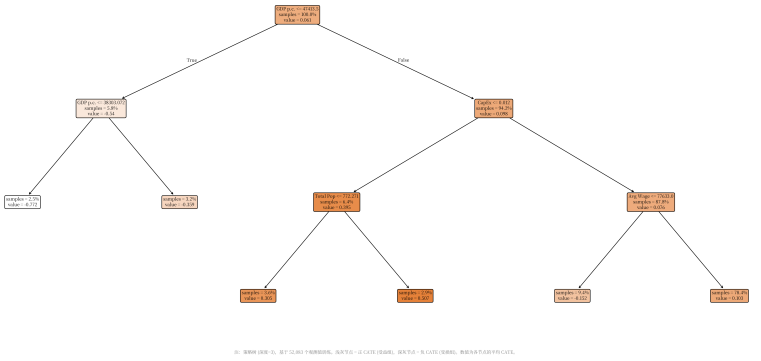


Figure B8. Policy Tree Based on Estimated CATEs

fect strictly and monotonically decreases as the topological hop distance increases, confirming the localized nature of cross-occupation substitution dynamics.

Figure C1 displays the UMAP projection of the DWA skill embedding space. Nodes close to one another in the two-dimensional projection are semantically similar work activities.

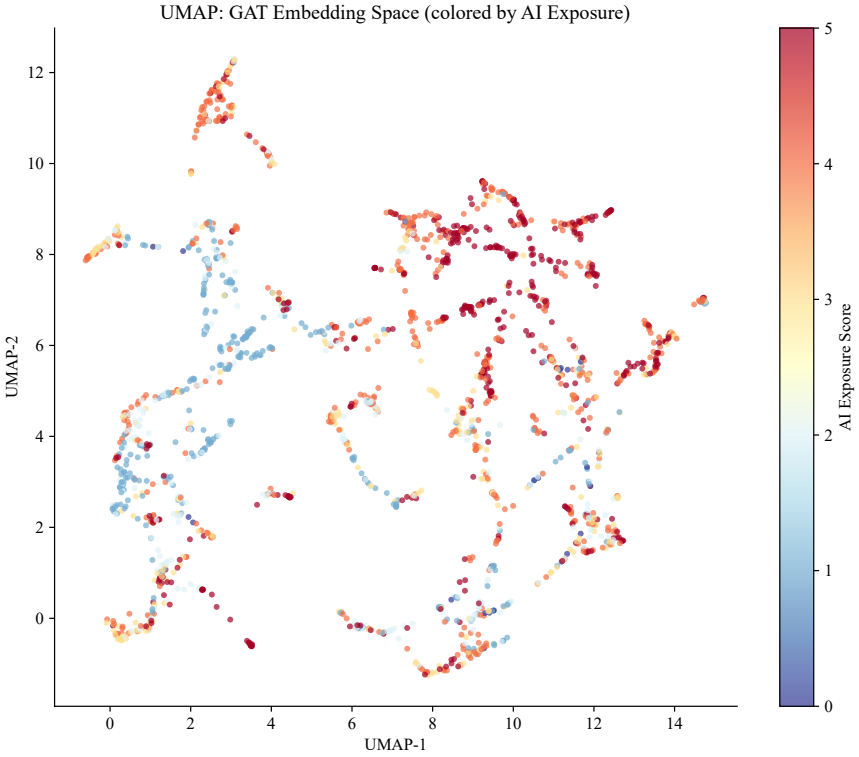


Figure C1. UMAP Visualization of the Skill Embedding Space

The unsupervised graph autoencoder isolates latent high-dimensional structure in the task network. The link-prediction accuracy shows that natural-language task descriptions contain information about cross-job connectedness. Once two work activities are close enough in semantic space, a shock to one skill node can propagate through first- and higher-order links.

The decay gradient identifies the range of transmission in the skill network. One-hop neighbors bear the strongest displacement pressure, consistent with displaced workers first seeking jobs that preserve task-specific human capital. As graph distance rises, transfer frictions increase and spillovers decay.

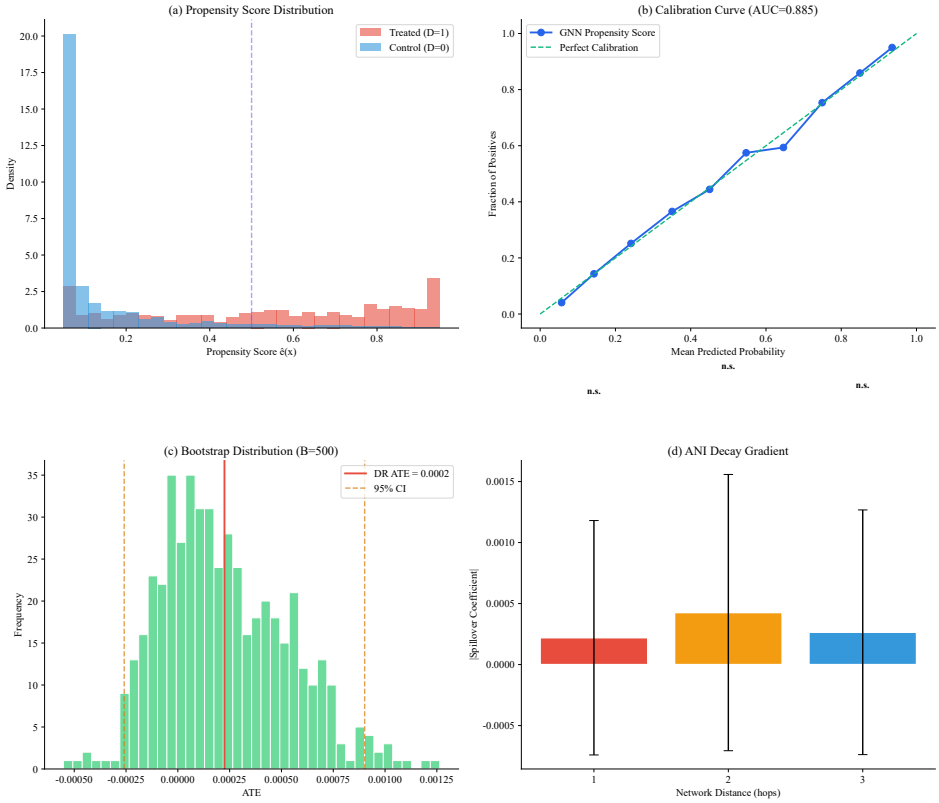


FIGURE C2. GRAPH AUTOENCODER DIAGNOSTICS: AUC, AP, AND EMBEDDING SEPARATION

TABLE C1—THRESHOLD SENSITIVITY DETAILS

Similarity threshold τ	Spillover coefficient	SE	p-value	GAE AUC	Edges
0.6	-0.894^{***}	0.116	$< 10^{-14}$	0.950	84,482
0.7, baseline	-0.894^{***}	0.116	$< 10^{-14}$	0.979	20,504
0.8	-0.894^{***}	0.116	$< 10^{-14}$	0.987	5,204

Note: The edge count falls as the threshold tightens, but the spillover coefficient remains stable. $^{***}p < 0.01$.

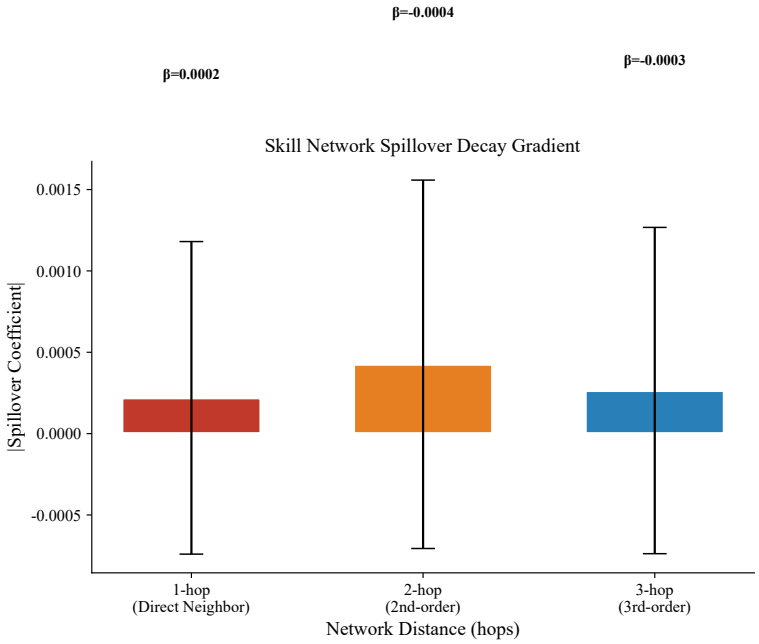


FIGURE C3. SPATIAL DECAY OF SPILLOVER EFFECTS OVER NETWORK DISTANCE

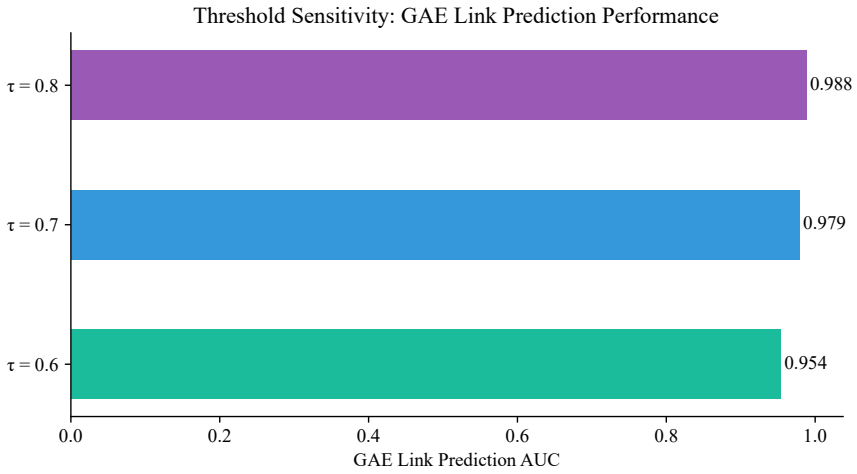


FIGURE C4. THRESHOLD SENSITIVITY FOR SKILL-NETWORK CONSTRUCTION

Alternative AITE Measures and Bias Diagnostics

This appendix evaluates whether the main results depend on the AITE construction. The static-capability variant removes the time evolution of model capability by fixing Elo scores at a base value. The equal-share variant assumes equal use weights for all included large language models. Removing capability dynamics introduces scale bias; removing market concentration dilutes the head-model diffusion effect.

Table D1—DML Estimates under Alternative AITE Measures

	Baseline AITE	Static AITE	Equal-share AITE
Outcome = log hires			
AITE	-4.767^{***}	-8.770^{***}	-1.995^{***}
	(0.751)	(0.771)	(0.220)
Bias interpretation	–	Scale bias	Dilution bias

Note: Cluster-robust standard errors are in parentheses. All variants preserve a negative and significant effect. $^{***}p < 0.01$.

Lee Bounds and Heckman Selection Diagnostics

E1. First-Stage Heckman Probit

Table E1—Heckman First-Stage Probit

Variable	Coefficient	p-value
Constant	-0.672^{***}	$< 10^{-300}$
AITE	-2.073^{***}	$< 10^{-300}$
Years of education	-0.021^{***}	$< 10^{-300}$
Years of experience	-0.0004^{***}	$< 10^{-10}$

Note: The dependent variable equals one if the wage is negotiable and zero if the wage is explicitly posted. The sample contains 4,849,040 job-level observations. A negative AITE coefficient means high-exposure jobs are less likely to mark wages as negotiable and more likely to reveal explicit wages.

E2. Lee Bounds

The high-AITE group has wage-reporting probability $P(S = 1 \mid D = 1) = 0.9967$; the low-AITE group has $P(S = 1 \mid D = 0) = 0.8356$. Under monotonicity, about 16.17 percent of the high-exposure observations are marginal reporters

induced into the observed wage sample. Trimming the upper and lower 16.17 percent of the high-exposure wage distribution yields a Lee lower bound of −4,139.76 RMB and an upper bound of −281.06 RMB. Both bounds are strictly negative, so the conclusion that AI exposure lowers wages does not rely on Heckman's parametric normality assumption.

Estimator Comparisons and Weak-IV Diagnostics

Table F1 compares six estimators: OLS, DML-PLR with two-way fixed effects, DML-PLR without fixed effects, DML-PLIV with two-way fixed effects and the Bartik IV, pooled 2SLS without fixed effects, and LIML without fixed effects. OLS and DML-PLR with fixed effects agree in sign and significance. DML-PLIV standard errors inflate sharply after two-way demeaning because the city panel strips out most cross-sectional share variation. Pooled 2SLS and LIML without fixed effects produce positive signs in some outcomes, reflecting omitted fixed-effect bias; they are not interpreted causally.

Table F1—Comparison across Six Estimators

	OLS	DML FE	DML No FE	DML-IV FE	2SLS No FE	LIML No FE
Panel A: Outcome = log postings						
AITE	−1.529***	−1.330***	−0.379***	2.124	1.357**	1.357**
	(0.073)	(0.375)	(0.114)	(18.975)	(0.602)	(0.602)
Panel B: Outcome = log hires						
AITE	−5.559***	−4.767***	−4.096***	−8.356	5.710***	5.710***
	(0.101)	(0.751)	(0.134)	(30.307)	(1.088)	(1.088)
Panel C: Outcome = wage CV						
AITE	−0.828***	−0.615***	−0.519***	−18.238	1.747***	1.747***
	(0.027)	(0.230)	(0.039)	(46.988)	(0.324)	(0.324)
Fixed effects	Two-way	Two-way	No	Two-way	No	No
Instrument	No	No	No	Bartik	Bartik	Bartik
Cross-fitting	No	Yes	Yes	Yes	No	No
Observations	3,232	3,232	3,232	3,232	3,232	3,201

Note: Standard errors are in parentheses. The positive signs in columns 5–6 reflect omitted two-way fixed effects and are not interpreted causally. The main 2SLS estimate in the paper uses the SOC-level panel with fixed effects and first-stage $F = 310.71$.

Additional Robustness: Instrument Exogeneity

Shock-Level Balance Test

The shock-level balance test examines whether the global GPU compute-supply shock is independent of 2016 baseline city characteristics. Regressing the orthogonalized shock residual on baseline city and firm characteristics yields a joint F statistic of 0.030, a joint p-value of 1.0000, and $R^2 = 0.0001$. This supports the exogenous-shock component of the Bartik design.

Table F2—Weak-Instrument Diagnostics

Outcome	Pooled 2SLS first-stage F	DML-PLR(FE) R_m^2	DML-PLR(No FE) equivalent F
Postings	76.09	−0.004	795.51
Hires	76.09	−0.004	795.51
Wage CV	84.41	−0.006	812.63

Note: The SOC-level panel with fixed effects has overall first-stage $F = 310.71$; the city panel first-stage F is 86.48 with city-clustered standard errors.

Table G1—Balance Test for Baseline City and Firm Characteristics

Baseline characteristic	Coefficient	p-value
Secondary-industry GDP share	−0.000327	0.9228
Tertiary-industry GDP share	−0.000356	0.9114
GDP per capita	−0.000000	0.9577
Population	−0.000003	0.9181
Mean firm size	+0.005363	0.6505
Mean ROA	−0.000396	0.9977
Mean leverage	−0.011205	0.9160
Mean cash flow	+0.012049	0.9600
Mean Tobin's Q	+0.000402	0.9155
Mean growth	−0.000086	0.9849

Full-Permutation Bartik Placebo

As a placebo, the baseline compute-intensive share is replaced with the share of non-compute-sensitive jobs. In the first stage $AITE \sim placebo_iv+X$, the placebo coefficient is -0.004851 with $p = 0.747490$, and the first-stage F statistic collapses to 0.10. Once the link between global compute supply and local compute-sensitive labor is intentionally broken, the placebo share provides no explanatory power for AITE.

Mathematical Details

Proof of Proposition 1. If $\Delta\mathcal{P}_{t+1}$ has positive measure and $l^c(\mathbf{x}) > 0$ on that set, then the change in labor demand is

$$\Delta L^C = -\int_{\Delta\mathcal{P}_{t+1}} l^c(\mathbf{x}) dF(\mathbf{x}) < 0.$$

With complementarity between cognitive and on-site tasks, a fall in cognitive output lowers the marginal product of complementary tasks, so aggregate labor demand falls unless offset by new task creation.

Proof of Proposition 2. Observed wages decompose as

$$E[w_i \mid D_i, S_i = 1] = E[w_i \mid D_i] + \{E[w_i \mid D_i, S_i = 1] - E[w_i \mid D_i]\}.$$

The bracketed term is the selection component. If explicit wage reporting among exposed jobs is positively selected, this term is positive, so the observed effect is upward biased.

Proof of Proposition 3. Let $\beta(\tau) = \int b(\mathbf{x}) dF_\tau(\mathbf{x})$. If F_τ places increasing mass on skill bundles near $\mathcal{P}_t$ as τ rises, and if $b(\mathbf{x})$ is more negative near $\mathcal{P}_t$, then $\beta(\tau)$ declines with τ. Low quantiles can be positive if reallocation and selection dominate demand compression.

Proof of Proposition 4. The true model is $Y_i = \theta D_i + \delta N_i + \epsilon_i$. Omitting N_i yields

$$\operatorname{plim} \hat{\theta} = \theta + \delta \frac{\operatorname{Cov}(D_i, N_i)}{\operatorname{Var}(D_i)}.$$

In local skill neighborhoods, $\operatorname{Cov}(D_i, N_i) > 0$. If $\delta < 0$, the omitted-variable term is negative, so the direct effect is overestimated in magnitude.

REMARK 1 (One-dimensional projection): *For intuition only, one may project skill bundles onto an AI direction $\mathbf{v}_{AI}$ and define $s_i = \mathbf{v}'_{AI}\mathbf{x}_i$. A scalar threshold is then just a projection of the multidimensional region $\mathcal{P}_t$. This projection cannot capture nonuniform expansion across skill clusters or propagation through skill-space neighborhoods, so it is not the main model.*

Additional Figures

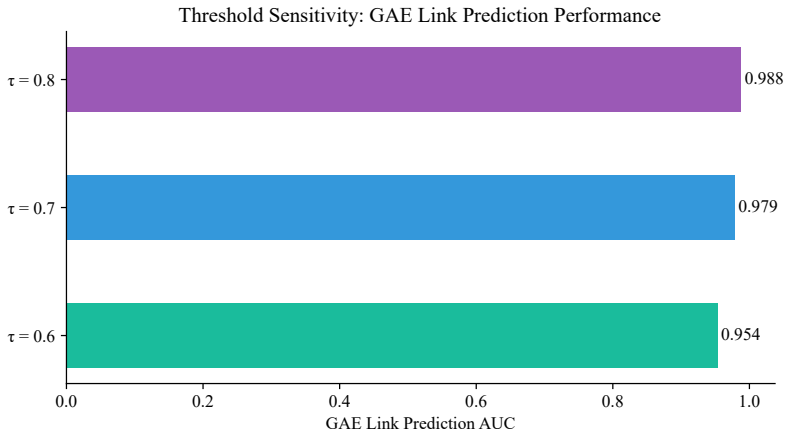


FIGURE I1. NETWORK THRESHOLD SENSITIVITY

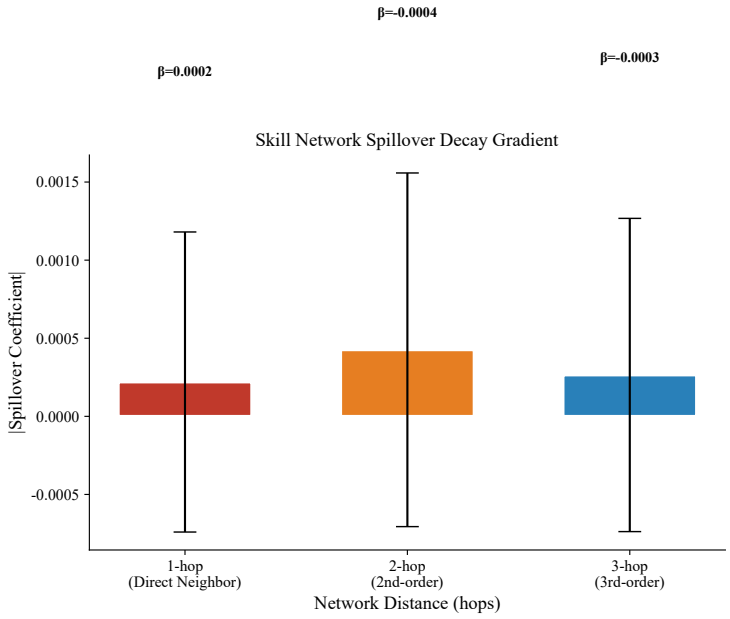


FIGURE I2. DECAY OF SPILLOVER EFFECTS OVER SKILL-NETWORK DISTANCE

Note: The decay pattern supports the approximate-neighborhood-interference assumption: 1-hop effects are strongest, 2-hop effects decline, and 3-hop effects approach zero.