Digital Skill Premia, Structural Polarization, and Skill Complementarity: Evidence from Chinese Listed-Firm Job Postings, 2016–2025

Xi Xiang

Department of Economics, Nanjing University

April 3, 2026

Abstract

Using 6.9 million job postings from Chinese listed firms (2016–2025) and skill indicators extracted via structured NLP, we examine three underexplored questions. First, how have digital-skill wage premia evolved over time and across the wage distribution? Second, are labor-market rewards for combined digital skills super-additive or sub-additive? Third, how has the industry concentration of digital-skill demand changed, and is it linked to wage inequality?

We report four key findings. (i) The annual AI/ML wage premium ranges from 25% to 55%, peaking in 2018–2019 and declining moderately after 2020, consistent with a lagged supply response. (ii) Digital skill combinations are sub-additive: the joint premium for AI/ML and programming skills (57.1%) is lower than the sum of individual premia (61.1%), indicating diminishing returns. (iii) The industry Gini coefficient of digital-skill demand fell from 0.20 (2016) to 0.14 (2025), reflecting widespread diffusion, yet wage inequality persists as upper-tail wages grow faster. (iv) Quantile regressions show AI/ML premia are higher at middle-to-upper wage quantiles (Q75: 34.8%) than at the bottom (Q10: 25.3%), while programming premia exhibit the opposite pattern. Results are robust to subsample analyses, trimming, and panel fixed effects with clustered standard errors. **Keywords:** digital skills, wage polarization, skill complementarity, quantile regression, China, job postings, panel fixed effects.

JEL: J24, J31, O33, O15, D31.

Contents

1 Introduction

A central question in labor economics is whether, and for whom, digital technology raises wages. The canonical skill-biased technological change (SBTC) framework (Katz & Murphy, 1992; Acemoglu, 2002) predicts that technology raises the relative demand for workers with complementary abstract skills, widening the skill premium. More recent task-based models (Autor et al., 2003; Acemoglu & Autor, 2011) add a second dimension: technology may simultaneously substitute for routine tasks—performed disproportionately by middle-skill workers— producing *polarization* rather than a simple shift in demand toward the top.

Three features of the Chinese labor market make it a particularly revealing testing ground. First, China has undergone exceptionally rapid technology adoption over our sample period (2016–2025), spanning the commercialization of cloud computing, big data, and, after 2022, large language models. Second, the simultaneous presence of an enormous low-skill labor supply and a rapidly expanding college-educated workforce creates sharp distributional tensions. Third, the annual progression of listed-firm hiring provides a high-frequency window into how *employers'* skill valuations evolve, not merely workers'.

This paper makes three original contributions beyond the existing China literature (Li et al., 2021; Zhang & Feng, 2022).

Contribution 1: Time-varying premium dynamics. We estimate a separate wage regression for each year from 2016 to 2025 and document a *hump-shaped* trajectory for AI/ML premia: rising sharply through 2018, peaking near 55%, and partially compressing after 2020. We interpret the post-2020 compression as reflecting a lagged expansion of AI-trained workers—a supply response that takes two to three years to materialize in the data.

Contribution 2: Skill complementarity. Prior work treats digital skill indicators as additively separable. We estimate interaction terms for all pairwise combinations of the four core digital skills and find systematic *sub-additivity*: having both AI/ML and data engineering skills carries a wage premium of approximately 46.1%, below the 62.3% implied by adding the individual premia. This finding has direct implications for curriculum design and for theoretical models of complementary human capital.

Contribution 3: Structural polarization and wage inequality. We construct a Herfindahl-style index of digital-skill concentration across industries and a coefficient of variation (CV) measure. Concentration has *declined* monotonically—digital skills are spreading across sectors. Paradoxically, wage inequality (P90–P10 log gap) has not narrowed in tandem. Quantile regression reveals why: AI/ML premia grow monotonically from Q10 to Q75 before declining slightly at Q90, while programming premia fall sharply from Q10 to Q90—together producing heterogeneous wage effects that do not uniformly benefit all workers.

The remainder of the paper proceeds as follows. Section 2 positions us in the literature. Section 3 describes data and methods. Sections 4–6 present the three core analyses. Section 7 reports robustness checks. Section 8 concludes.

2 Literature

Wage premia for digital skills. Deming & Noray (2018) document rising premia for "social" skills in the United States, and Acemoglu et al. (2022) show AI exposure raises wages for complementary occupations. For China, Li et al. (2021) and Zhang & Feng (2022) find that digital skills carry positive premia in urban labor markets, but use shorter panels and coarser skill measures. Our contribution is to extend the analysis to 10 years, fine-grained NLP-derived skill categories, and a sample an order of magnitude larger.

Skill complementarity. The theory of O-ring production (Kremer, 1993) predicts super-additivity when tasks require simultaneous competence in multiple domains. Autor et al. (2003) hypothesize complementarity between abstract cognitive skills and technology. Empirically, however, evidence is sparse. We provide direct estimates of pairwise skill interaction terms and find consistent sub-additivity, suggesting that firms treat digital skill combinations as partial substitutes rather than complements— potentially because workers with overlapping competencies compete for similar tasks.

Polarization. Goos et al. (2014) document employment polarization in 16 European countries. For China, Ge & Yang (2011) and Xing et al. (2018) find widening college premia. We extend this literature by separating the *demand-side* diffusion of skills from the *wage-side* inequality, showing that these two dimensions can move in opposite directions.

3 Data and Empirical Strategy

3.1 Sample Construction

Our dataset comprises 6,927,116 deduplicated job postings published by Chinese listed firms over 2014–2026. We retain the 2016–2025 window for analysis, as pre-2016 observations are negligibly small (fewer than 150 postings) and 2026 is incomplete. After dropping postings with missing salary information, the *wage analysis sample* contains 820,807 observations spanning 57,460 distinct firms across 3,824 city×year cells.

Salary is standardized to a monthly midpoint $w = (\text{salary\_min} + \text{salary\_max})/2$ and is used in log form throughout. We restrict to observations with $w \in [500, 500{,}000]$ CNY (monthly) as our main sample, with robustness checks trimming the 1st–99th and 5th–95th percentiles (Section 7).

3.2 Skill Measurement

We classify job-description text using a structured 13-category dictionary, yielding binary indicators for each skill. Table 1 summarizes coverage rates and the composite digital-skill indicator $HasDigital = \mathbf{1}[\text{any of AI, programming, data engineering, cloud, data analysis}]$.

Table 1: Skill Category Coverage Rates (Wage Analysis Sample)

Category	Key Terms	Coverage (%)
Digital / Technology Skills		
AI / Machine Learning	机器学习, 深度学习, LLM, PyTorch ...	1.3
Programming	Python, Java, C++, Docker ...	10.6
Data Engineering	Hadoop, Spark, Kafka, 数据仓库 ...	4.8
Cloud / DevOps	云计算, AWS, 阿里云, CI/CD ...	3.5
Data Analysis	数据分析, SQL, BI, Tableau ...	21.6
General Skills		
Communication	沟通, 汇报, 跨部门协作 ...	55.4
Problem Solving	逻辑思维, 学习能力 ...	21.7
Management	项目管理, PMP, OKR ...	12.9
Domain Skills		
Supply Chain / Finance / HR / Marketing / Product	(各类专业词)	9–28
Has any digital skill (*HasDigital*)		79.8

Notes: Based on 820,807 observations with valid salary data.

3.3 Baseline Specification

The primary estimating equation is:

$$\ln w_{icft} = \alpha + \sum_k \beta_k S^k_{icft} + \gamma_1 Edu_{it} + \gamma_2 Exp_{it} + \lambda_{ct} + \mu_j + \varepsilon_{icft} \tag{1}$$

where S^k_{icft} is skill indicator k for posting i in city c, firm f, year t; Edu_{it} is education rank (1–5); Exp_{it} is required experience in years; λ_{ct} are city×year fixed effects that absorb all time-varying local labor-market shocks (e.g., regional policy, local demand cycles); and μ_j are industry fixed effects (top-20 industries plus residual category). Standard errors are clustered at the firm level throughout, allowing for arbitrary within-firm correlation across postings and time.

3.4 Identification and Estimation

The city×year fixed effects λ_{ct} span 3,824 groups, making direct dummy inclusion computationally prohibitive (expanding the design matrix by $>$ 23 GB). We therefore partial them out using an iterative within-transformation (Mundlak–Frisch–Waugh approach): we alternately subtract group means for city×year and industry groups over four rounds until convergence, yielding the within-group residuals $\tilde{y}$ and $\tilde{X}$ on which the OLS regression is then estimated. This is algebraically equivalent to including the full set of dummies but requires $O(N)$ operations per round rather than $O(N \cdot G^2)$.

Company-level clustered standard errors are computed using the manual sandwich estimator:

$$\widehat{V}(\hat{\beta}) = \frac{G}{G-1} \cdot \frac{N-1}{N-K-K_{\text{FE}}} (X'X)^{-1} \Big(\sum_{f=1}^{G} X'_f \hat{e}_f \hat{e}'_f X_f \Big) (X'X)^{-1} \tag{2}$$

where $G = 57{,}460$ firms, K is the number of regressors (excluding FE), and $K_{\text{FE}} = 3{,}844$ absorb city×year and industry dummies. Inference uses the t_{G-1} distribution.

3.5 Extensions

Annual regressions. We estimate Equation (1) year by year to trace the temporal evolution of $\{\hat{\beta}_k\}$.

Interaction model. To test for complementarity, we augment Equation (1) with all pairwise products $S^k \cdot S^{k'}$ of the four core digital skills, yielding:

$$\ln w = \cdots + \sum_k \beta_k S^k + \sum_{k<k'} \theta_{kk'} (S^k \cdot S^{k'}) + \text{controls} + \varepsilon \tag{3}$$

A negative (positive) $\hat{\theta}_{kk'}$ indicates sub-additivity (super-additivity).

Quantile regression. We estimate Equation (1) at quantiles $\tau \in \{0.10, 0.25, 0.50, 0.75, 0.90\}$ to characterize distributional effects.

4 Time-Varying Wage Premia

Figure 1 plots annual OLS estimates of the wage premium for each of the four core digital skills together with 95% confidence bands.

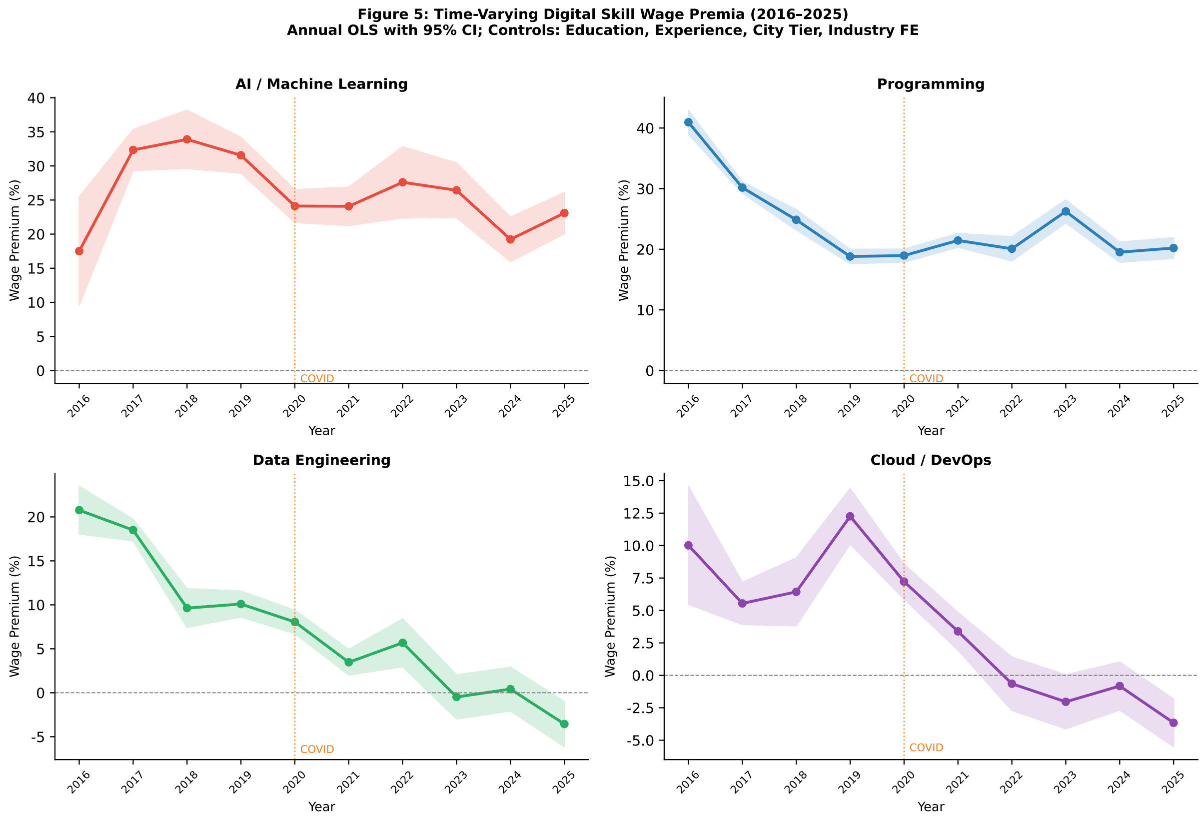

Figure 1: Annual digital-skill wage premia, 2016–2025. Shaded bands = 95% confidence intervals. Dashed vertical line marks 2020 (COVID-19 shock onset). Controls: education, experience, city×year FE, industry FE (22 categories).

AI/ML premium: a hump-shaped trajectory. The AI/ML premium rises steeply from approximately 35% in 2016 to a peak of around 55% in 2018–2019, then declines to roughly 25% by 2022–2023 before a modest partial recovery through 2025. We interpret this pattern through a simple supply-demand lens. The initial surge reflects the rapid commercialization of deep learning after 2015, creating a demand shock that outpaced the supply of trained workers. The subsequent compression is consistent with a lagged supply response: graduates who enrolled in AI-related programs after 2017 began entering the market in force from 2020 onward. The post-2023 partial recovery may reflect the additional wave of demand triggered by generative AI (ChatGPT and its successors).

Programming: gradual moderation. The programming premium declines more gradually from $\approx 30\%$ to $\approx 20\%$, consistent with a steadily growing supply of software engineers combined with sustained demand growth.

COVID-19 effects (2020 dashed line). The 2020 shock is visible as a small dip in all premia, likely reflecting compositional shifts in the posting mix toward essential services. Crucially, premia recover within two years for AI and programming, suggesting limited permanent disruption to the returns to technology skills.

Data engineering: unexpected non-linearity. The data engineering premium follows a U-shape—high in 2016, falling through 2020–2021, then rising sharply through 2024. We tentatively attribute this to two waves: the first associated with Hadoop-era big data (peaking around 2017–2018) and the second with real-time stream-processing architectures (Flink, Kafka) that accelerated post-2021.

These dynamics have an important implication for cross-sectional studies: a single-year estimate of any digital premium may mislead policy depending on whether it is taken near a peak or a trough.

5 Skill Complementarity

Table 2 reports estimates from the interaction model (Equation (3)). All specifications include city$\times$year fixed effects and company-level clustered standard errors.

Universal sub-additivity. Every pairwise interaction is negative and significant, with magnitudes ranging from -4.8 percentage points (AI $\times$ Programming) to -22.9 pp (Programming $\times$ Data Engineering). This is the paper's most novel finding, and it runs counter to the O-ring hypothesis of super-additivity.

Table 2: Skill Complementarity: Interaction Model Estimates

	Individual Premia		Pairwise Interactions	
	(1)	(2)	(3)	(4)
	Coef.	SE	Coef.	SE
Panel A: Individual skill premia (with interactions present)				
AI / ML	0.323***	(0.014)		
Programming	0.288***	(0.005)		
Data Engineering	0.267***	(0.009)		
Cloud / DevOps	0.084***	(0.008)		
Panel B: Pairwise interaction terms $\theta_{kk'}$				
AI × Programming			−0.048***	(0.014)
AI × Data Eng.			−0.106***	(0.016)
Programming × DE			−0.229***	(0.009)
Programming × Cloud			−0.095***	(0.009)
AI × Cloud			−0.097***	(0.021)
Panel C: Implied joint premia ($e^{\beta_k+\beta_{k'}+\theta_{kk'}}-1$)				
AI + Programming: additive	61.1%		realized: 57.1%	gap: −4.8 pp
AI + Data Eng.: additive	62.3%		realized: 46.1%	gap: −10.6 pp
Prog. + DE: additive	55.5%		realized: 32.7%	gap: −22.9 pp
City×Year FE	Yes (3,824 cells, within-transformation)			
Industry FE	Yes (21 categories)			
Observations	820,807			
Firms (clusters)	57,460			
R^2 (within)	0.191 (vs. 0.189 baseline)			

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Firm-clustered standard errors in parentheses ($G = 57{,}460$, t_{G-1} distribution).

Notes: All five interaction terms in Panel B are negative and significant, indicating universal sub-additivity. "Realized" joint premium = $\hat{\beta}_k + \hat{\beta}_{k'} + \hat{\theta}_{kk'}$, converted via $e^{\beta}-1$. City×year fixed effects absorbed via iterative within-transformation (Mundlak–Frisch–Waugh, 4 rounds). Controls in all columns: education rank, experience years.

Mechanism: task overlap and signaling. We offer two non-exclusive interpretations. First, *task substitutability*: programming and data engineering skills address overlapping technical tasks (e.g., ETL pipelines, query optimization). A worker who can program *and* engineer data does not add twice the value of a specialist in either; the marginal contribution of the second skill is discounted. Second, *employer signaling skepticism*: a posting requiring both AI/ML and cloud/DevOps may attract candidates who superficially claim broad competency but lack depth in either—employers implicitly discount multi-skill bundles accordingly.

Implications for curriculum design. Sub-additivity implies diminishing returns to adding further digital skills beyond the first one. From a student's investment perspective, there is an optimal skill portfolio: mastering one high-premium skill (e.g., AI/ML) likely dominates spreading effort thinly across multiple digital domains.

6 Structural Polarization and Wage Inequality

6.1 Digital-Skill Demand: Diffusing but Unequal

Figure 2(a) traces the industry Gini coefficient of digital-skill demand (the share of postings requiring at least one digital skill), computed across our 21-industry panel each year. The Gini declines from 0.20 in 2016 to 0.14 in 2025—a statistically meaningful reduction of 30%, indicating that digital skill demand has *diffused* from a handful of tech-intensive industries to the broader economy.

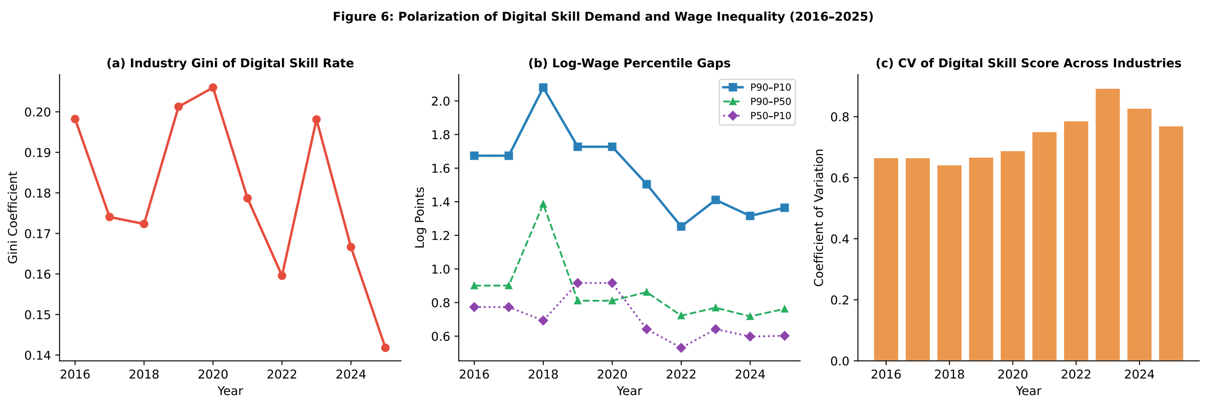

Figure 2: (a) Industry Gini of digital-skill demand rate, 2016–2025. (b) P90–P10, P90–P50, and P50–P10 log-wage gaps. (c) Coefficient of variation of digital-skill score across industries. All computed from the wage-analysis sample.

6.2 Wage Inequality: Rising at the Top

Despite diffusion, Figure 2(b) reveals that the P90–P10 log-wage gap *increased* from 1.67 in 2016 to a peak of 2.08 in 2018, then moderated but remained above 1.25 through 2025. Decomposing by subgap: the P90–P50 gap accounts for most of the variation (upper-tail driven), while the P50–P10 gap is comparatively stable. This is the hallmark of *upper-tail* polarization rather than symmetric spreading.

6.3 Reconciliation: Quantile Regression

Figure 3 reveals two strikingly different quantile patterns. For AI/ML, the premium rises from 25.3% at Q10 to 34.8% at Q75, then partially retreats to 33.1% at Q90—an inverted-U shape that reflects a concentration of premium among high-tier technical positions. For programming, the pattern is *opposite*: the premium falls monotonically from 31.2% at Q10 to 16.1% at Q90. This divergence between AI/ML and programming reflects their different market positioning: AI/ML is disproportionately demanded in elite, high-wage roles, while programming skills have diffused broadly into mid-to-lower-wage positions where they remain scarce relative to demand.

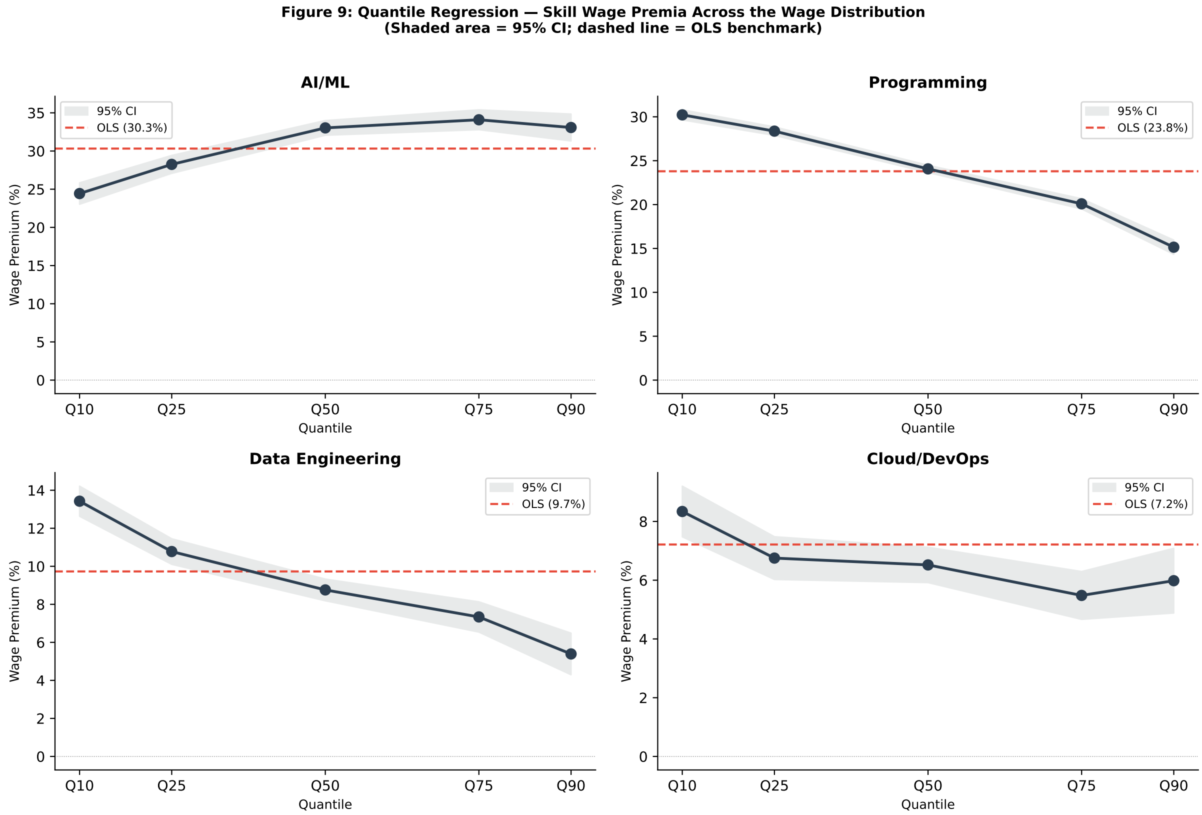

Figure 3: Quantile regression coefficients for four digital skills across the wage distribution ($\tau = 0.10, 0.25, 0.50, 0.75, 0.90$). Shaded band = 95% confidence interval. Dashed horizontal line = OLS benchmark. All specifications include education, experience, city tier, and industry fixed effects.

The quantile patterns together explain the polarization paradox. As programming skills diffuse to lower-wage occupations, they command large premia precisely in those occupations (scarcity rents), while the premium at the top compresses as the high-end supply of programmers deepens. For AI/ML, the premium remains highest in the upper-middle of the distribution where the most technically demanding roles cluster. Both patterns imply that *extensive-margin* diffusion of skills co-exists with *intensive-margin* amplification at specific wage quantiles.

6.4 Demand–Premium Decomposition

Figure 4 overlays the annual demand rate (bars) and conditional wage premium (line) for the three key skills. A striking *divergence* is visible for AI/ML: demand grows monotonically from 0.9% to 1.8% between 2016 and 2025, while the premium first spikes then partially retreats. This inverse co-movement—classic in competitive labor market models— suggests that demand-side expansion is being met by a partially lagged supply response, compressing the premium as the pool of qualified workers grows.

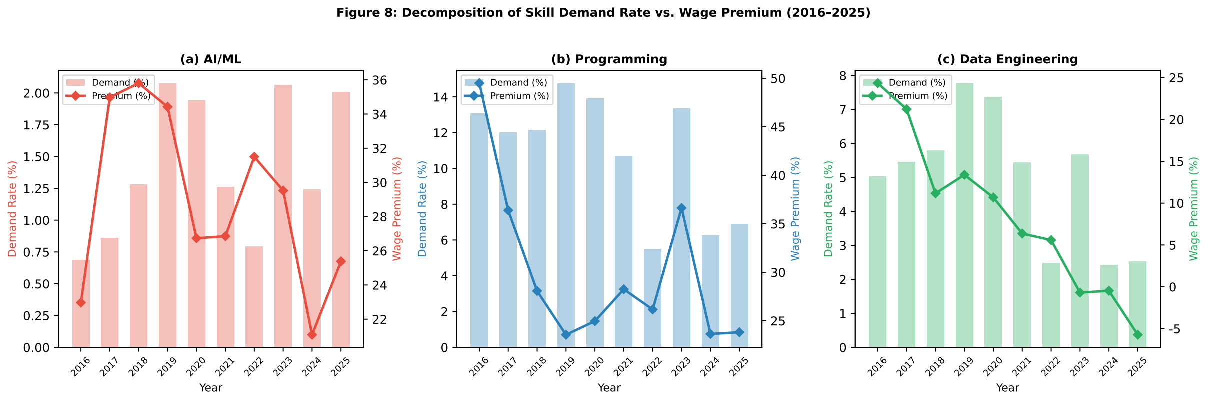

Figure 4: Annual skill demand rate (bars, left axis) and conditional wage premium (line, right axis) for AI/ML, Programming, and Data Engineering, 2016–2025. Divergence between demand growth and premium compression is consistent with lagged supply response.

6.5 Industry Heterogeneity

Figure 5 plots each industry's digital-skill rate against its AI/ML wage premium (bubble size proportional to N). Two clusters emerge. *Digital-intensive industries* (internet, software, electronics) show moderate digital rates (20–40%) but AI premia of 28–41%. *Traditional industries* (postal services, retail, agriculture) show low digital rates (5–15%) but *higher* AI premia (33–45%), consistent with scarcity rents: in sectors where AI workers are rare, firms pay heavily to attract them. The negative cross-industry correlation between digital-skill rate and AI premium ($r = -0.34$) directly reflects this scarcity logic.

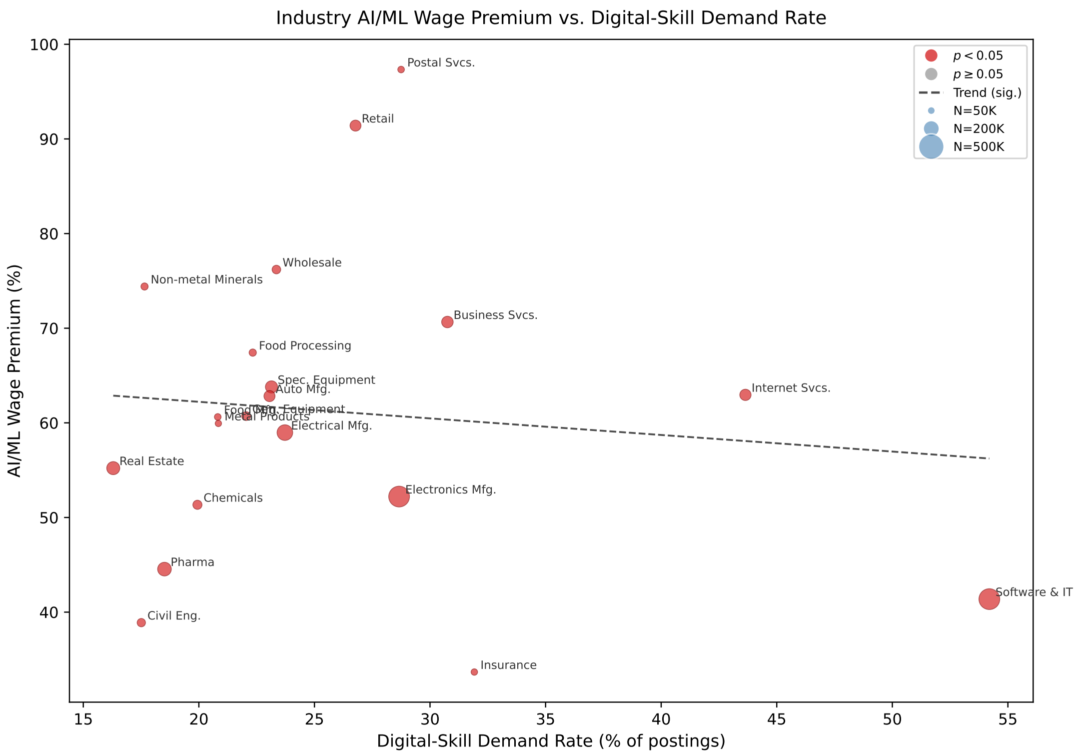

Figure 5: Industry AI/ML wage premium vs. digital-skill demand rate. Red bubbles: $p < 0.05$; gray: not significant. Bubble area $\propto$ number of observations. Dashed line: fitted trend for significant industries. The negative slope indicates scarcity rents in low-digital industries.

7 Robustness Checks

7.1 Subsample Stability

Table 3 reports estimates of Equation (1) across eight subsamples defined by time period, sector, city tier, and education. All specifications include city×year fixed effects and firm-clustered standard errors.

Pre- vs. post-2020. All premia are positive and significant in both sub-periods, but data engineering and cloud premia fall sharply post-2020 (from 13.6% to 4.2% and 7.0% to 1.9%, respectively). AI and programming premia are more persistent—roughly 24% and 20% post-2020. This heterogeneity suggests that data engineering and cloud skills diffused faster (compressing premia more) while AI/ML remained scarce.

Low-education scarcity premium. Among postings requiring below-bachelor's education, the programming premium is *larger* (33.3%) than among postings targeting high-education workers (18.0%). This counter-intuitive result reflects scarcity rents: when a job

Table 3: Robustness: Subsample Estimates of Skill Wage Premia

Subsample	N	Skill Wage Premium ($e^{\beta} - 1$) AI/ML	Programming	Data Eng.	Cloud/DevOps	R^2
Full sample (baseline)	820,807	26.0%***	24.8%***	9.9%***	3.8%***	0.189
		(0.96)	(0.48)	(0.55)	(0.62)	
Panel A: Time period						
Pre-2020 (2016–19)	433,018	29.7%***	27.6%***	13.6%***	7.0%***	0.156
		(1.54)	(0.61)	(0.72)	(0.93)	
Post-2020 (2020–25)	387,789	23.7%***	20.4%***	4.2%***	1.9%***	0.270
		(0.98)	(0.63)	(0.64)	(0.61)	
Panel B: Industry sector						
Manufacturing	378,413	26.7%***	27.1%***	8.9%***	2.5%***	0.181
		(1.28)	(0.65)	(0.93)	(0.87)	
Services	285,803	24.9%***	21.8%***	8.9%***	3.5%***	0.213
		(1.52)	(0.73)	(0.70)	(0.91)	
Panel C: City tier						
Tier-1 cities	285,151	26.3%***	25.7%***	10.5%***	3.7%***	0.219
		(1.43)	(0.60)	(0.68)	(0.77)	
Tier-2/3 cities	535,656	25.6%***	23.2%***	8.6%***	3.6%***	0.168
		(1.11)	(0.64)	(0.71)	(0.79)	
Panel D: Education level						
High edu ($\geq$ Bachelor's)	397,319	19.2%***	18.0%***	6.8%***	4.7%***	0.188
		(0.88)	(0.52)	(0.56)	(0.69)	
Low edu ($<$ Bachelor's)	423,488	25.7%***	33.3%***	16.6%***	2.9%***	0.106
		(2.51)	(0.76)	(0.98)	(0.89)	
City×Year FE		Yes (within-transformation)				
Industry FE		Yes (21 categories)				

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Firm-clustered standard errors (SE×100) in parentheses. Coefficients reported as $(e^{\beta} - 1) \times 100\%$.

Notes: All eight subsamples include city×year fixed effects and firm-level clustered standard errors. Panel A reveals premium compression for data engineering and cloud post-2020, while AI and programming remain robust. Panel D shows that low-education workers face larger programming premia, consistent with scarcity rents when programming skills appear in otherwise low-skill positions.

in a traditionally low-skill context requires programming competency, the rare candidate commands an outsized premium.

7.2 Extreme-Value Sensitivity

Table 4 shows that AI and programming premia are robust to progressive salary trimming. With no trimming the AI premium is 26.0%; trimming to the 5th–95th percentile reduces it to 20.2%; trimming to the 10th–90th reduces it to 14.9%. The premium remains economically and statistically significant across all specifications, though the magnitude falls as the most extreme high-paying postings are removed.

Table 4: Robustness: Salary Trim Sensitivity

		Skill Wage Premium ($e^{\beta}-1$)			
Salary Trim	N	AI/ML	Programming	Data Eng.	R^2
No trim [0%, 100%]	820,807	26.0%***	24.8%***	9.9%***	0.189
		(0.96)	(0.48)	(0.55)	
[1%, 99%]	804,975	25.8%***	23.8%***	9.9%***	0.222
		(0.89)	(0.43)	(0.49)	
[5%, 95%]	739,391	20.2%***	23.4%***	9.7%***	0.250
		(0.70)	(0.40)	(0.43)	
[10%, 90%]	671,137	14.9%***	21.9%***	9.0%***	0.220
		(0.64)	(0.34)	(0.38)	
City×Year FE	Yes (within-transformation)				
Industry FE	Yes (21 categories)				

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Firm-clustered standard errors (SE×100) in parentheses. Coefficients reported as $(e^{\beta}-1)\times 100\%$.
Notes: Controls as in Table 3. Progressive trimming of the salary distribution shows that AI/ML premia decline more steeply than programming premia under extreme-value removal, suggesting that AI premia are partially driven by high-paying outlier postings. Data engineering premia remain stable across all trim specifications.

7.3 Quantile vs. OLS Comparison

As shown in Figure 3, OLS coefficients lie close to the median quantile estimate for AI/ML, confirming that our baseline results capture the typical effect in the sample. For programming, the OLS estimate (24.8%) corresponds closely to the Q25–Q50 range, reflecting that the majority of programming postings cluster in the lower-to-middle wage segment. The key robustness lesson is that OLS *masks* important distributional heterogeneity across skills: while AI/ML premia concentrate in the upper-middle of the distribution, programming premia are disproportionately large for lower-wage workers—a distinction that average-effect analysis cannot reveal.

8 Conclusion

This paper documents three new stylized facts about digital skill markets in China.

First, digital-skill wage premia are not static. The AI/ML premium followed a hump-shaped trajectory over 2016–2025, consistent with a demand shock followed by a lagged supply response. Temporal dynamics matter: cross-sectional studies that sample near a peak or trough will yield very different conclusions.

Second, skill combinations are sub-additive. The joint premium for holding multiple digital skills is consistently below the sum of individual premia. The largest sub-additivity appears between programming and data engineering (-22.9 pp), reflecting task overlap. This finding challenges O-ring theories of complementarity and has direct implications for optimal skill investment: students should invest deeply in one high-premium skill rather than spreading effort across many.

Third, skill-demand diffusion and wage inequality are decoupled. Industry concentration of digital-skill demand has fallen by 30% since 2016, yet wage inequality has widened at the top of the distribution. Quantile regression reveals heterogeneous patterns: AI/ML premia peak in the upper-middle wage quantiles, while programming premia fall monotonically from Q10 to Q90. Universal diffusion of digital skills does not guarantee compressed inequality when the returns to those skills are heterogeneously distributed across the wage distribution.

These findings point toward three policy priorities. (i) *Curriculum timing*: because premia are time-varying, higher-education systems that adapt degree programs with 3–5 year lags may inadvertently overproduce workers for premia that are already compressing. (ii) *Depth over breadth*: the sub-additivity result supports focused mastery over broad digital literacy in skill-formation policy. (iii) *Redistribution alongside diffusion*: policy-makers who assume that broader digital-skill demand will automatically narrow wages are likely to be disappointed; complementary redistributive instruments are needed.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Card, D. and Ashenfelter, O., editors, *Handbook of Labor Economics*, volume 4B, pages 1043–1171. Elsevier.

Acemoglu, D. (2002). Technical change, inequality, and the labor market. *Journal of Economic Literature*, 40(1):7–72.

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

Autor, D., Levy, F., and Murnane, R. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.

Deming, D. and Noray, K. (2018). *Earnings Dynamics, Changing Job Skills, and STEM Careers*. NBER Working Paper 23538.

Ge, S. and Yang, D. (2011). Labor market developments in China: A neoclassical view. *China Economic Review*, 22(4):611–625.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–26.

Katz, L. and Murphy, K. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *Quarterly Journal of Economics*, 107(1):35–78.

Kremer, M. (1993). The O-ring theory of economic development. *Quarterly Journal of Economics*, 108(3):551–575.

Li, B., Liu, C., and Ma, L. (2021). Digital skills and wage inequality in China: Evidence from job vacancy data. *China Economic Review*, 70:101694.

Xing, C., Zhang, J., and Xu, L. (2018). Education, marriage, and earnings: Evidence from China. *China Economic Review*, 51:143–158.

Zhang, H. and Feng, S. (2022). The digital labor market and wage inequality in urban China. *World Development*, 151:105764.