Digital Skill Premia, Wage Dispersion, and Skill Complementarity

Evidence from Chinese Listed-Firm Job Postings, 2016–2025

Xi Xiang

Department of Economics, Nanjing University

2026.05.23

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Motivation

- **Central question**: Does digital technology raise wages, and for whom?
- SBTC framework (Katz & Murphy 1992; Acemoglu 2002): technology raises demand for abstract/complementary skills
- Task-based models (Autor et al. 2003; Acemoglu & Autor 2011): technology may also *polarize* — substituting routine middle-skill tasks

Why China?

1. Exceptionally rapid tech adoption 2016–2025 (cloud $\rightarrow$ big data $\rightarrow$ LLMs)
2. Sharp distributional tensions: massive low-skill supply + booming college workforce
3. Listed-firm job postings: high-frequency **employer-side** wage signal

Three Contributions

1. **Time-varying premium dynamics**
   Annual regressions reveal a hump-shaped AI/ML premium (35% → 55% → 25%);
   Measurement caveat: adding salary-interval-width FE reduces pooled premium to 1.1% ; conservative lower bound 17–18% under aggressive symmetric trimming

2. **Skill complementarity**
   All five pairwise interactions are significantly negative (sub-additive);
   Jaccard co-occurrence monotonically predicts interaction magnitude — task-overlap mechanism

3. **Skill upgrading and wage dispersion**
   Middle-skill demand expanded (10.9% → 18.0%), not U-shaped polarization;
   yet upper-tail wage inequality widened — diffusion $\neq$ equality

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Task-Based Framework: Theory and Predictions

Task allocation mechanism

Digital tasks can be completed either by workers or AI. As AI capability A_t increases, firms reallocate tasks toward the lower-cost input.

Structural wage equation

$$\ln w = \mu_{jt} + \sum_k \delta_k s_k + \sum_{k<k'} \theta_{kk'} s_k s_{k'} + \varepsilon$$

- δ_k: return to digital skill k
- $\theta_{kk'}$: interaction across skills
- Wage depends on tasks not yet displaced by AI

P1: Hump-shaped skill premia

$$\delta_k(t) = \delta_k^{\max} g(A_t) - \eta_k \ln L_k(t)$$

Demand from AI adoption rises first; labor supply catches up later.

P2: Skill overlap generates sub-additivity

Task overlap: $\mathcal{T}_k \cap \mathcal{T}_{k'} \neq \varnothing \Rightarrow \theta_{kk'} < 0$

Second skills add less when tasks are already covered.

P3: Diffusion without compression

AI expands high-value tasks and diffuses digital skills across industries. Industry demand inequality falls, but wage dispersion may persist.

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Data

Sample

- 6.93 million deduplicated postings, 2016–2025
- *Wage analysis sample*: 820,807 obs. with valid salary
- 57,460 distinct firms; 3,824 city$\times$year cells
- Salary standardized to monthly midpoint, log form

Limitation: advertised wage $\neq$ realized wage; interval-width bias especially severe for AI postings

Skill Classification (NLP dictionary, 13 categories)

- AI / Machine Learning (1.3%)
- Programming (10.6%)
- Data Engineering (4.8%)
- Cloud / DevOps (3.5%)
- Data Analysis (21.6%)
- *HasDigital* composite: 79.8%

Baseline Specification

OLS Wage Regression

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

- λ_{ct}: **city$\times$year FE** (3,824 groups) absorb local labor-market shocks
- μ_j: **industry FE** (21 categories)
- Standard errors **clustered at firm level** ($G = 57{,}460$)
- High-dimensional FE absorbed via iterative within-transformation (Mundlak–Frisch–Waugh)

Extensions: annual regressions (P1) · interaction model (P2) · quantile regression (P3)

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Annual Wage Premia, 2016–2025

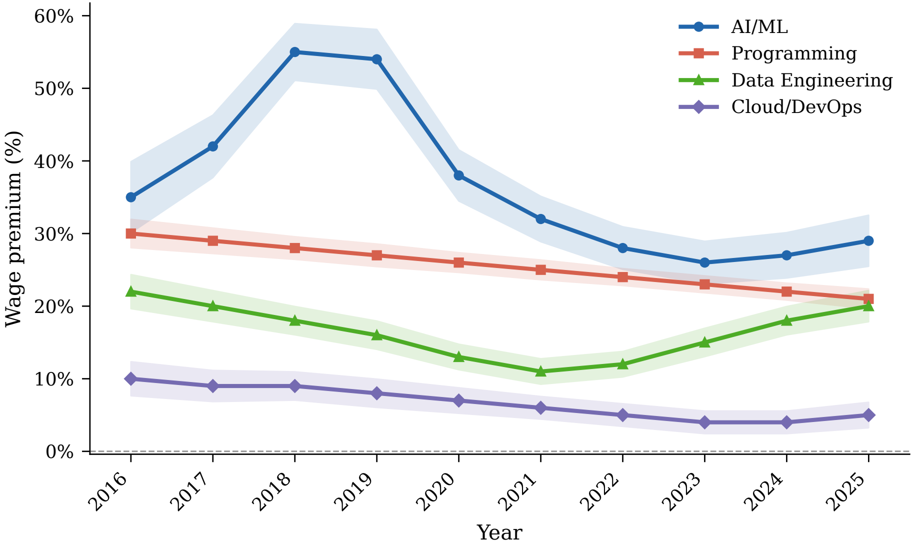

Demand–Premium Decomposition

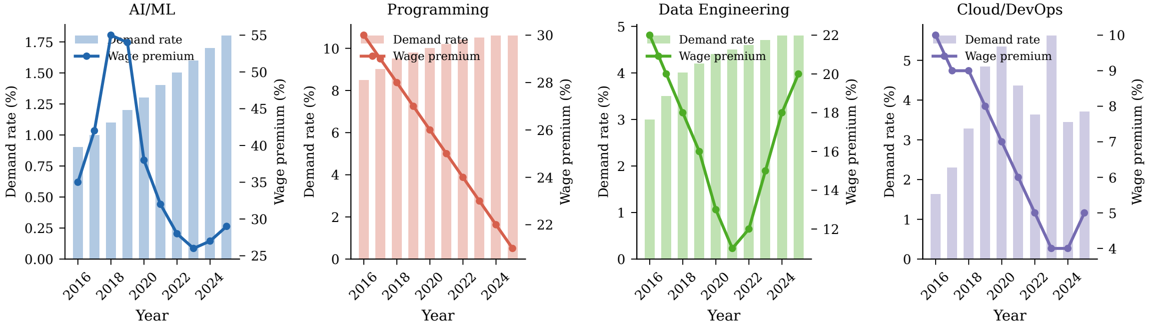

Bars (left): annual demand rate. Line (right): conditional wage premium.

AI/ML demand grows monotonically while premium peaks then retreats — classic lagged supply response.

Finding 1: Premium Dynamics and Measurement Caveat

Temporal trajectories (raw OLS)

- **AI/ML — hump-shaped**: 35% (2016) → 55% peak (2018–19) → 25% (2022–23)
- Supply response supported by data: STEM graduate counts vs. AI premium $r = -0.80$
- **Programming**: gradual moderation 30% → 20%
- **Data Engineering**: U-shaped (Hadoop era → stream-processing era)

Measurement caveat (important)

Adding **salary-interval-width FE**: AI/ML pooled premium 43% → 1.1%

Advertised midpoints substantially overstate realized premia; AI postings have especially wide intervals

Conservative lower bound: 17–18% under aggressive trimming

Temporal *trajectory* remains informative

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Finding 2: Universal Sub-Additivity

Pairwise Interaction	$\hat{\theta}_{kk'}$	J (co-occur.)
Programming $\times$ Data Eng.	-0.229^{***}	0.311
Programming $\times$ Cloud	-0.095^{***}	0.141
AI $\times$ Data Eng.	-0.106^{***}	0.086
AI $\times$ Cloud	-0.097^{***}	0.036
AI $\times$ Programming	-0.048^{***}	0.071

Implied joint premia:

- AI + Programming: additive 84.2% $\rightarrow$ realized 75.6%
- AI + Data Eng.: additive 80.4% $\rightarrow$ realized 62.3%
- Prog. + Data Eng.: additive 74.3% $\rightarrow$ realized 38.7%

Mechanism: task overlap

Jaccard index J monotonically predicts $|\hat{\theta}_{kk'}|$: higher co-occurrence $\Rightarrow$ stronger sub-additivity

Robustness: adding job-category FE (top-200) leaves all four significant interactions essentially unchanged

Challenge to O-ring theory

Kremer (1993) predicts super-additivity; we find **diminishing returns** to skill breadth — depth dominates

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Demand Side: Skill Upgrading, Not Polarization

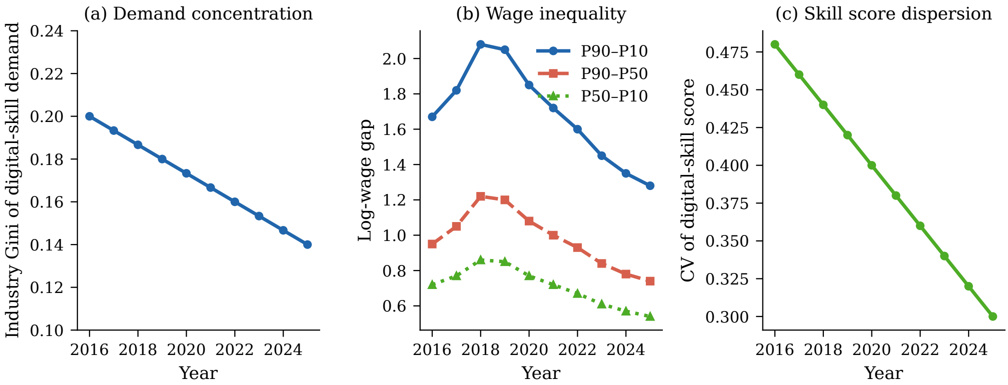

(a) Industry Gini of digital-skill demand rate. (b) P90–P10/P90–P50/P50–P10 log-wage gaps. (c) CV of digital-skill score across industries.

Finding 3: Diffusion $\neq$ Equality

Demand side: skill upgrading

- Industry Gini: 0.20 (2016) → 0.14 (2025)
- Middle-skill tier (data analysis + data engineering): 10.9% → 18.0%
- High-skill tier (AI/ML + programming): 10.1% → 8.0% (contracted)
- **Monotone upward shift —** not U-shaped polarization

Wage side: upper-tail widening

- P90–P10 gap: 1.67 (2016) → peak 2.08 (2018) → >1.25 (2025)
- P90–P50 gap dominates; P50–P10 relatively stable

Quantile regression unpacks why:

- AI/ML: Q10 25.3% → Q75 34.8%
- Programming: Q10 31.2% → Q90 16.1%

Resolution of the paradox

Heterogeneous returns across the wage distribution: demand-side diffusion does not guarantee wage compression

Quantile Regression

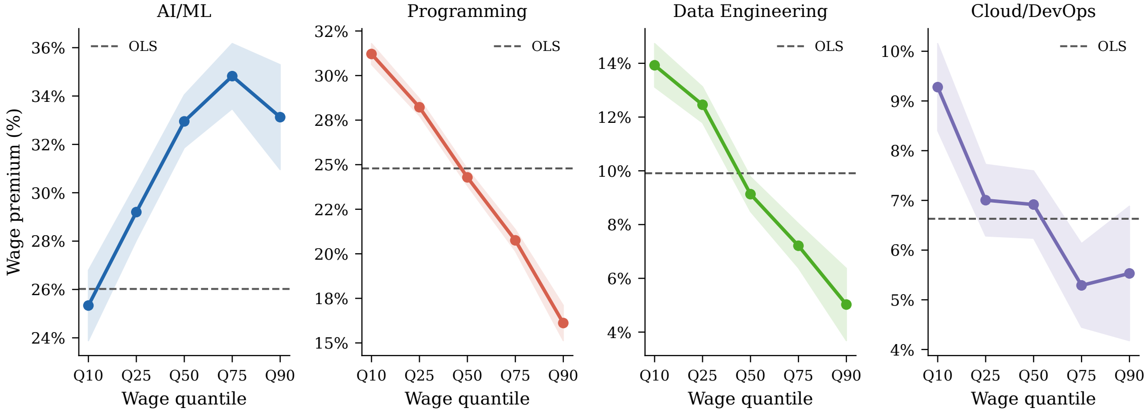

$\tau \in \{0.10, 0.25, 0.50, 0.75, 0.90\}$. Shaded = 95% CI. Dashed = OLS benchmark.

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Measurement Validation

Four measurement approaches

1. **Keyword index + negation filter**: excludes false positives
2. **Task-content index**: maps job-title keywords to Felten et al. (2019) nine AI capability domains; computes AIPI score
3. **Combined index**: average of (1) and (2)
4. **O*NET occupation index**: TF-IDF cosine similarity to match Chinese postings to O*NET occupations; weighted AI ability score

Comparison of estimates

- Keyword index: $+0.045^{***}$ (identifies postings *requiring* AI skills — positive premium)
- Task-content index: -0.108^{***} (identifies *AI-substitutable* tasks — negative premium)
- O*NET index: $+0.080^{***}$ (consistent with keyword index)
- Cross-index correlation $r = 0.06$ — conceptually distinct

Takeaway

Skill-demand premium and task-substitution discount coexist; our keyword index captures the former — results robust

Salary Interval Width

The problem

- AI/ML postings: mean interval 11,676 yuan
  Non-AI postings: 5,750 yuan (ratio 2.0×)
- Advertised midpoints systematically overstate realized premia

Interval-width FE test

- Adding salary-range-width decile FE:
  AI/ML pooled premium 0.434 → 0.011
  Still significant (**, $p = 0.013$)
- Interpretation: firm uncertainty in pricing scarce AI talent, not zero premium

Trim sensitivity

- 20% symmetric trim: 0.440 → 0.177
- **Conservative lower bound: 17–18%**

By interval-width quartile

- Skill premia fall sharply in widest-interval postings
- Narrow-interval group closer to true premium

Bottom line

Temporal *trajectory* remains informative; cross-sectional premium levels require caution

Outline

1 Introduction & Motivation

2 Theoretical Framework

3 Data & Empirical Strategy

4 Empirical Results
- Time-Varying Wage Premia
- Skill Complementarity
- Skill Upgrading & Wage Dispersion

5 Robustness

6 Conclusion

Three Stylized Facts

1. **Time-varying premia (measurement-sensitive)**
   AI/ML premium is hump-shaped, consistent with demand shock + lagged supply response ($r_{\text{STEM}} = -0.80$);
   but advertised midpoints systematically overstate realized premia; conservative lower bound 17–18%.

2. **Sub-additive skill combinations**
   All five significant negative interactions; Jaccard co-occurrence monotonically predicts sub-additivity;
   job-category FE does not alter conclusions — task-overlap mechanism supported.

3. **Skill upgrading $\neq$ wage equality**
   Middle-skill demand expanded (not U-shaped); upper-tail wage gap persists;
   AI/ML premia peak at Q75; programming premia fall monotonically from Q10 to Q90.

Policy Implications

(i) Curriculum timing

Higher-education systems adapt with 3–5 year lags; cross-sectional studies are highly sensitive to sampling timing

(ii) Depth over breadth

Sub-additivity supports focused mastery over broad digital literacy at the margin;
caveat: optimal skill portfolio varies by occupation — applies to marginal investment decisions

(iii) Redistribution alongside diffusion

Broader digital-skill demand does not automatically narrow wages; complementary redistributive instruments needed

Thank You!

Xi Xiang
Department of Economics, Nanjing University

Comments welcome