Weak bundle and strong bundle: How AI redraws job boundaries?

Luis Garicano
(LSE & CEPR)

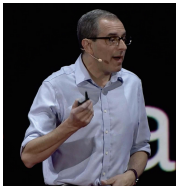
Joint with
Jin Li (HKU), Yanhui Wu (HKU), Leo Yang Yang (HKBU), Zhenyu Zhu (HKU)

Santiago, November 2026

Debate about the impact of AI on the future of work

"There will come a point where no job is needed. You can have a job if you want to have a job for personal satisfaction, but AI will be able to do everything."





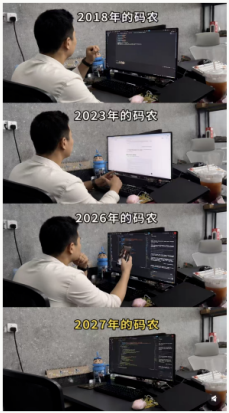
David Autor
(MIT)

"The industrialized world is **awash in jobs, and it's going to stay that way**. [...] AI will change the labor market, but not in the way Musk and Hinton believe. Instead, **it will reshape the value and nature of human expertise**."

Disappearing jobs

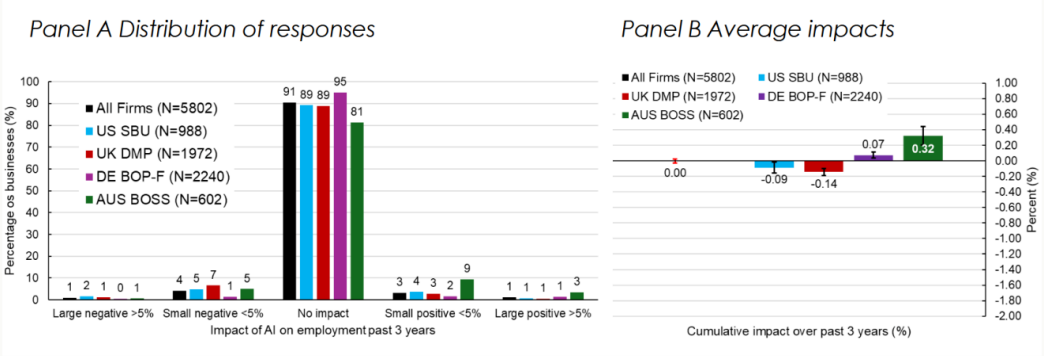






The coder in 2018, 2023, 2026, and 2027.

Impact of AI on recent employment



Yotzov et al. (2026), Figure 8.

Yotzov et al. (2026), Figure 8. Firm surveys in the US, UK, Germany and Australia: most firms report no employment effect of AI over the past three years, and the average cumulative effect is close to zero.

Task-based approach: AI exposure

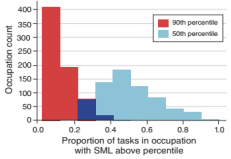


FIGURE 1. FREQUENCY COUNTS OF OCCUPATIONAL TASK PROPORTIONS ABOVE NINETIETH, SEVENTY-FIFTH, AND FIFTIETH PERCENTILES

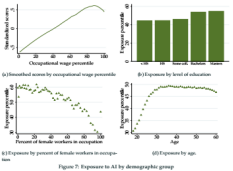
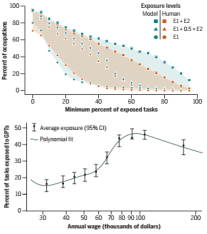
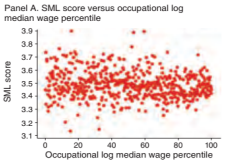


Figure 7: Exposure to AI by demographic group

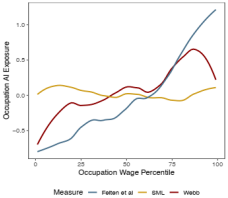
Webb 2020, "What Can Machines Learn and What Does It Mean for Occupations and the Economy?



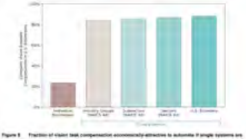
Eloundou, Manning, Mishkin & Rock 2024, "GPTs are GPTs: Labor market impact potential of LLMs"



Brynjolfsson & Mitchell, 2018, "What Can Machines Learn and What Does It Mean for Occupations and the Economy?



Acemoglu, Autor, Hazell, & Restrepo 2022, "Artificial Intelligence and Jobs: Evidence from Online Vacancies"



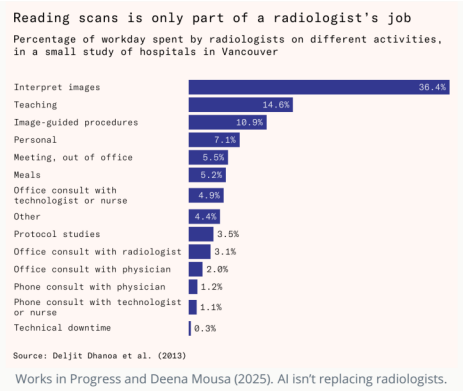
© David Autor and Neil Thompson, 2024

Svanberg, Li, Fleming, Goehring & Thompson 2024, "Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?"

Brynjolfsson, Mitchell and Rock (2018); Webb (2020); Eloundou, Manning, Mishkin and Rock (2024); Brynjolfsson and Mitchell (2017); Acemoglu, Autor, Hazell and Restrepo (2022); Svanberg, Li, Fleming, Goehring and Thompson (2024).

AI exposure $\neq$ AI replacement

AI exposure: proxy for being automated at the task level; but labor markets price jobs, not tasks.



Works in Progress, Deena Mousa (2025), "AI isn't replacing radiologists." Reading scans is about a third of a radiologist's workday.

Why was Hinton's prediction wrong?





In 2016, Geoffrey Hinton:

"Stop training radiologists?"

By the end of 2025,

- 700+ FDA-cleared radiology AI models (about 75% of all medical AI devices).
- Yet U.S. radiology residency spots hit a record high and radiologists' pay rose to about $520,000 on average.

Strength of bundle

The effect of AI displacement depends not only on the proportion of tasks taken by AI but also on the strength of a task bundle.

Weak bundle:
Easy to break with little loss of value.



Strong bundle:
Value depreciates sharply if broken up.

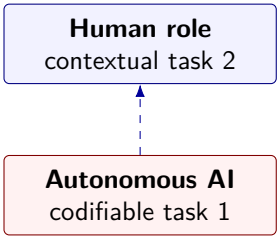


Plan

- **A price theory of task bundling**
- Evidence from labor markets
- Messy jobs as protection of employment

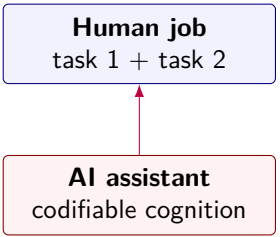
How to model bundle strength?

Weak bundle: low separation loss



Separate agents

Strong bundle: high separation loss



AI stays inside the bundle

The same exposure can produce very different labor-market responses.

The model: technology and time allocation

A unit of service requires two tasks, 1 and 2. Final output is Cobb–Douglas:

$$y = q_1^{\alpha} q_2^{1-\alpha}, \qquad \alpha \in (0,1).$$

- Worker skills: $S = [\underline{s}_1, \overline{s}_1] \times [\underline{s}_2, \overline{s}_2] \subset \mathbb{R}^2_+$, density $g(s_1, s_2) > 0$.
- Time allocation: $t_1 + t_2 \leq 1$, $t_1 \geq 0$, $t_2 \geq 0$.
- Gross bundled output (pre-AI benchmark)

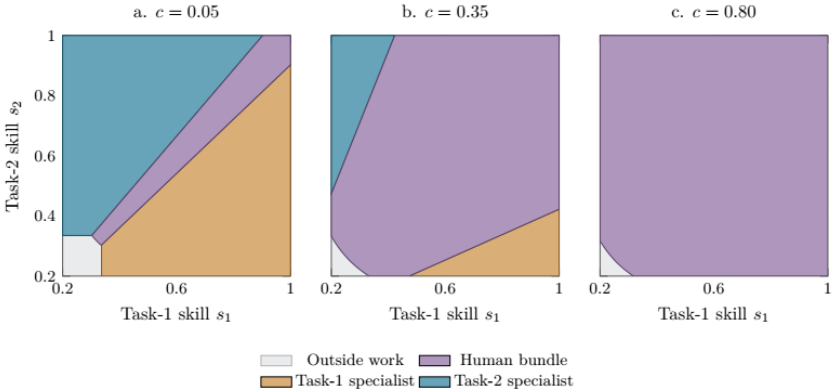
$$y_B^0(t, s) = (t_1 s_1)^{\alpha} (t_2 s_2)^{1-\alpha}.$$

- Gross unbundled output (pre-AI benchmark)

$$y_U^0(s) = (1 - c)\, s_1^{\alpha} s_2^{1-\alpha}, \qquad 0 \leq c < 1.$$

The coordination loss c summarizes shared context, liability, and knowledge spillovers lost when the two tasks are supplied by separate agents.

Coordination cost and occupational sorting



Pre-AI benchmark with human task-1 and task-2 specialists. As c rises (left to right), the bundled human job takes over the skill space.

Bundled and unbundled production with AI

Gross bundled output

$$y_H^g(t_1, t_2, s, m) = \left[\mu_H(m)^{1-\eta}(t_1 s_1)^{\eta}\right]^{\alpha}(t_2 s_2)^{1-\alpha}, \qquad \eta \in (0,1).$$

Optimal bundled time allocation

$$t_1^* = \frac{\alpha\eta}{\kappa}, \qquad t_2^* = \frac{1-\alpha}{\kappa}, \qquad \kappa \equiv \alpha\eta + 1 - \alpha = 1 - \alpha(1-\eta).$$

Gross unbundled output

$$y_U(s_2, m; c) = (1-c)\,\mu_A(m)^{\alpha} s_2^{1-\alpha}.$$

The effective-machine frontier

$$\Gamma(m) \equiv \frac{\mu_A(m)^{\alpha}}{\mu_H(m)^{\alpha(1-\eta)}}, \qquad \frac{d}{dm}\Gamma(m) > 0.$$

Human earnings and sorting

Under competitive pricing of the final-good technology,

$$p_1 q_1 = \alpha P y, \qquad p_2 q_2 = (1-\alpha) P y.$$

Human revenue shares. Fix (P, m, c). Net human earnings are

$$w_H(P, s, m) = \kappa\, P\, y_H(s, m), \qquad w_U(P, s_2, m; c) = (1-\alpha)\, P\, y_U(s_2, m; c).$$

Since $\kappa = 1 - \alpha(1-\eta) > 1-\alpha$, the bundled human keeps a strictly larger share of revenue.
Sorting cutoff. Fix (P, m, c). There exists a cutoff

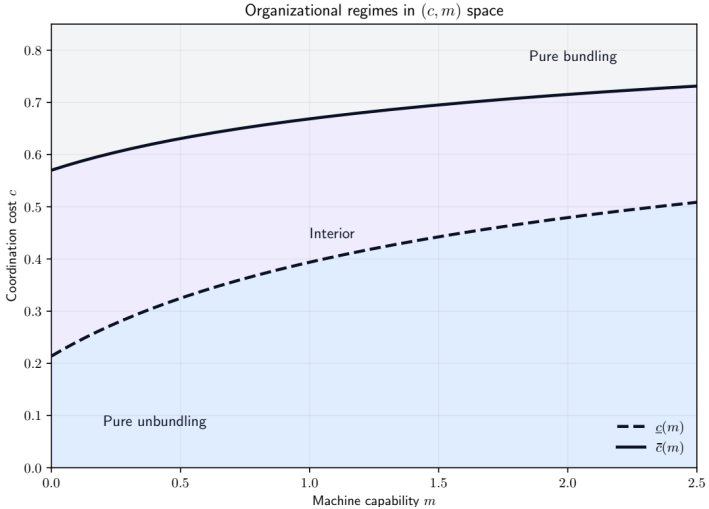
$$s_1^*(m, c) = \left[\frac{(1-\alpha)(1-c)}{\kappa \Omega_H}\, \Gamma(m) \right]^{1/(\alpha\eta)}$$

such that

$$w_U(P, s_2, m; c) \geq w_H(P, s, m) \quad \Longleftrightarrow \quad s_1 \leq s_1^*(m, c).$$

s_1^* is increasing in m and decreasing in c.

Occupational choice



Regimes in (c, m) space. Below $\underline{c}(m)$: pure unbundling; above $\bar{c}(m)$: pure bundling; between: both forms coexist. Better AI shifts both cutoffs up.

Participation in competitive equilibrium (fixed P)

Reservation utility $\bar{U} > 0$; the relevant relative price is $\tau \equiv \bar{U}/P$.

Worker welfare

$$V(s; P, \bar{U}, m, c) = \max\{\bar{U},\ w_H(P, s, m),\ w_U(P, s_2, m; c)\}.$$

A competitive equilibrium with fixed product price is a measurable partition

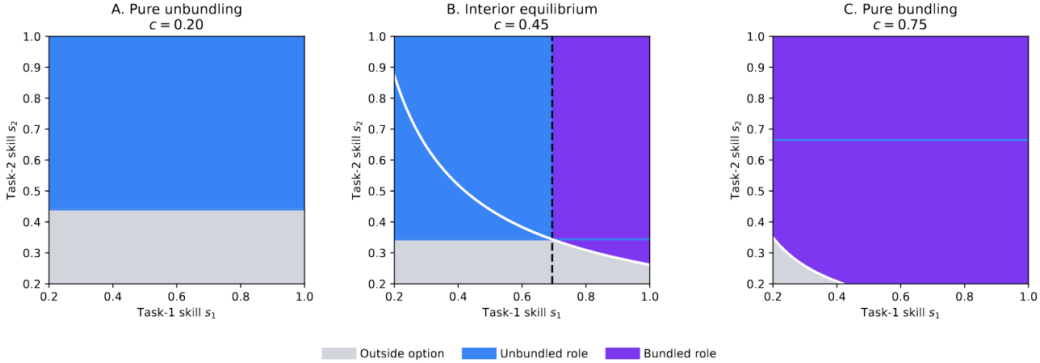
$$\big(\mathcal{O}(\tau, m, c), \mathcal{U}(\tau, m, c), \mathcal{H}(\tau, m, c)\big).$$

$$\mathcal{U}(\tau, m, c) = \{(s_1, s_2) \in S :\ s_1 < s_1^*(m, c),\ s_2 \geq b_U(\tau, m, c)\},$$

$$\mathcal{H}(\tau, m, c) = \{(s_1, s_2) \in S :\ s_1 \geq s_1^*(m, c),\ s_2 \geq b_H(s_1; \tau, m)\}.$$

Accessibility. For $s_1 \leq s_1^*$: $b_U(\tau, m, c) \leq b_H(s_1; \tau, m)$. Once task 1 is peeled away, all that matters is whether the worker can supply task 2 .

Coordination cost and occupational choice



Occupational choice in skill space at $m = 0.8$. Gray: outside option; blue: unbundled (residual) role; purple: bundled role. In the interior case the vertical line is s_1^*, the horizontal line is b_U, the curve is $b_H(s_1)$.

General equilibrium (endogenous P)

Suppose the final good faces inverse demand $P(Y) = Y^{-1/\varepsilon}$, so $\tau = \bar{U}\, Y^{1/\varepsilon}$ is an equilibrium object. Unbundling releases the time a worker spent on task 1: task-2 capacity rises by $\kappa/(1-\alpha) > 1$. For a fixed worker type,

$$\frac{d \ln w_U^*}{dm} = \underbrace{\frac{d \ln P^*}{dm}}_{\text{price compression } (<0)} + \underbrace{\alpha \frac{d \ln \mu_A}{dm}}_{\text{direct productivity}}, \qquad \frac{d \ln w_H^*}{dm} = \frac{d \ln P^*}{dm} + \alpha(1-\eta)\frac{d \ln \mu_H}{dm}.$$

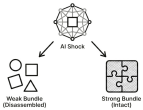
In pure unbundling, human employment falls iff $\frac{1}{\varepsilon}\frac{d \ln Y^*}{dm} > \alpha \frac{d \ln \mu_A}{dm}$.

Weak bundle: AI can redraw the job, release human time, and lower entry requirements. The response is large, but can be expansionary or contractionary depending on demand.

Strong bundle: AI remains inside the job. The capacity and entry shocks are muted, so employment and task content respond less.

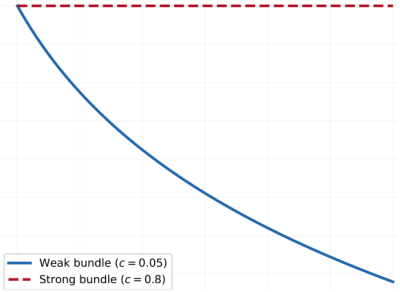
Tradeoff between productivity and employment

Weak bundle leads to split of tasks and gains from specialization.



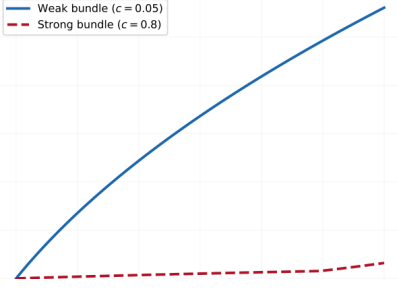
Strong bundle protects human jobs and gains from AI–human co-pilot.

Total human employment



Frontier AI capability m

Aggregate output



Frontier AI capability m

Two conditions. The weak-bundle decline requires inelastic demand ($\varepsilon < 1$; benchmark $\varepsilon = 0.85$). The strong bundle is protected until m reaches the interior regime ($m \approx 2.68$ in the benchmark), after which earnings and labor income fall.

Plan

- A price theory of task bundling
- **Evidence from labor markets**
- Messy jobs as protection of employment

How do we take the theory to data?

The central econometric prediction is not a universal employment decline. It is an interaction: the response to exposure should be attenuated by bundle strength.

	Weak bundle	Strong bundle
Task content of postings	Exposed task leaves; residual-only postings increase	Perimeter changes little
AI-tool requirements	Larger rise: buy more machine services	Smaller rise
Posted wages	Wage in narrower job falls: composition and reallocation, plus price compression	Price compression only; wages hold up better
Entry requirements	Fall: the residual role drops the task-1 requirement	Unchanged
Hiring (posting counts)	Larger response; sign depends on demand elasticity	Smaller response

But above all, we need to come up with a measure of bundle strength (and/or coordination cost).

Empirical Strtegy: From the model to the regression

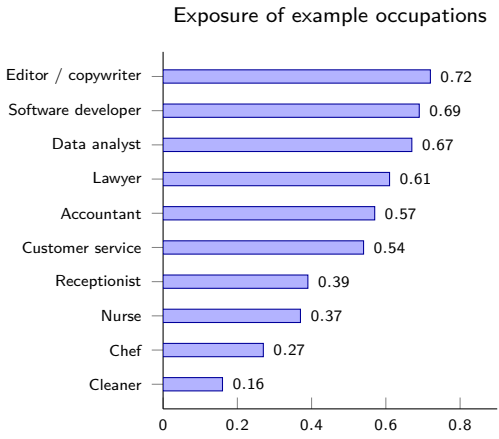
- Exposure $E_o \leftrightarrow \alpha$, the weight of codifiable work in the job. Bundle strength $M_o \leftrightarrow c$.
- Shock. The May 2024 price war cut the price of API tokens by an order of magnitude; DeepSeek-R1 (January 2025) raised quality at a lower price. Both lower the price of *autonomous* machine services (μ_A) more than of chat assistance (μ_H): $\Gamma(m)$ rises. The sorting cutoff s_1^* increases, by more the lower is c.
- Job Postings. Exposed task drops, residual rises: $\rho_U = N_U/(N_H + N_U)$; exposed-task share $= (1 - \rho_U)\times$ task-1 content of bundled posting
- Occupation-average posted wage $\bar{w} = W^{pop}/(N_H + N_U)$:including a within-form term, a composition term (lower-s_2 entrants), and a reallocation term, negative whenever $\bar{w}_U < \bar{w}_H$.
- Attenuation. For $c \geq \bar{c}(m)$ the response of all these objects to m is zero; only for $c < \bar{c}(m)$ it is signed. Hence the interaction coefficient makes the difference.

Data

- O*NET 31.0 Dataset
  - Survey-based descriptors of about 900 U.S. occupations: work context, work activities, knowledge
  - Used to measure the three sources of bundle strength
- Job Posting Dataset
  - Online job postings in China from five major platforms: 51job, Zhaopin, BOSS Zhipin, Liepin, 58.com
  - January 2023–July 2026: about 210 million postings after removing within-month duplicates
  - Each posting is parsed into tasks and requirements (a catalog of 65,200 items), together with wage, education, experience, location, and industry
  - About 1,200 occupations, each mapped to an O*NET-SOC occupation (1,214 of 1,265 matched, to 395 O*NET codes); white-collar sample of 952 occupations

Measure GenAI exposure

- Score every task in the catalog (44,479 tasks): can current large language models substantially accelerate or absorb its core output?
  - e.g., “conducting data analysis” 0.75; “greeting walk-in customers” 0.20.
- Occupation exposure E_o: mention-weighted average score over the tasks in its postings, 2023 Q4 (before the shock).
- GenAI-exposed tasks: score ≥ 0.6 and on-site presence < 0.3.
- Validation: correlates 0.73–0.79 with Eloundou et al. (2024) and Felten et al. (2021).



Measure bundle strength

- Decompose bundle strength into three sources, and use O*NET data to measure each source.

Source	**Description**	**O*NET Data**
Shared context	Task execution depends on information relevant to other tasks	Work Context
Liability	Dividing work can complicate responsibility for its consequences	Work Context
Across-task spillovers	Tasks draw on overlapping bodies of knowledge	Work Activities & Knowledge

- Conduct principal component analysis (PCA) on the three source measures and use the first principal component as an index of bundle strength.

Bundle strength: construction (background)

Shared context

- relational embeddedness (contact with others; face-to-face discussions; work with a group; deal with external customers; physical proximity)
- task discretion (determine tasks and priorities; freedom to make decisions; frequency of decision making).

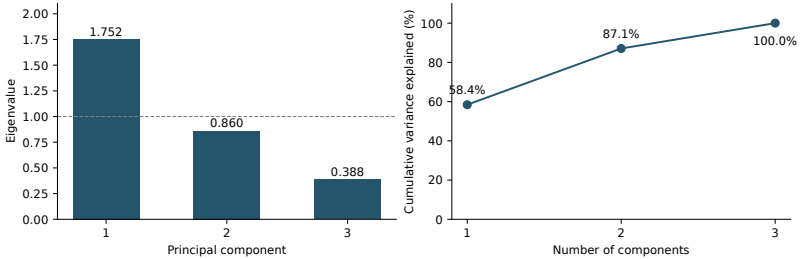
Liability

- consequence of error;
- impact of decisions on co-workers or company results;
- responsibility for health and safety, and for work outcomes, of other workers.

Across-task spillovers =knowledge overlap across an occupation's 41 work activities: each activity's knowledge profile over 33 domains, cosine similarity between activity pairs, averaged with importance weights.
Each index is standardized across O*NET occupations and matched with Chinese occupation.

PCA

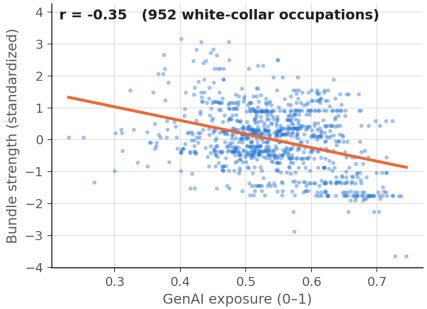
Nothing depends on PCA: an equal-weight average of the three standardized indices has correlation 0.996 with the first component, so the measure is robust to the weighting.



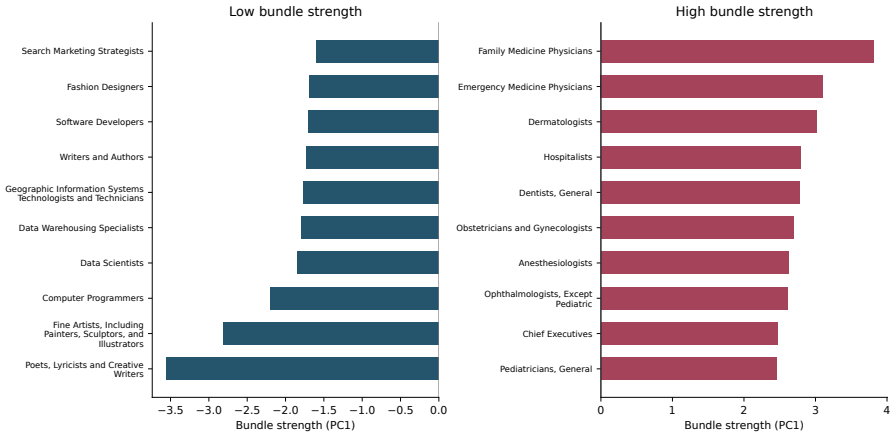
Loading on PC1	
Shared context	0.884
Liability	0.774
Across-task spillovers	0.609
Variance explained	58.4%

Bundle strength vs. GenAI exposure

- In the white-collar sample (952 occupations), bundle strength is moderately correlated with GenAI exposure ($r = -0.35$), but nearly uncorrelated with log wages (-0.06) and college-degree requirements (-0.08).
- Results are not driven by nonlinear exposure: with exposure^2 $\times$ month controls, the interaction stays significant.



Which occupations are weak and strong bundles?



Ten lowest and ten highest distinct O*NET occupations by bundle strength (PC1) in the white-collar sample. Weak bundles: search marketing, fashion design, software development, writing, data warehousing, data science, programming, performing arts. Strong bundles: physicians, dentists, obstetricians, anesthesiologists, chief executives, pediatricians.

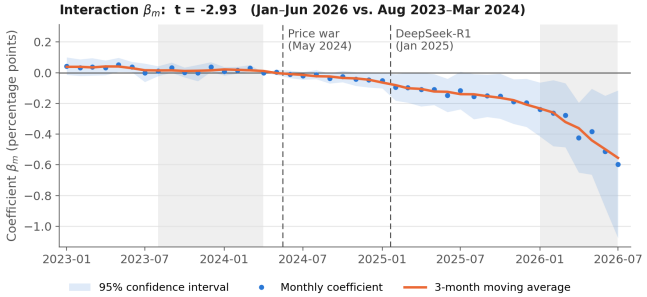
Econometric specification

$$y_{ut} = \alpha_u + \lambda_{pt} + \sum_{m \neq 2024\text{-}04} \mathbf{1}[t = m](\beta_m E_o M_o + \gamma_m E_o + \delta_m M_o) + \varepsilon_{ut}.$$

- Unit: occupation $\times$ platform; α_u and λ_{pt} are unit and platform $\times$ month FE
- E_o: GenAI exposure of occupation o; M_o is bundle strength; (standardized N(0,1)
- Reference period is April 2024, just before LLM price war
- β_m: how a one-SD increase in bundle strength changes the association between exposure and the outcome change from April 2024 to month m, per SD of exposure
- Sign warning. β_m is the strong-minus-weak difference β_m the opposite sign to the weak-bundle response: negative where weak bundles rise

Exposure × bundle strength: AI adoption

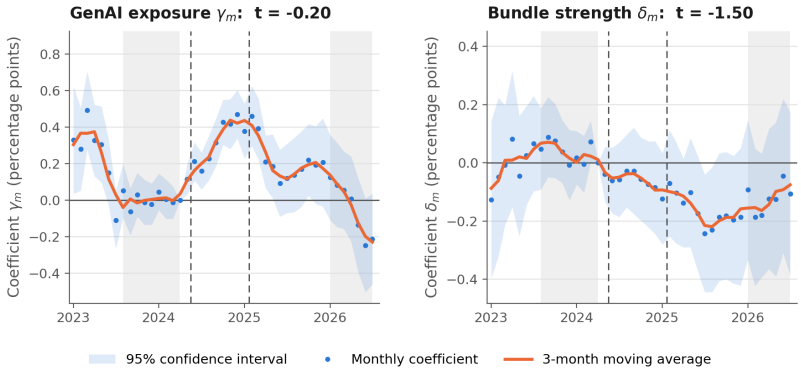
Share of postings requiring GenAI tools



- After the price war, exposed weak-bundle occupations added generative-AI tool requirements to their postings faster than equally exposed strong-bundle occupations: 0.37 pp more per SD of bundle strength by 2026 ($\beta_m = -0.37$, $t = -2.93$; survives detrending).

Exposure and bundle strength: task composition

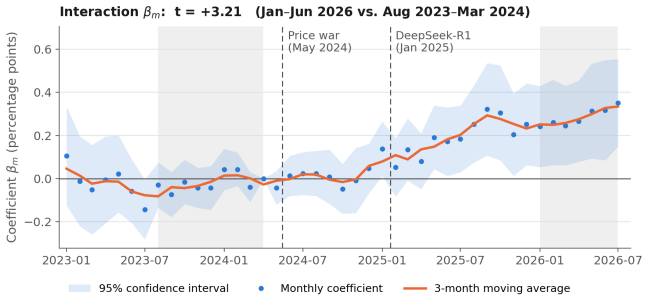
Share of GenAI-exposed tasks in postings



- Exposure by itself does not predict any change in the share of AI-exposed tasks that postings list ($t = -0.20$), and neither does bundle strength by itself ($t = -1.50$).

Exposure × bundle strength: task composition

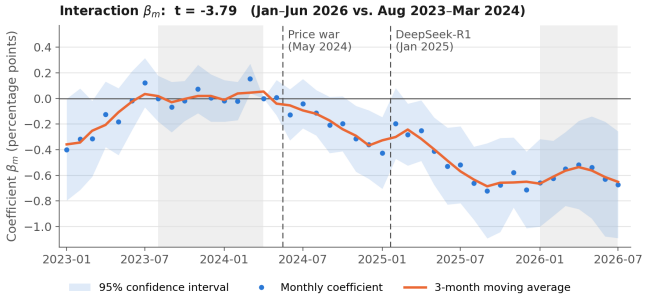
Share of GenAI-exposed tasks in postings



- Among exposed occupations, postings in weak bundles lost AI-exposed task content relative to postings in strong bundles: 0.29 pp per SD of bundle strength ($\beta_m = +0.29$, $t = 3.21$); not distinguishable from the extrapolated pre-trend ($t = 0.35$).

Exposure × bundle strength: residual human role

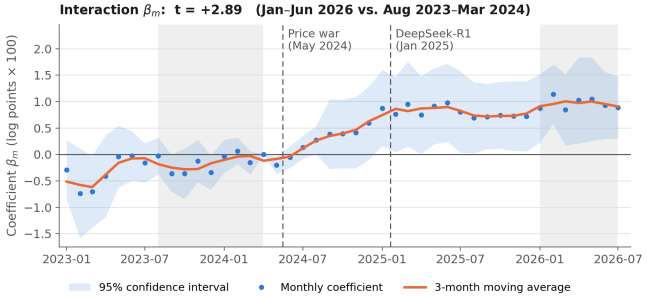
Share of postings with no GenAI-exposed task



- Among exposed occupations, the share of postings made up only of tasks AI cannot absorb rose in weak bundles relative to strong bundles: 0.60 pp more per SD of bundle strength ($\beta_m = -0.60$, $t = -3.79$; survives detrending). The line falls because β_m is the strong-minus-weak difference.

Exposure × bundle strength: return of human role

Posted wage within education × experience category



- Holding the education and experience mix of postings fixed, posted wages in exposed weak-bundle occupations fell relative to exposed strong-bundle occupations: 1.1 log points per SD of bundle strength ($\beta_m = +1.14$, $t = 2.89$); not distinguishable from the extrapolated pre-trend ($t = 0.71$).

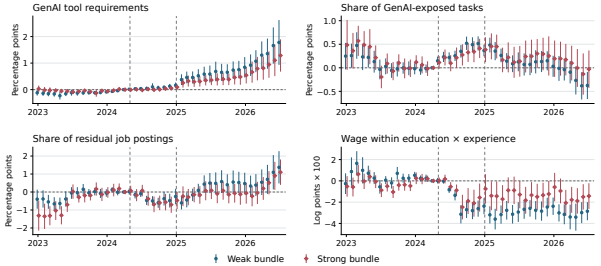
Summary of estimates: interaction β, change 2026H1 vs. pre-war

Outcome	Change	SE	t	Pretrend t	Detrended t
GenAI-tool requirement share (pp)	−0.367	0.125	−2.93	1.18	−3.32
Exposed-task share (pp)	0.294	0.092	3.21	1.43	0.35
Residual-posting share (pp)	−0.599	0.158	−3.79	0.78	−3.25
Advertised wage (log × 100)	1.144	0.396	2.89	1.13	0.71
Bachelor's degree or above (pp)	0.118	0.176	0.67	0.45	0.09
At least three years of experience (pp)	0.232	0.214	1.08	0.01	0.29
New postings (log × 100)	−1.979	1.859	−1.06	0.59	−0.98

Change: mean of the monthly interaction coefficients over Jan–Jun 2026 minus their mean over Aug 2023–Mar 2024 ("late minus pre" in the figure titles). Pretrend compares the second half of the pre window with the first. Detrended t tests the late coefficient against the extrapolated linear pretrend. SE clustered by O*NET code (245–255 clusters; 754–881 occupations).

Two results survive detrending (residual postings, tool requirements); the task-content and wage interactions do not on their own establish a shock-induced change.

Sensitivity: median split (weak vs. strong)



Within-group exposure gradients per 1 SD of exposure; reference April 2024; 95% CIs clustered by O*NET occupation. Dashed lines: May 2024 price war, January 2025 DeepSeek-R1.

Same design with M_o replaced by High_o ($= 1$ above the median of bundle strength): separate exposure gradients β_m^H, β_m^L for strong and weak bundles.

Change 2026H1 vs. pre-war in the exposure gradient, weak vs. strong (t on the difference):

- Tool requirements: 1.29 vs 0.86 ($t = -1.18$)
- Exposed-task share: -0.11 vs 0.05 ($t = 0.75$)
- Residual postings: 0.58 vs 0.26 ($t = -0.63$)
- Wage: -3.31 vs -1.19 ($t = 2.49$)

Posted wages fall with exposure in both groups, by less in strong bundles. The split discards 39% of the index's variation; the continuous interaction is the main specification.

What we do not (yet) find

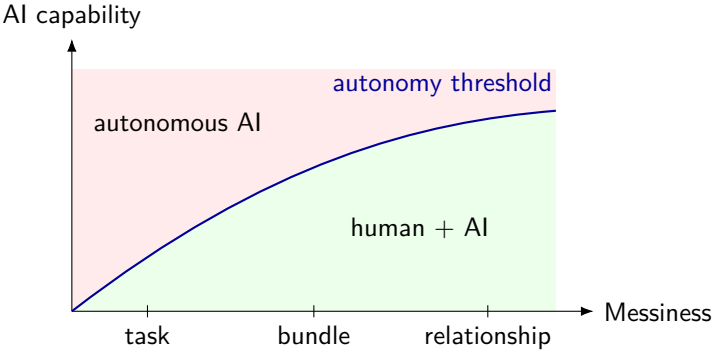
- Entry requirements. The accessibility result predicts that residual postings in exposed weak bundles drop codifiable-skill requirements. Bachelor's-or-above: $+0.12$ pp ($t = 0.67$); three-plus years of experience: $+0.23$ pp ($t = 1.08$). Signs compatible, estimates imprecise.
- Posting volume. The employment test is the price-theoretic payoff: fewer total postings with more residual ones is the inelastic case, more is the Jevons case. Interaction -1.98 log points ($t = -1.06$): inconclusive.
- Within-occupation wage gaps. The experience premium rises more with exposure in strong bundles ($+0.54$, $t = 2.10$; static $t = 3.45$; detrended $t = 2.03$). Dispersion (P90–P10, P75–P25), the education premium, and the mixed-task minus residual gap show no robust interaction.

Plan

- A price theory of task bundling
- Evidence from labor markets
- **Messy jobs as protection of employment**

Autonomy threshold and messy jobs

The **autonomy threshold** is the point at which a task requires human judgment, accountability, or contextual sensitivity that AI cannot reliably provide.



Three tiers of jobs by messiness

Tier 1

Single verifiable tasks
What determines the threshold?

Tier 2

Strong task bundles
What determines whether a bundle is strong?

Tier 3

Bundles of relationships
Who owns the relationship?

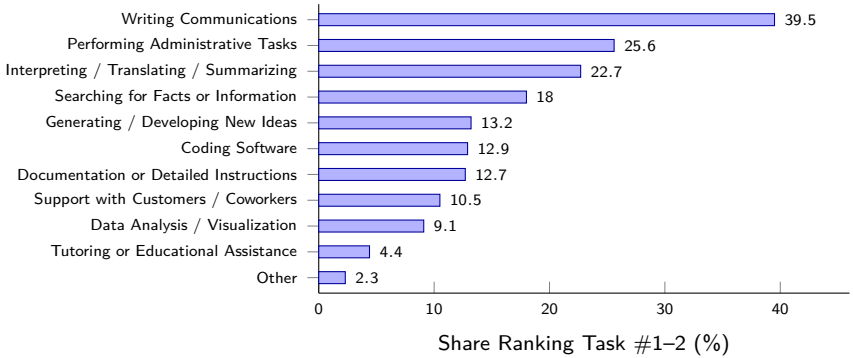
Roughly:

- Tier 1: coder, junior contract lawyer, journalist working off the wire.
- Tier 2: nurse, litigation lawyer, tech sales, journalist doing "shoe-leather" reporting.
- Tier 3: entrepreneur, senior manager, consultant doing AI implementation.

Silicon Valley got it right only at Tier 1

Clean, verifiable tasks

AI can learn quickly and accurately and cross the autonomy threshold.



Three tiers of jobs, three economies

Algorithmic economy

Tier 1

Cognition at near-zero cost.
Who captures value?

The messy middle

Tier 2

Most careers.
Bundled work that survives.

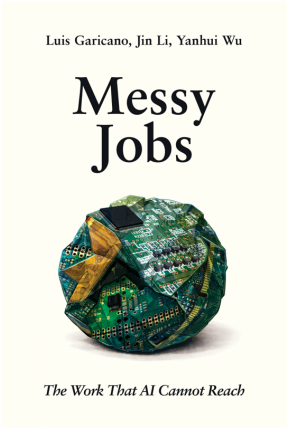
Relational work

Tier 3

Authority, implementation, authenticity.

AI cannot handle messy jobs!

https://messyjobs.ai/

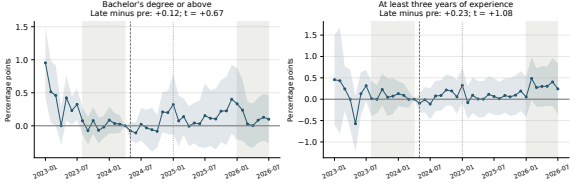


Backup: robustness and caveats

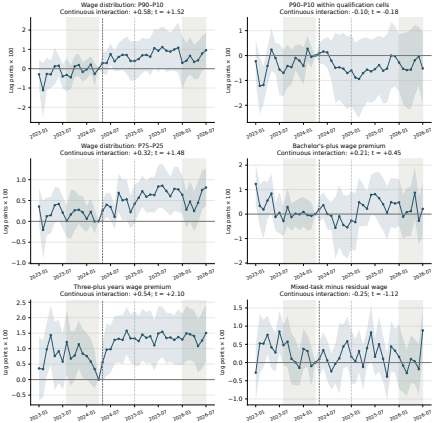
- The shock is national, so the design is a continuous difference-in-differences; parallel trends rest on the sixteen-month pre-period. Pretrend and detrended contrasts are reported for every outcome.
- Bundle strength and exposure correlate at -0.35; ceilings and mean reversion could generate an interaction. Exposure^2 $\times$ month controls address this.
- Postings are firm demand for a job form, not worker choice; in the competitive model the two coincide, so the test is of the organizational margin with worker sorting inferred.
- Posted wages are offers, not realized pay; employment stocks are not observed.
- The five-platform sample was chosen in earlier exploration, not preregistered; coefficient-level intervals do not account for that selection or for multiple outcomes.
- O*NET 31.0 is an August 2026 snapshot; the crosswalk (BGE-M3 similarity, then an LLM pick) is many-to-one and not fully reviewed; inference clusters at the O*NET code.

Backup: requirements and within-occupation wage gaps

Education and experience requirements, β_m



Within-occupation wage gaps, β_m



Backup: taking stock

Proved

The share wedge $\kappa > 1 - \alpha$; the sorting cutoff and the three regimes in (c, m); accessibility of the residual role; fixed-price expansion; displacement under inelastic demand; the attenuation of all posting responses at $c \geq \bar{c}(m)$.

Interpreted

Exposure as α, bundle strength as c; the price war as a rise in Γ; tool requirements as purchased machine services; the wage result as composition and reallocation, not a cut.

Conjectured

That the common wage decline reflects inelastic demand; that an "accelerate" vs "absorb" split of the exposure score would separate μ_H from μ_A; that strong-bundle protection ends at the interior boundary, at a date the data do not yet reach.

Exposure alone predicts nothing about task content; exposure interacted with separability does.