Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries*

Luis Garicano† Jin Li‡ Yanhui Wu§ Leo Yang Yang¶
Zhenyu Zhu‖

November 3, 2026

Abstract

This paper studies how the effect of AI on an occupation depends not just on which tasks AI can perform (AI exposure) but also on how costly it is to unbundle those tasks from the job (bundle strength). We build a two-task model in which AI can either assist one task inside a bundled job or supply that task autonomously while a human supplies the residual task. We show that, in weak-bundle occupations, AI automates some tasks and narrows the boundary of the job. In strong-bundle occupations where tasks are not independently reallocable, AI improves performance inside the job, but does not remove the human from the bundle. We test the theory with more than 200 million Chinese job postings during January 2023–July 2026. We measure bundle strength based on O*NET indices for shared context, liability, and across-task knowledge spillovers. We find that among the higher AI-exposed occupations, lower bundle strength is associated with an increase in the use of AI as a job requirement, a decline in the share of AI-exposed-task content, and a rise in human-only postings. We also find that while greater AI exposure is associated with a wage decline, strong-bundle occupations receive a wage premium within the same education and experience category.

*We used GenAI tools for assistance with the proofs, coding, references, building the figures and tables, English drafting and checking the syntax and grammar. All remaining errors are our own.

†London School of Economics.
‡University of Hong Kong.
§University of Hong Kong.
¶Hong Kong Baptist University.
‖University of Hong Kong.

1 Introduction

A growing empirical literature measures the labor market impact of artificial intelligence by looking at the exposure of individual tasks to GenAI (e.g., Eloundou et al., 2024). The idea is that occupations with more tasks exposed to AI receive a larger shock and hence are likely to suffer lower employment and lower earnings. In this paper we argue that jobs are not tasks, but bundles of tasks, and that the strength of the bundle determines the impact of the shock on wages and employment.

Consider the job of a radiologist. A radiologist not only classifies scans. She triages cases, communicates with other physicians, trains residents, takes responsibility for difficult calls, and signs a diagnosis that allows others to take decisions on treatment. The market pays for a bundle, not for a discrete task, in the same way that when a driver buys a car she pays for the bundle that is the car, not just for the sum of the speed, acceleration, safety and other characteristics.

In our theory, the survival of the bundle is the result of the tradeoff between the benefit of specialization in a single task and the cost of coordination between different tasks. As Becker and Murphy (1992) argued, a generalist who performs several tasks inside one job avoids coordination losses at the cost of doing some components less efficiently. When a bundle is strong, breaking the bundle is very costly. As a result, AI assists the human, but the human still sells the full bundle and retains a larger share of revenue. When the bundle is weak, so that the cost of splitting it is small, AI replaces some tasks, the human role becomes narrower, and the labor share of output decreases. The question then is not only to what extent AI can perform a given task, but whether splitting that task from the job and performing it separately destroys enough value that it is more efficient to keep the job intact.

We model jobs as a bundle of two tasks in a competitive equilibrium. Task 1 is prone to AI replacement while Task 2 is intrinsically human. For simplicity, we call Task 1 the verifiable task, as LLMs are capable of providing verifiable cognition such as drafting, classifying, calculating, summarizing, and coding a bounded problem. We refer to Task 2 as contextual, as human presence is necessary to handle relational, physical, or political tasks. Frontier AI capability improves machine performance on Task 1 but not Task 2.

Production can stay bundled, with one human supplying both tasks and using machine assistance on Task 1, or become unbundled, with AI supplying Task 1 autonomously and a human supplying only Task 2. In bundled production, a worker allocates time between Task 1 and Task 2, while in unbundled production, a worker devotes all available time to Task 2. Jobs differ in their strength of bundling, defined by the coordination cost of

unbundling tasks. Workers are heterogeneous in two dimensions of skills, corresponding to the two task types.

The model produces four main results. First, the labor share is larger when the job remains bundled than when it is unbundled, since the human contributes to both tasks in the bundle but only to Task 2 in the unbundled case. Second, as AI advances, the share of jobs that are unbundled increases. Third, unbundling produces a discrete jump in labor supply, since workers go from allocating their work to both tasks to working only on one, freeing a lot of time. If demand is inelastic, this increase in supply reduces price and displaces the least productive workers. This is a form of the task-substitution mechanism, originating with Autor et al. (2003) and extended and studied by Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018, 2019). Fourth, strong bundles protect workers from the AI shock. It is still the case that AI raises performance on Task 1, but since the task remains inside the bundle, the worker keeps the larger revenue share, and the capacity shock is mitigated.

We take the theory to data using a large dataset of Chinese job postings. It contains more than 200 million ads posted on five recruitment platforms between January 2023 and July 2026. The ads contain a list of responsibilities, hiring requirements, and a wage. Inside our window there were two large shocks to the cost of autonomous AI: the price war that began in May 2024, when the main Chinese labs cut the price of API access by an order of magnitude; and the release of DeepSeek-R1 in January 2025. Both led, in terms of our model, to a rise in the relative advantage of autonomous AI.

To measure exposure we use the tasks listed in the job posting: we score each of the 44,479 tasks on the extent to which a language model can substantially accelerate or absorb it. From it we can construct the exposure of the occupation, as the mention-weighted average of those scores in the last quarter of 2023, before the shock.

To measure the strength of the bundle we rely on O*NET. For each occupation we build indices of the three sources of coordination cost that the model points to, namely shared context, liability, and across-task knowledge spillovers, and then take their first principal component. We show that the two measures are related but distinct. Among the white-collar occupations we study, their correlation is -0.35.

Our empirical design relies on comparing occupations by exposure, by bundle strength, and by their interaction, before and after May 2024, within platform-by-month and platform-occupation cells. The interaction is the object the theory is about. According to the theory, if we compare two occupations with high exposure, we expect the one in which the bundle is less strong to experience a larger rise in postings that require the use of generative AI tools, a decline in the share of exposed-task content, and a rise in postings that contain

only residual, human tasks. We find that posted wages decline with exposure in both weak and strong bundles, but by less in strong ones: holding the education and experience mix of postings fixed, a one standard deviation increase in bundle strength increases the exposure gradient of posted wages by about one log point. We find that education and experience requirements show no differential movement. We find that two of these results, on residual postings and on tool requirements, survive the extrapolation of pre-period trends; but the task-content and wage results do not. We therefore interpret the estimates as evidence on the restructuring of advertised jobs, not as causal effects of AI on employment or earnings.

We believe that several recent empirical findings fit our hypothesis better than the "exposure leads to job loss" hypothesis. Early evidence suggests that generative AI changes the composition of work more than the level of employment. Humlum and Vestergaard (2025) combine GenAI adoption surveys and Danish administrative records and find occupational switching and task restructuring but precise null effects on earnings and hours two years after ChatGPT. They also find the content of the work changed: users spend the time saved by GenAI on other tasks inside their jobs. Gathmann et al. (2024) combine patent-based AI exposure measures with German survey and administrative data and also find that AI shifts the task content of jobs with small displacement effects, and Hampole et al. (2025) that exposed tasks lose labor demand by less when they are a small part of the job, because workers move time to the rest of it. This is what one would expect if AI is useful on one task inside a job but not for the whole bundle: in some occupations the bundle survives and AI removes a weak link; in others the job reorganizes around a narrower human role. Losses appear where the bundle is weakest: on freelance platforms, where a job is a single task, postings for writing and coding fell by about a fifth within months of ChatGPT and freelancers lost work and income (Demirci et al., 2025; Hui et al., 2024); in U.S. payroll data, firms hire fewer young workers in the most exposed occupations while older workers in the same occupations are unaffected (Brynjolfsson et al., 2025; Crane and Soto, 2026). A junior job is formed by the list of separate tasks as in the task-based model (in a consulting firm, doing the PowerPoint, producing the Excel spreadsheet), while the senior job is a bundle of the simpler tasks with tasks such as judgment and responsibility. These tasks require understanding the underlying work and hence do not allow cognition to be split off. We find that, consistent with the theory, the two cases are different: when the bundle is weak, the job narrows and the exposed task disappears from the advertised job, as residual-only postings increase, and the posted wage on the narrower job decreases, while where the bundle is strong, tasks stay together and wages hold up better.

The gains from AI are also asymmetric: weaker and less-experienced workers see larger short-run productivity gains in writing, customer service, and coding (Brynjolfsson et al.,

2023; Noy and Zhang, 2023; Peng et al., 2023). In the model, that comes from relaxing the Task 1 bottleneck: workers whose weak dimension was verifiable cognition benefit most when Task 1 is assisted or replaced.

Our paper is closest to a recent literature that takes the job bundle, rather than the task, as the relevant unit of analysis. Autor and Thompson (2025) treat jobs as bundles and use that structure to forecast how automation reshapes expertise requirements and hence wages. Freund and Mann (2025) study job transformation within a surviving bundle: when automation removes one task, the remaining human tasks become more valuable. We focus on a novel, prior question: the composition of the bundle itself is endogenous. AI can either improve performance inside the existing job or cause the existing job to break apart, depending on coordination costs. They explain heterogeneous wage effects across workers inside transformed occupations while we explain heterogeneous task-content, wage, and employment effects across occupations with similar AI exposure but different bundle strength. Exposure is an incomplete approach not just because workers differ in skill portfolios, but also because jobs differ in how *splittable* they are. Hosseini Maasoum and Lichtinger (2026) likewise move beyond exposure measures by showing that AI changes occupational entry barriers and effective labor supply by automating some tasks and making others easier. In our model, entry depends on the strength of the bundle: if the job is unbundled, entry is easier as workers no longer need to be strong in both skills.

More broadly, our paper connects to the task-based automation literature cited above and to the empirical literature that measures AI exposure at the occupation or task level, which already finds that generative AI is more likely to automate some tasks within occupations than occupations outright (Gmyrek et al., 2023, 2025). While this literature treats tasks as a list of separable units, our unit of analysis is the job bundle. In the weak-bundle case, our model nests the standard displacement mechanism of the task-based model. Bundling is the force that can shut that mechanism down. Our empirical design also nests the exposure approach: the exposure gradient matters where bundle strength does not matter, and its interaction with bundle strength tests whether it does matter.

Our treatment of bundle strength draws on the organizational-design literature on specialization and coordination. Rosen (1978) first studied how technology and the distribution of worker skills determine job bundles and assignment, very relatedly to his analysis of hedonic prices in markets for differentiated goods. Rosen (1983) studies how specialization increases the utilization of knowledge. In this latter context, Becker and Murphy (1992) formalized the specialization-coordination tradeoff, and Dessein and Santos (2006) studied the trade-off between ex ante coordination and ex post adaptation. Garicano (2000) shows how communication allows an increase in the utilization of knowledge via the formation of

hierarchies and Garicano and Rossi-Hansberg (2006) introduce assignment of heterogeneous workers to hierarchical layers. Ide and Talamàs (2025) build on this framework to analyze how AI replacing workers at different hierarchical levels (as a problem solver or as a worker) and with different ability levels affects the wage distribution.

The paper proceeds as follows. Section 2 introduces the model and obtains the allocation of workers to jobs. Section 3 characterizes the competitive equilibrium with fixed product prices. Section 4 introduces the demand side. Section 5 takes the model to the Chinese posting data. Section 6 discusses interpretation and concludes. Appendix A contains proofs; Appendices B and C report sample sizes, additional estimates, and a rank-split specification.

2 Model and organizational sorting

We study the sorting of workers by skill level between bundled and unbundled production as a function of coordination costs and the level of artificial intelligence.

2.1 Technology and time allocation

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

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

where q_i is delivered quality on task i.

Task 1 is codifiable cognition. Task 2 is the part of the work that remains intrinsically human. A worker has type $s = (s_1, s_2)$ drawn from

$$S = [\underline{s}_1, \overline{s}_1] \times [\underline{s}_2, \overline{s}_2] \subset \mathbb{R}^2_+$$

with continuous density $g(s_1, s_2)$ that is strictly positive on S.

Each worker is endowed with one unit of working time. If both tasks are bundled in the job, the worker must allocate time across the two tasks:

$$t_1 + t_2 \leq 1, \qquad t_1 \geq 0, \qquad t_2 \geq 0. \tag{2}$$

The level of AI technology is m for machine. We let

$$\mu_H(m) > 0, \qquad \mu_A(m) > 0, \tag{3}$$

denote the effective machine services available on task 1 in bundled and unbundled pro-

duction. Throughout, μ_H and μ_A are continuously differentiable and strictly increasing in m.

There are two organizational forms.

Production with bundled tasks. One human supplies both tasks inside one job and combines time with machine assistance on task 1. Gross bundled output is

$$y_H^g(t_1, t_2, s, m) = \left[\mu_H(m)^{1-\eta}(t_1 s_1)^{\eta}\right]^{\alpha}(t_2 s_2)^{1-\alpha}, \qquad \eta \in (0,1). \tag{4}$$

The parameter η is the human share inside task 1 when the two tasks remain bundled. The task-1 input is itself a Cobb–Douglas composite of machine services (weight $1-\eta$) and human time-weighted skill (weight η).

Production when tasks are unbundled. Autonomous AI supplies task 1 and a human supplies only task 2. Splitting the two tasks across agents destroys a fraction $c \in [0,1)$ of value. Because the machine handles task 1, the human devotes the full time endowment to task 2. Gross unbundled output is then:

$$y_U(s_2, m; c) = (1-c)\mu_A(m)^{\alpha} s_2^{1-\alpha}. \tag{5}$$

The coordination loss c summarizes the cost of the loss of shared context, accountability, and knowledge spillovers between tasks, and measurement problems when the tasks are split.

Lemma 1 (Optimal bundled time allocation)**.** *Fix (s, m). A bundled worker uses the full time endowment and allocates it as*

$$t_1^* = \frac{\alpha\eta}{\kappa}, \qquad t_2^* = \frac{1-\alpha}{\kappa}, \tag{6}$$

where

$$\kappa \equiv \alpha\eta + 1 - \alpha = 1 - \alpha(1-\eta). \tag{7}$$

Substituting (6) *into* (4) *yields maximized bundled output*

$$y_H(s, m) = \Omega_H \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta} s_2^{1-\alpha}, \tag{8}$$

where

$$\Omega_H \equiv \left(\frac{\alpha\eta}{\kappa}\right)^{\alpha\eta}\left(\frac{1-\alpha}{\kappa}\right)^{1-\alpha}. \tag{9}$$

Proof. See Appendix A. □

The time split follows from the Cobb–Douglas structure. Because the exponents on the two tasks in (4) are $\alpha\eta$ (for the human component of task 1) and $1-\alpha$ (for task 2), the worker allocates time in proportion to those exponents. The constant $\kappa = \alpha\eta + 1 - \alpha$ is just the normalizing sum, and $\Omega_H \in (0,1)$ captures the efficiency loss from splitting one unit of time across two tasks. In unbundled production that bottleneck disappears because the human no longer spends time on task 1.

2.2 Competitive downstream pricing and human earnings

Let $P > 0$ denote the product price. Under competitive pricing of the final-good technology (1), task revenues satisfy

$$p_1 q_1 = \alpha P y, \qquad p_2 q_2 = (1-\alpha) P y. \tag{10}$$

This follows from the Cobb–Douglas form: each task receives its output-elasticity share of revenue.

In bundled production the human supplies all of task 2 and the human component of the task-1 composite. Since the task-1 composite is itself Cobb–Douglas with human share η, competitive pricing inside that composite assigns share η of task-1 revenue to the human and share $1-\eta$ to machine services. In unbundled production the human supplies only task 2.

Lemma 2 (Human revenue shares)**.** *Fix (P, m, c). Net human earnings are*

$$w_H(P, s, m) = \kappa\, P\, y_H(s, m), \tag{11}$$

$$w_U(P, s_2, m; c) = (1-\alpha)\, P\, y_U(s_2, m; c). \tag{12}$$

Hence the worker keeps a larger fraction of the gross revenue in the bundled case than the unbundled one.

Proof. A bundled human earns the full task-2 revenue share $(1-\alpha)Py$ plus the human part of the task-1 revenue, $\eta \cdot \alpha P y = \alpha\eta P y$. The sum is $[\alpha\eta + 1 - \alpha]Py = \kappa P y$. Since y_H already incorporates the optimal time allocation, we have $w_H = \kappa P y_H$. An unbundled human earns only the task-2 share $(1-\alpha)Py_U$. Since $\kappa = 1 - \alpha(1-\eta) > 1-\alpha$ whenever $\eta > 0$, the bundled human keeps a strictly larger fraction of gross revenue. □

Note that the reduced-form objects in the baseline are the mappings $m \mapsto \mu_H(m)$ and $m \mapsto \mu_A(m)$. The share comparison in Lemma 2 is not imposed. It is derived from competitive pricing of the downstream task structure conditional on those effective machine services.

2.3 Sorting between bundled and unbundled production

Define the effective-machine frontier

$$\Gamma(m) \equiv \frac{\mu_A(m)^{\alpha}}{\mu_H(m)^{\alpha(1-\eta)}}. \tag{13}$$

Γ measures the relative advantage of autonomous machine cognition over the machine-assisted bundle.

Assumption 1 (Relative autonomous effectiveness rises at the relevant margin)**.** *For the range of m under study,*

$$\frac{d}{dm}\Gamma(m) > 0. \tag{14}$$

This says that better frontier AI helps autonomous production more than it helps machine-assisted bundled production. It holds whenever μ_A grows sufficiently fast relative to μ_H.

One interpretation is that frontier systems increasingly complete bounded, verifiable subtasks with little human time, while within-job assistance remains limited by prompting, monitoring, verification, or institutional sign-off. The coordination wedge c captures a different margin: the value destroyed when the two tasks are separated across agents.

Proposition 1 (Sorting cutoff)**.** *Fix (P, m, c). There exists a cutoff*

$$s_1^*(m, c) = \left[\frac{(1-\alpha)(1-c)}{\kappa\Omega_H}\Gamma(m)\right]^{1/(\alpha\eta)} \tag{15}$$

such that

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

Hence, conditional on entering the focal occupation, workers with low task-1 skill weakly prefer unbundled production, while workers with high task-1 skill weakly prefer bundled production. Moreover, $s_1^*(m, c)$ is increasing in m and decreasing in c.

Proof. See Appendix A. □

The intuition is straightforward. A worker with high s_1 has a comparative advantage in bundled production, where task-1 skill earns a return. A worker with low s_1 gains little from the bundle and prefers to let AI handle task 1 autonomously, freeing all time for task 2. Better AI (higher m) raises Γ and therefore s_1^*, pushing more workers into unbundling. Higher coordination cost c lowers s_1^*, keeping more workers bundled.

Note that the low-c benchmark nests task substitution. When c is low enough, task 1 is separated from the human role. The human then loses both the task itself and the revenue attached to it: the human claim on gross revenue narrows from κ to $1-\alpha$. That is a form of the task-substitution logic. What bundling adds is an organizational force that can block this reallocation when c is high.

The cutoff also yields an occupation-level distinction between weak and strong bundles.

Assumption 2 (Pure unbundling is feasible)**.** *For the range of m under study,*

$$\kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}\overline{s}_1^{\alpha\eta} \leq (1-\alpha)\mu_A(m)^{\alpha}. \tag{16}$$

This ensures that autonomous AI is productive enough that some workers prefer unbundling even at the highest task-1 skill in the population.

Corollary 1 (Occupation-level cutoffs)**.** *For each m, define*

$$\underline{c}(m) \equiv 1 - \frac{\kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}\overline{s}_1^{\alpha\eta}}{(1-\alpha)\mu_A(m)^{\alpha}}, \tag{17}$$

$$\bar{c}(m) \equiv 1 - \frac{\kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}\underline{s}_1^{\alpha\eta}}{(1-\alpha)\mu_A(m)^{\alpha}}. \tag{18}$$

Then $0 \leq \underline{c}(m) \leq \bar{c}(m) < 1$, and:

(i) if $c \leq \underline{c}(m)$, every worker who enters the focal occupation weakly prefers unbundled production;

(ii) if $c \geq \bar{c}(m)$, every worker who enters the focal occupation weakly prefers bundled production;

(iii) if $c \in (\underline{c}(m), \bar{c}(m))$, both organizational forms are privately optimal for some worker types.

Figure 1 plots the regime boundaries in (c, m) space. The dashed curve is $\underline{c}(m)$ and the solid curve is $\bar{c}(m)$. Below the dashed curve every entering worker prefers unbundling; above the solid curve every entering worker prefers bundling; between the two curves the occupation splits internally. Better machine capability shifts both cutoffs upward. As a result, pure unbundling expands, pure bundling shrinks, and the interior region moves upward but narrows in c-space.

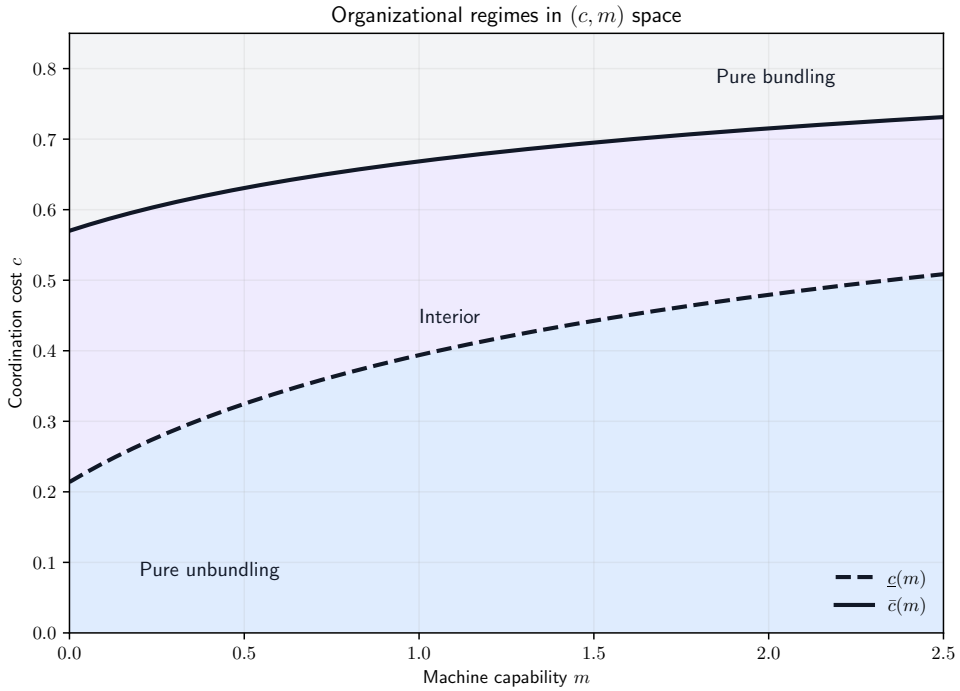


Figure 1: Organizational regimes in (c, m) space. The dashed curve is $\underline{c}(m)$ and the solid curve is $\bar{c}(m)$. Below $\underline{c}(m)$ the occupation is in pure unbundling; between the two curves it is in interior equilibrium; above $\bar{c}(m)$ it is in pure bundling. Benchmark: $S = [0.2, 1]^2$, g uniform, $\alpha = 0.50$, $\eta = 0.75$, hence $\kappa = 0.875$ and $\Omega_H \approx 0.55$, $\mu_H(m) = 1 + m$, and $\mu_A(m) = 1.5 + 1.5m$. The organizational cutoffs do not depend on the demand closure.

3 Participation and competitive equilibrium with a fixed product price

We now introduce the outside option and the fixed-price participation margin. The logic is recursive: Section 2 determined which inside form a worker prefers; this section asks whether that preferred inside option beats outside work. That distinction matters because unbundling can make the residual human role easier to enter even before prices adjust.

3.1 Participation and occupational choice

Workers can also stay outside the focal occupation and obtain reservation utility $\bar{U} > 0$. In the fixed-price benchmark, the relevant relative price is

$$\tau \equiv \frac{\bar{U}}{P}. \tag{19}$$

A higher τ means that outside work is more attractive relative to the focal occupation.

Given $(P, \bar{U}, m, c)$, worker welfare is

$$V(s; P, \bar{U}, m, c) = \max\{\bar{U},\, w_H(P, s, m),\, w_U(P, s_2, m; c)\}. \tag{20}$$

A worker enters bundled production if bundled earnings weakly exceed both outside work and unbundled work. A worker enters unbundled production if unbundled earnings weakly exceed both outside work and bundled work.

With P and $\bar{U}$ taken as given, equilibrium is the allocation implied by occupational choice, participation, and competitive pricing.

Definition 1 (Competitive equilibrium with fixed product price)**.** Fix $(P, \bar{U}, m, c)$ and let $\tau = \bar{U}/P$. A competitive equilibrium with fixed product price is a measurable partition

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

of the type space S, where O is the set of workers who choose the outside option, $\mathcal{U}$ is the set of workers who choose unbundled production, and H is the set of workers who choose bundled production, such that each worker type is assigned to an option that attains the maximum in (20). Equivalently, for every $s = (s_1, s_2) \in S$,

(i) if $s \in O(\tau, m, c)$, then

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

(ii) if $s \in \mathcal{U}(\tau, m, c)$, then

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

(iii) if $s \in H(\tau, m, c)$, then

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

On indifference sets, any measurable tie-breaking rule is admissible.

Given Proposition 1, this equilibrium allocation can be characterized by two participation thresholds. A worker who stays bundled enters if and only if bundled earnings weakly exceed the outside option, which yields

$$s_2 \geq b_H(s_1; \tau, m) \equiv \left(\frac{\tau}{\kappa \Omega_H \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta}}\right)^{1/(1-\alpha)}. \tag{21}$$

A worker who prefers unbundling enters if and only if unbundled earnings weakly exceed the outside option, which yields

$$s_2 \geq b_U(\tau, m, c) \equiv \left(\frac{\tau}{(1-\alpha)(1-c)\mu_A(m)^\alpha} \right)^{1/(1-\alpha)}. \tag{22}$$

Using the convention that knife-edge types with $s_1 = s_1^*(m, c)$ are assigned to bundled production, the active workers in each organizational form are

$$\mathcal{U}(\tau, m, c) = \Big\{ (s_1, s_2) \in S : \; s_1 < s_1^*(m, c), \; s_2 \geq b_U(\tau, m, c) \Big\}, \tag{23}$$

$$H(\tau, m, c) = \Big\{ (s_1, s_2) \in S : \; s_1 \geq s_1^*(m, c), \; s_2 \geq b_H(s_1; \tau, m) \Big\}. \tag{24}$$

Workers not in $\mathcal{U}(\tau, m, c) \cup H(\tau, m, c)$ choose the outside option.

These sets determine employment and gross output in each organizational form:

$$N_U(\tau, m, c) \equiv \int_{\mathcal{U}(\tau,m,c)} g(s)\, ds, \qquad Y_U(\tau, m, c) \equiv \int_{\mathcal{U}(\tau,m,c)} y_U(s_2, m; c) g(s)\, ds, \tag{25}$$

$$N_H(\tau, m, c) \equiv \int_{H(\tau,m,c)} g(s)\, ds, \qquad Y_H(\tau, m, c) \equiv \int_{H(\tau,m,c)} y_H(s, m) g(s)\, ds. \tag{26}$$

Lemma 3 (Existence and uniqueness under fixed P)**.** *For every $(P, \bar{U}, m, c)$, a competitive equilibrium exists and is unique up to zero-measure tie sets.*

Because the product price is given, there is no goods-market-clearing equation to solve. Once $(P, \bar{U}, m, c)$ are fixed, the thresholds b_H and b_U pin participation and the sorting cutoff s_1^* pins organizational choice.

3.2 Accessibility of the unbundled role

Proposition 2 (Accessibility of the unbundled role)**.** *For every (τ, m, c), if $s_1 \leq s_1^*(m, c)$, then*

$$b_U(\tau, m, c) \leq b_H(s_1; \tau, m),$$

with strict inequality for every $s_1 < s_1^*(m, c)$. Therefore, conditional on preferring unbundling, the participation condition for the residual human role is weakly easier than the participation condition for bundled production.

Once task 1 is peeled away, the market no longer asks whether the worker can supply the whole bundle. It asks only whether the worker can supply what remains. The task-1 requirement drops out, so the role opens to worker types who could not profitably enter the bundled occupation.

3.3 Fixed-price expansion of the residual role

Proposition 3 (Fixed-price expansion of the residual role)**.** *Suppose the fixed-price equilibrium is interior, so that*

$$s_1^*(m,c) \in (\underline{s}_1, \overline{s}_1), \qquad b_U(\tau, m, c) \in (\underline{s}_2, \overline{s}_2).$$

Then

$$\frac{dN_U(\tau,m,c)}{dm} = \underbrace{\int_{b_U}^{\overline{s}_2} g(s_1^*, s_2) \frac{ds_1^*(m,c)}{dm}\, ds_2}_{\text{sorting effect: } (+)} - \underbrace{\int_{\underline{s}_1}^{s_1^*} g(s_1, b_U) \frac{db_U(\tau,m,c)}{dm}\, ds_1}_{\text{threshold effect: } (-)}. \tag{27}$$

Moreover,

$$\frac{d \ln b_U(\tau,m,c)}{dm} = -\frac{\alpha}{1-\alpha} \frac{d \ln \mu_A(m)}{dm} < 0. \tag{28}$$

Hence better machine capability expands unbundled employment in the fixed-price benchmark.

The first term of (27) captures workers shifted from bundled to unbundled production as s_1^* rises with m. The second captures new entrants whose task-2 skill was previously below the participation threshold. Both terms are positive at fixed prices, so unbundled employment expands unambiguously in the partial-equilibrium benchmark. This isolates the organizational and participation margins. It does *not* yet include the labor-saving capacity shock generated by inelastic demand.

4 Endogenous prices, the capacity shock, and displacement

The fixed-price model isolates the organizational margin. But it suppresses the labor-saving mechanism that matters most in weak-bundle occupations. At fixed prices, better AI expands the residual human role. To make clear when that expansion turns into displacement, the product price must respond when each surviving worker can suddenly process more machine output.

4.1 Demand and endogenous-price equilibrium

Suppose the final good faces inverse demand

$$P(Y) = Y^{-1/\varepsilon}, \qquad \varepsilon > 0. \tag{29}$$

Then the relative value of the outside option is no longer a parameter but an equilibrium object:

$$\tau \equiv \frac{\bar{U}}{P} = \bar{U} Y^{1/\varepsilon}. \tag{30}$$

Given τ, the active sets are still (23) and (24), so aggregate output is

$$Y(\tau, m, c) = Y_H(\tau, m, c) + Y_U(\tau, m, c).$$

An endogenous-price equilibrium is therefore a fixed point of

$$\tau = \bar{U}\, Y(\tau, m, c)^{1/\varepsilon}. \tag{31}$$

We denote endogenous-price equilibrium objects by a superscript $*$.

Proposition 4 (Endogenous-price equilibrium)**.** *For every $(\bar{U}, m, c)$, an endogenous-price competitive equilibrium exists and is unique.*

The organizational frontier is unchanged by demand closure. Lemma 2, Proposition 1, and Corollary 1 compare private payoffs inside an occupation at a common product price, so they survive any goods-market closure. What changes is participation.

Figure 2 shows occupational choice in skill space for three representative points, one from each regime. In pure unbundling (left panel), all active workers are in the residual role; the horizontal threshold b_U separates participants from outsiders, and task-1 skill is irrelevant. In pure bundling (right panel), both skills matter: the curved threshold $b_H(s_1)$ falls with s_1, so workers stronger in codifiable cognition can enter with lower task-2 skill. The interior panel combines both structures within one occupation.

4.2 The time-split capacity shock and displacement

The time-allocation microfoundation reveals why unbundling can displace labor even when output per remaining worker rises.

In bundled production the worker is a bottleneck to herself: by Lemma 1, the worker devotes only

$$t_2^* = \frac{1-\alpha}{\kappa}$$

of the workday to task 2. In unbundled production AI takes over task 1 and the worker devotes all time to task 2. The worker's task-2 capacity therefore rises by the factor

$$\frac{1}{t_2^*} = \frac{\kappa}{1-\alpha} > 1. \tag{32}$$

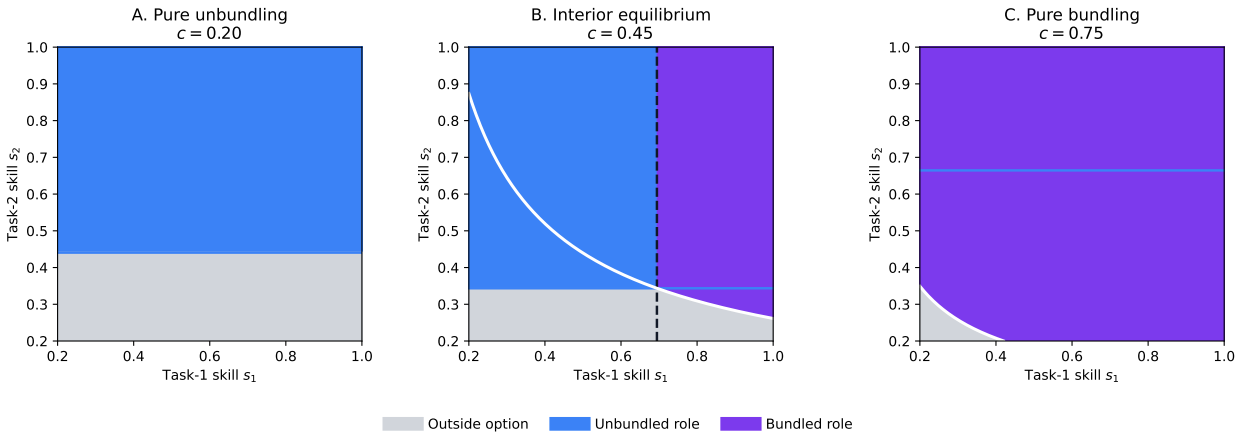


Figure 2: Occupational choice in skill space at $m = 0.8$ under the endogenous-price benchmark used in Figure 3. The three panels show representative coordination costs from pure unbundling, interior equilibrium, and pure bundling. Gray types stay outside the occupation, blue types supply the residual human role in unbundled production, and purple types stay in bundled production. In the interior panel, the dashed vertical line is the sorting cutoff $s_1^*(m, c)$, the blue horizontal line is the unbundled participation threshold $b_U(\tau^*, m, c)$, and the white curve is the bundled participation threshold $b_H(s_1; \tau^*, m)$. Benchmark: $S = [0.2, 1]^2$, g uniform, $\alpha = 0.50$, $\eta = 0.75$, $\mu_H(m) = 1 + m$, $\mu_A(m) = 1.5 + 1.5m$, inverse demand $P(Y) = Y^{-1/\varepsilon}$ with $\varepsilon = 0.8$, and $\bar{U} = 0.60$.

This is a worker-level capacity shock. Each surviving unbundled worker produces more task-2 output than a bundled worker did, because the time formerly spent on task 1 is now free. If demand absorbs the extra output, wages hold; if not, price falls and the marginal worker is displaced.

Under endogenous prices, the unbundled participation threshold is

$$b_U^*(m, c) = \left(\frac{\bar{U}}{(1-\alpha) P^*(m, c)(1-c)\mu_A(m)^\alpha} \right)^{1/(1-\alpha)}. \tag{33}$$

Using (29),

$$\frac{d \ln b_U^*(m, c)}{dm} = \frac{1}{1-\alpha} \left(\underbrace{\frac{1}{\varepsilon} \frac{d \ln Y^*(m, c)}{dm}}_{\text{price compression}} - \underbrace{\alpha \frac{d \ln \mu_A(m)}{dm}}_{\text{direct productivity gain}} \right). \tag{34}$$

The first term raises the threshold (more output depresses the price, making the occupation less attractive). The second lowers it (better AI raises the value of the residual role). The sign depends on which force dominates.

Proposition 5 (Displacement in the weak-bundle benchmark)**.** *Suppose the endogenous-*

price equilibrium is in pure unbundling and $b_U^*(m,c) \in (\underline{s}_2, \overline{s}_2)$. *Then*

$$\frac{dN_U^*(m,c)}{dm} = -\left(\int_{\underline{s}_1}^{\overline{s}_1} g(s_1, b_U^*(m,c))\, ds_1\right) \frac{db_U^*(m,c)}{dm}. \tag{35}$$

Therefore human employment falls if and only if

$$\frac{1}{\varepsilon}\frac{d\ln Y^*(m,c)}{dm} > \alpha \frac{d\ln \mu_A(m)}{dm}. \tag{36}$$

When (36) *holds, price compression dominates the direct machine-productivity effect, the participation threshold rises, and the bottom tail of workers is displaced.*

Proposition 5 is the clearest statement of the mechanism. In weak-bundle occupations the human loses task 1 and the labor share narrows from κ to $1-\alpha$. At the same time, each surviving worker can devote all time to task 2. If demand is inelastic ($\varepsilon < 1$), that per-worker capacity shock lowers the price enough that fewer workers remain profitable. This is the standard displacement channel nested inside the model.

The interior regime adds a second force. Better machine capability shifts the sorting cutoff toward unbundling, which tends to expand the residual role, but price compression can work in the opposite direction.

Proposition 6 (Unbundled employment in interior equilibrium)**.** *Suppose the endogenous-price equilibrium is interior and satisfies*

$$s_1^*(m,c) \in (\underline{s}_1, \overline{s}_1), \qquad b_U^*(m,c) \in (\underline{s}_2, \overline{s}_2).$$

Then

$$\frac{dN_U^*(m,c)}{dm} = \underbrace{\int_{b_U^*}^{\overline{s}_2} g(s_1^*, s_2)\frac{ds_1^*(m,c)}{dm}\, ds_2}_{\textit{sorting effect}} - \underbrace{\int_{\underline{s}_1}^{s_1^*} g(s_1, b_U^*)\frac{db_U^*(m,c)}{dm}\, ds_1}_{\textit{threshold effect}}. \tag{37}$$

The first term is a sorting effect: higher m *shifts workers from bundled to unbundled production. The second is a threshold effect. In the fixed-price benchmark the threshold effect is positive because* $db_U/dm < 0$. *Under inelastic demand it becomes negative whenever price compression makes* $db_U^*/dm > 0$.

Figure 3 shows the two aggregate objects that matter most under inelastic demand. Panel A plots population labor income,

$$W^{pop}(m,c) \equiv P^*(m,c)\Big[\kappa Y_H^*(m,c) + (1-\alpha)Y_U^*(m,c)\Big], \tag{38}$$

which equals labor income per person because the worker population has unit mass. Panel B plots total human employment. In the weak-bundle region, moving right lowers employment: each remaining worker can process more task-2 work, price falls, and the bottom tail is displaced. In the strong-bundle region, employment is much more stable because the occupation stays bundled.

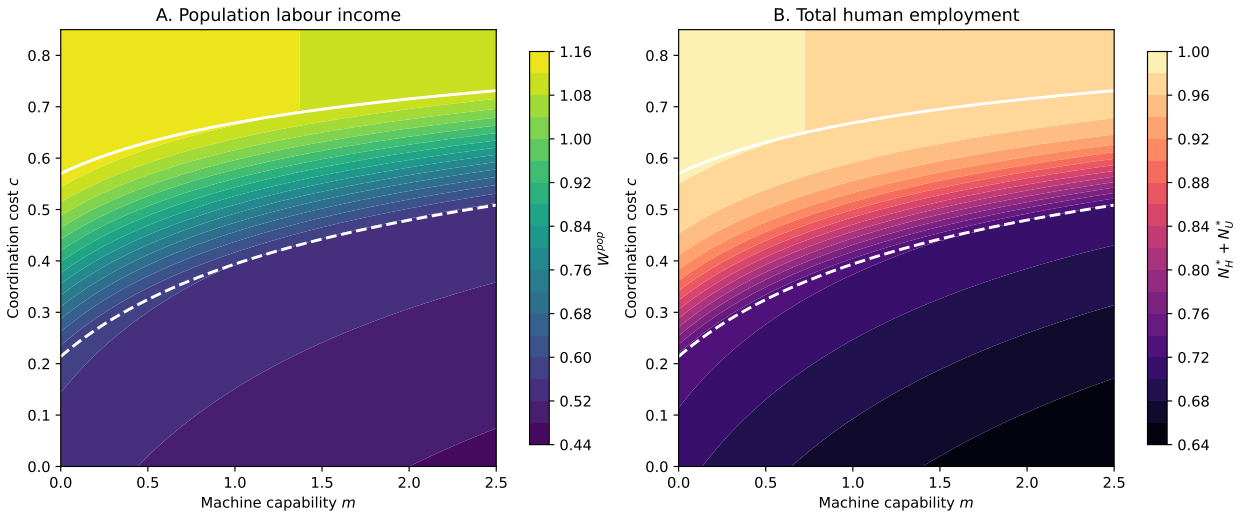


Figure 3: Population labor income and human employment in (c, m) space under inelastic demand. Panel A plots population labor income $W^{pop} = P^*[\kappa Y_H^* + (1-\alpha)Y_U^*]$. Panel B plots total human employment $N_H^* + N_U^*$. The white dashed and solid curves are the organizational cutoffs $\underline{c}(m)$ and $\bar{c}(m)$. In the weak-bundle region, higher m reduces employment because the time-release capacity shock lowers the product price enough to raise the participation threshold. In the strong-bundle region, the occupation remains bundled and employment is much more stable. Benchmark: $S = [0.2, 1]^2$, g uniform, $\alpha = 0.50$, $\eta = 0.75$, $\mu_H(m) = 1 + m$, $\mu_A(m) = 1.5 + 1.5m$, inverse demand $P(Y) = Y^{-1/\varepsilon}$ with $\varepsilon = 0.8$, and $\bar{U} = 0.60$.

4.3 Fixed-type earnings and real losers

The endogenous-price model also changes the interpretation of worker-level earnings schedules. With fixed prices, a decline in average unbundled earnings could come entirely from composition. With inelastic demand, even a fixed worker type can lose because the product price falls.

Corollary 2 (Fixed-type earnings under endogenous prices)**.** *For any bundled type* $s = (s_1, s_2)$,

$$w_H^*(s, m, c) = \kappa P^*(m, c)\Omega_H \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta} s_2^{1-\alpha}, \tag{39}$$

so

$$\frac{d \ln w_H^*(s,m,c)}{dm} = \underbrace{\frac{d \ln P^*(m,c)}{dm}}_{\text{price effect } (<0)} + \underbrace{\alpha(1-\eta)\frac{d \ln \mu_H(m)}{dm}}_{\text{productivity effect } (>0)}. \tag{40}$$

For any unbundled type s_2*,*

$$w_U^*(s_2,m,c) = (1-\alpha)P^*(m,c)(1-c)\mu_A(m)^{\alpha}s_2^{1-\alpha}, \tag{41}$$

so

$$\frac{d \ln w_U^*(s_2,m,c)}{dm} = \underbrace{\frac{d \ln P^*(m,c)}{dm}}_{\text{price effect } (<0)} + \underbrace{\alpha\frac{d \ln \mu_A(m)}{dm}}_{\text{productivity effect } (>0)}. \tag{42}$$

Hence fixed-type earnings fall whenever price compression dominates the direct productivity effect. The model therefore generates real losers, not only composition effects among survivors.

Figure 4 provides a direct comparison of the two regimes. For a weak-bundle occupation ($c = 0.05$) and a strong-bundle occupation ($c = 0.80$), it traces employment, population labor income, fixed-type earnings, and aggregate output as AI capability m rises.[1] In the weak bundle, Panel D reflects the combination of two forces. Relative to bundled production, unbundling creates a discrete time-release capacity gain because the worker no longer allocates time to Task 1. That output expansion drives down the product price, reducing employment by roughly 7% and population labor income by 12%. In the strong bundle, the occupation stays bundled until m is high enough that some workers choose to unbundle.

[1]The benchmark for Figure 4 uses $\eta = 0.85$ and $\mu_H(m) = 1 + 0.2m$, to sharpen the contrast between the two regimes; the qualitative pattern holds for the baseline parameters used in the previous figures.

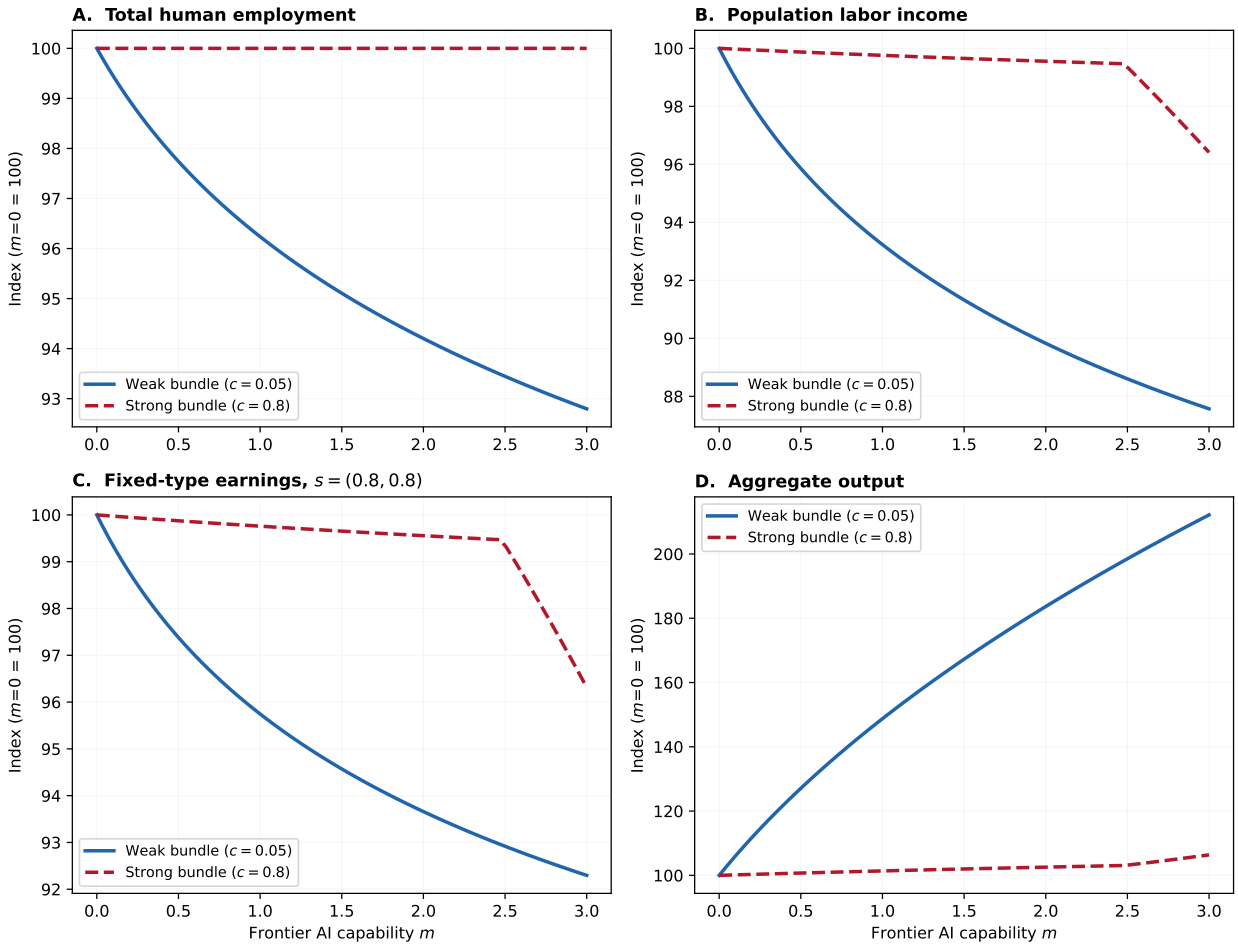


Figure 4: How bundling protects workers. Each panel compares a weak-bundle occupation ($c = 0.05$, solid blue) with a strong-bundle occupation ($c = 0.80$, dashed red) as frontier AI capability m rises. Panel A: total human employment. Panel B: population labor income W^{pop}. Panel C: chosen earnings for a fixed worker type $s = (0.8, 0.8)$. Panel D: aggregate output. All indices are normalized to 100 at $m = 0$. In the weak bundle, the time-release capacity gain associated with unbundling and continued improvements in autonomous AI drive output up by over 100%, depressing the product price enough to reduce employment by 7% and population labor income by 12%. In the strong bundle, the occupation stays bundled until $m \approx 2.68$, after which earnings and labor income drop precipitously. This kink indicates that the occupation has entered the interior regime. Until that point, output rises only modestly, employment is essentially unchanged, and labor income declines by less than 4%. Benchmark: $S = [0.2, 1]^2$, g uniform, $\alpha = 0.50$, $\eta = 0.85$, $\mu_H(m) = 1 + 0.2m$, $\mu_A(m) = 1.5 + 2m$, inverse demand $P(Y) = Y^{-1/\varepsilon}$ with $\varepsilon = 0.85$, and $\bar{U} = 0.40$.

5 Empirical evidence from Chinese job postings

The model's predictions concern two characteristics of an occupation and their interaction. AI exposure corresponds to the weight of codifiable work in the job, the tasks a machine can accelerate or absorb. Bundle strength corresponds to the coordination cost c. Proposition 1 and Corollary 1 say that a rise in the relative advantage of autonomous machine cognition moves Task 1 out of the human job in weak bundles and leaves it in place in strong ones; Lemma 2 says that the human claim on revenue narrows from κ to $1-\alpha$ when it does. Job postings are the natural place to look for this margin. They record the tasks, the requirements, and the posted wage of each advertised job, so a change in what firms post reveals an organizational change before it appears in employment stocks. This section constructs the two measures, describes the data and the design, and reports the results. The estimates are associations around a national shock, and we report their sensitivity to pre-period trends throughout.

5.1 Measuring AI exposure

We measure GenAI exposure from the task content of Chinese job advertisements. The unit of classification is the recruitment corpus's custom occupation, rather than an occupation in China's official occupational directory. The task-based approach is related to Eloundou et al. (2024), but our implementation uses a Chinese posting-derived catalog and continuous task scores, not their occupation-level exposure estimates.

The catalog contains 44,479 responsibility-task entries. Each entry has a continuous LLM-exposure score $e_k \in [0,1]$, assessing whether LLMs or agents can substantially accelerate or absorb its core output. These are model-assessed capability scores, not observations of realized adoption or displacement. Let $a_{jk} = 1$ when advertisement j mentions catalog task k, counting each distinct task at most once within an advertisement. We fix occupational exposure using responsibility mentions in October–December 2023:

$$\widetilde{E}_o = \frac{\sum_{j \in \mathcal{J}_{o,2023Q4}} \sum_{k \in \mathcal{K}_j} a_{jk} e_k}{\sum_{j \in \mathcal{J}_{o,2023Q4}} \sum_{k \in \mathcal{K}_j} a_{jk}}, \qquad E_o = \frac{\widetilde{E}_o - \overline{\widetilde{E}}}{\mathrm{sd}(\widetilde{E})}. \tag{43}$$

Here $\mathcal{K}_j$ contains recognized responsibility entries with an observed exposure score. The aggregation weights tasks by their frequency of mention; it is neither an unweighted mean over an official task list nor a mean that first gives every advertisement equal weight. Requirement entries are excluded. Standardization gives equal weight to distinct eligible occupations within each outcome's regression sample.

The frozen exposure measure was constructed from the original baseline advertisement rows before the deduplication revision. All outcome panels below use deduplicated advertisements. Keeping this distinction explicit makes the current results reproducible; duplication can still affect the baseline exposure weights. The catalog and its capability assessments were developed retrospectively, so a baseline posting window does not make the entire measurement procedure prospective.

For the task-content outcomes, an exposed responsibility task ("Task 1") has $e_k \geq 0.6$ and an on-site-presence score below 0.3. A residual posting has at least one recognized responsibility task and no such exposed task. These labels describe advertised content; they do not establish that a particular worker or job actually lost a task.

5.2 Measuring bundle strength

Following the aforementioned theoretical distinction between weak and strong bundles, we summarize three sources of difficulty in separating tasks: shared context, liability, and across-task knowledge spillovers. We use the O*NET 31.0 descriptor files (National Center for O*NET Development, 2026), construct the three indices at the O*NET-SOC level, map them to the custom Chinese occupations used in the regressions, and extract their first principal component. This is a proxy for bundle strength, not a direct estimate of the model's proportional output loss. Release 31.0 is an August 2026 snapshot, so these descriptors should not be interpreted as a directly observed pre-shock measure.

5.2.1 Shared context

Shared context combines relational embeddedness (RE), the extent to which relevant information resides in interpersonal interaction, and task discretion (TD), the latitude to determine tasks and decisions. Table 1 lists the constituent Work Context items. For each item, $z_s(\cdot)$ denotes standardization across O*NET occupations with observed values. We standardize the item means again:

$$\mathrm{RE}_s = z_s(\mathrm{mean}_{k\in\mathcal{R}_s}\, z_s(X_{sk}))\,, \quad \mathrm{TD}_s = z_s(\mathrm{mean}_{k\in\mathcal{D}_s}\, z_s(X_{sk}))\,, \quad \mathrm{SC}_s = z_s\left(\frac{\mathrm{RE}_s + \mathrm{TD}_s}{2}\right). \tag{44}$$

We use the CX scale and treat values marked "Recommend Suppress" as missing. RE requires at least four of its five items and TD at least two of its three items. Higher SC represents greater reliance on shared interpersonal context and discretion.

Table 1: Work Context items used to construct shared context and liability

Index	Element ID	O*NET element
RE	4.C.1.a.4	Contact With Others
	4.C.1.a.2.l	Face-to-Face Discussions with Individuals and Within Teams
	4.C.1.b.1.e	Work With or Contribute to a Work Group or Team
	4.C.1.b.1.f	Deal With External Customers or the Public in General
	4.C.2.a.3	Physical Proximity
TD	4.C.3.b.8	Determine Tasks, Priorities and Goals
	4.C.3.a.4	Freedom to Make Decisions
	4.C.3.a.2.b	Frequency of Decision Making
LI	4.C.3.a.1	Consequence of Error
	4.C.3.a.2.a	Impact of Decisions on Co-workers or Company Results
	4.C.1.c.1	Health and Safety of Other Workers
	4.C.1.c.2	Work Outcomes and Results of Other Workers

5.2.2 Liability

Liability captures responsibility for the consequences of work, including harm to others and consequences for coworkers or the organization. We average the standardized values of the four LI items in Table 1, requiring at least three observed items, and standardize the resulting mean:

$$\mathrm{LI}_s = z_s(\mathrm{mean}_{k\in\mathcal{L}_s}\, z_s(X_{sk}))\,. \tag{45}$$

Higher values indicate greater responsibility and more consequential errors. The same suppression rule applies as for shared context.

5.2.3 Across-task spillovers

Across-task spillovers measure the extent to which different activities use overlapping knowledge. We use the importance ratings of 41 Work Activities and 33 Knowledge domains for 910 O*NET occupations. Unlike the Work Context indices, this calculation retains flagged ratings. First, for each knowledge domain j, we estimate a nonnegative least-squares projection on activities, centering every column across occupations:

$$K_{sj} - \bar{K}_j = \sum_{k=1}^{41}(A_{sk} - \bar{A}_k)b_{kj} + u_{sj}, \qquad b_{kj} \geq 0. \tag{46}$$

The vector b_k is a common knowledge profile for activity k. Second, we measure the similarity between activity profiles using the occupation's own knowledge-importance ratings as

weights:

$$\varsigma_{sk\ell} = \frac{\sum_j K_{sj} b_{kj} b_{\ell j}}{\sqrt{\sum_j K_{sj} b_{kj}^2}\sqrt{\sum_j K_{sj} b_{\ell j}^2}}. \tag{47}$$

Finally, we average these similarities over distinct activity pairs, weighting by the product of the two activity-importance ratings:

$$\mathrm{SP}_s = z_s \left(\frac{\sum_{k \neq \ell} A_{sk} A_{s\ell} \varsigma_{sk\ell}}{\sum_{k \neq \ell} A_{sk} A_{s\ell}} \right). \tag{48}$$

Pairs with an undefined cosine similarity are omitted from numerator and denominator. The calculation uses all activity pairs and does not use GenAI exposure labels. Higher SP indicates more overlap in the knowledge used across activities, a potential source of losses when tasks are separated.

5.2.4 Crosswalk

The current empirical estimates use a many-to-one crosswalk from custom recruitment occupations to O*NET-SOC. On the Chinese side, we summarize each occupation using its five most common posting titles and ten most frequent responsibility-task labels in 2023Q4. On the O*NET side, we use the occupation title, up to eight reported job titles, the description, and up to ten core task statements. A multilingual sentence encoder (BGE-M3) produces title and task similarities; their equally weighted mean selects eight candidate O*NET occupations for each Chinese occupation.

DeepSeek-chat, at temperature zero, then selects the single candidate with the closest core tasks or rejects all candidates. Of 1,265 candidate Chinese occupations, 1,214 receive a match and 51 are unmatched. The matched occupations map to 395 distinct O*NET codes. Each matched Chinese occupation inherits its selected code's SC, LI, and SP values. No equal division over multiple matches is used in these estimates, and we do not describe the frozen mapping as fully manually reviewed. This crosswalk differs from the separately developed crosswalk for the official Chinese occupational directory. Keeping the mapping fixed here preserves consistency with the reported regressions; mapping quality remains a source of measurement uncertainty. Because occupations can share an O*NET code, inference clusters at that code rather than treating every custom occupation as an independent measure of bundle strength.

5.2.5 Principal component analysis

We perform PCA on the correlation matrix of SC, LI, and SP across the 1,214 matched Chinese occupations, with equal weight per custom occupation. We orient PC1 so that all three loadings are positive and standardize its scores. Figure 5 shows eigenvalues of 1.752, 0.860, and 0.388; PC1 explains 58.4% of total variance. Table 2 reports correlations between standardized inputs and component scores, rather than the eigenvector coefficients used to form the scores. For PC1 these loadings are 0.884 for shared context, 0.774 for liability, and 0.609 for spillovers. The corresponding eigenvector coefficients are 0.668, 0.585, and 0.460.

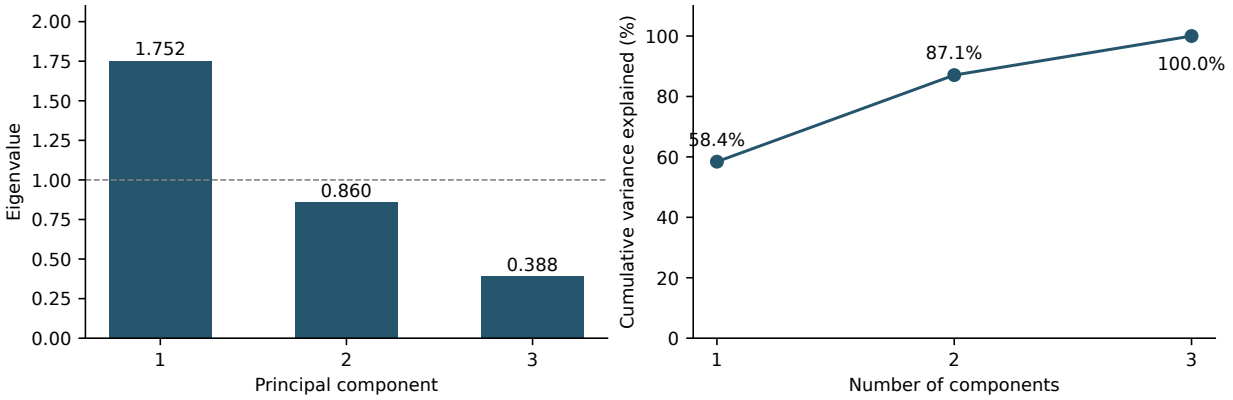


Figure 5: PCA on the three O*NET indices mapped to Chinese custom occupations. Each of the 1,214 occupations receives equal weight.

Table 2: Principal component loadings

	PC1	PC2	PC3
Shared context	0.884	-0.078	-0.461
Liability	0.774	-0.515	0.368
Across-task spillovers	0.609	0.767	0.201
Eigenvalue	1.752	0.860	0.388
Variance explained (%)	58.4	28.7	12.9

Notes: Loadings equal eigenvector coefficients multiplied by the square root of the component's eigenvalue. Signs of PC2 and PC3 have no substantive meaning.

Figure 6 separates the measurement universe from the white-collar universe used for estimation. PC1 and baseline GenAI exposure are negatively correlated: $r = -0.11$ among all 1,214 mapped occupations and $r = -0.35$ among the 952 white-collar occupations. They therefore capture distinct, but not statistically independent, occupational characteristics. Figure 7 illustrates the highest and lowest PC1 values among distinct mapped white-collar

O*NET occupations; English O*NET titles label the inherited scores. The empirical panels impose further outcome-specific sample restrictions.

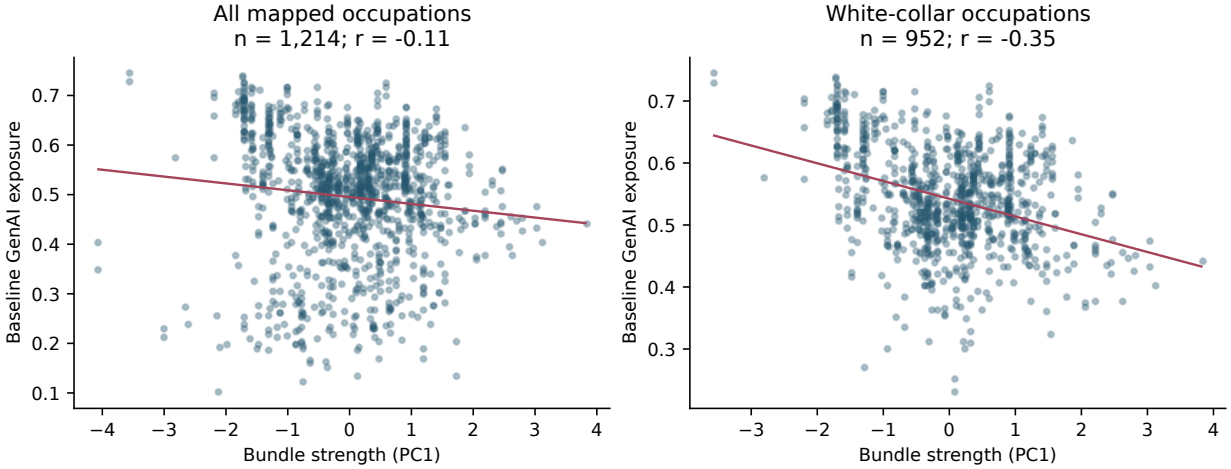


Figure 6: Bundle strength and baseline GenAI exposure. Points are custom Chinese occupations; overlapping points can share an O*NET code. Lines are unweighted linear fits.

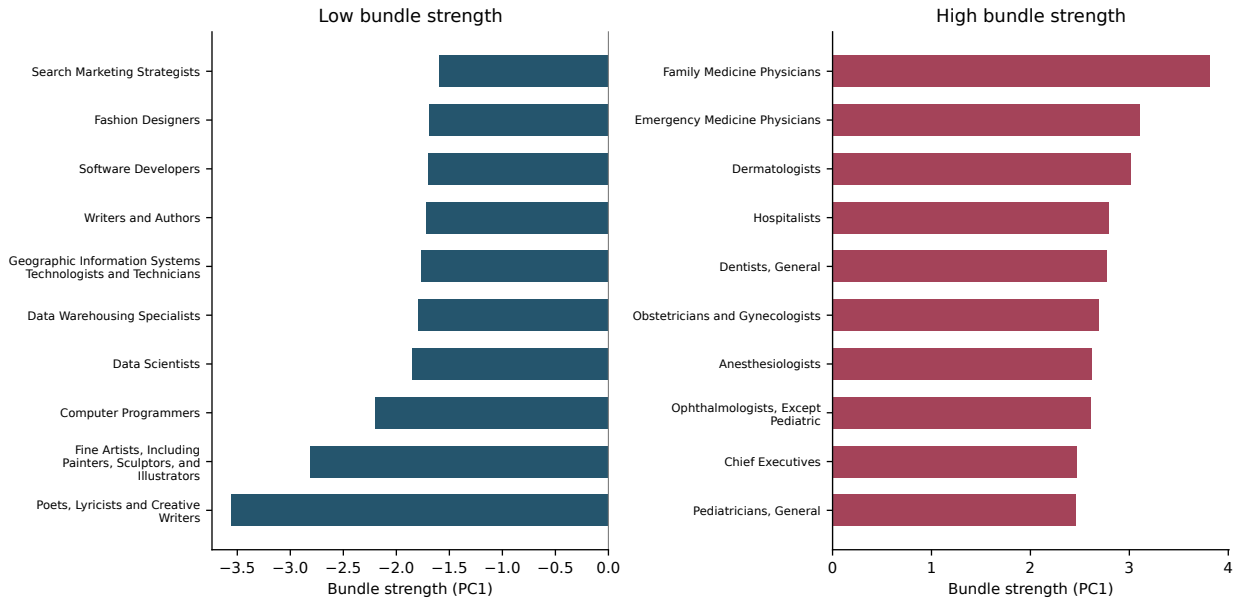


Figure 7: Occupations with low and high measured bundle strength in the white-collar universe. Ten lowest and ten highest distinct mapped O*NET occupations are shown. Scores come from PCA on Chinese custom occupations; labels do not imply a separate regression on US postings.

5.3 Data and outcomes

We use Chinese recruitment advertisements from 51job, 58.com, BOSS Zhipin, Liepin, and Zhilian. The estimation window is January 2023–July 2026, with 2022 used as a lookback period to identify newly appearing advertisements. The deduplicated corpus contains approximately 234 million platform–fingerprint–month records including the lookback year, of which approximately 210 million fall in 2023–2026 before the analysis-specific filters.

Deduplication retains one record per platform, content fingerprint, and month. The fingerprint combines occupation, location, industry, employer attributes, salary bounds, qualifications, and matched task and requirement lists. It has no employer identifier, so distinct employers with identical observed attributes can be merged; duplicates across platforms are not removed. Recruitment volume counts only first appearances of a fingerprint on its platform, while composition and wage outcomes use all monthly deduplicated records.

We restrict occupations to O*NET major groups 11, 13, 15, 17, 19, 21, 23, 25, 27, 29, 41, and 43. To limit changes in platform coverage, we drop platform-months with posting volume below 20% of that platform's median monthly volume during April 2023–April 2024. The main composition outcomes hold province-by-two-digit-industry weights fixed at their October 2023–April 2024 values. Baseline cells require at least five postings, and a monthly observation must retain at least half of its baseline weight and meet the 100-posting threshold.

The exposed-task share uses recognized responsibility-mention weights; other composition shares use posting weights. These controls address observed composition, not unobserved employer turnover.

The four central outcomes are: the share of postings mentioning GenAI-tool use; the fraction of recognized responsibility mentions classified as exposed tasks; the share of recognized-task postings containing no exposed task; and mean log advertised monthly wages. GenAI-tool mentions are a requirement or task-content proxy, not a direct measure of actual technology use. Wages are the midpoint of the reported salary bounds, restricted to RMB 1,000–200,000. The mean-wage outcome fixes education-by-experience composition at baseline, with at least five baseline salaries and three current salaries per cell, at least 50% weight coverage, and at least 100 current salaries. It does not additionally fix province-by-industry weights.

We also report new-posting volume, the share requiring at least a bachelor's degree, and the share requiring at least three years of experience. Qualification shares are calculated among postings with the relevant field recorded; they require at least three recorded values per contributing composition cell and a platform-month field-reporting rate of at least 40%. Missing requirements are not recoded as low requirements in these shares. Appendix Table 6 reports outcome-specific estimation samples.

5.4 Econometric specifications

Let $u = (p, o)$ index platform–occupation units, α_u unit effects, and λ_{pt} platform-by-month effects. Both exposure E_o and continuous bundle strength M_o are standardized over distinct eligible occupations for each outcome, assigning each occupation equal weight in this standardization. Neither measure is dichotomized in the main specifications.

The static specification is

$$y_{ut} = \alpha_u + \lambda_{pt} + \text{Post}_t(\beta E_o M_o + \gamma E_o + \delta M_o) + \varepsilon_{ut}, \tag{49}$$

where $\text{Post}_t = 1$ from May 2024 onward, the cutoff used for the model-price-war analysis. The monthly event study replaces these post interactions by month-specific coefficients:

$$y_{ut} = \alpha_u + \lambda_{pt} + \sum_{m \neq 2024\text{-}04} \mathbf{1}[t = m](\beta_m E_o M_o + \gamma_m E_o + \delta_m M_o) + \varepsilon_{ut}. \tag{50}$$

April 2024 is the reference month. Regressions are weighted by the logarithm of one plus the platform–occupation's baseline mean posting count; standard errors are clustered by mapped O*NET code. The coefficient β_m measures how the change in the exposure gradient

varies with a one-standard-deviation increase in bundle strength. It is not an unconditional AI treatment effect or a comparison of two pre-defined types of occupations.

This specification imposes a linear interaction in the measured index. At any fixed value M, the model-implied exposure gradient is $\gamma_m + \beta_m M$; evaluating that expression at different index values does not require a new binary regression. A median-split specification instead changes both the interaction and the bundle-by-month controls. We report that distinct sensitivity exercise in Appendix C, without treating its groups as theoretical organizational regimes.

For a concise long-run comparison, we report the mean coefficient in January–June 2026 minus its mean in August 2023–March 2024. This event-study contrast differs from the static coefficient, which pools all observed months before and after the May 2024 cutoff. We check pretrends using December 2023–March 2024 minus August–November 2023, and report a sensitivity check that subtracts the linear trend estimated over August 2023–March 2024 and extrapolated to the late window. The covariance of the estimated trend is included in the contrast's uncertainty.

5.5 Main results

Table 3 reports the continuous static interactions. Exposed-task shares and fixed-qualification wages have positive interactions, while residual-posting shares and GenAI-tool requirements have negative interactions; each is statistically distinguishable from zero at the conventional 5% level. The monthly estimates in Figure 8 show when these gradients emerge. Their late-minus-pre contrasts, rather than the static averages over all post-cutoff months, are the principal summary of the event studies.

Table 3: Static post-May-2024 estimates

Outcome	Continuous DDD	SE	t
GenAI-tool requirement share	-0.169	0.061	-2.80
Exposed-task share	0.171	0.066	2.62
Residual-posting share	-0.369	0.112	-3.30
Advertised wage (log)	0.873	0.346	2.52
Bachelor's degree or above	-0.020	0.162	-0.12
At least three years of experience	0.006	0.152	0.04
New postings (log)	-1.496	1.218	-1.23

Notes: Estimates are β in equation (49). E and M are continuous and standardized. SE denotes the O*NET-clustered standard error. Share coefficients are percentage points; log outcomes are multiplied by 100. Each row uses its outcome-specific panel.

Between the pre and late windows, the exposure-by-bundle interaction is +0.294 per-

centage points for exposed-task content ($t = 3.21$) and -0.599 points for residual-posting shares ($t = -3.79$; Table 4). Thus, among more exposed occupations, lower measured bundle strength is associated with a relative reduction in exposed-task content and a relative increase in postings containing only residual tasks. These are changes in advertised job composition, not observations of individual workers losing tasks or employers reorganizing a particular job.

The corresponding interaction is -0.367 percentage points for GenAI-tool requirements ($t = -2.93$), indicating a larger increase in the exposure gradient at lower bundle strength. This pattern concerns advertised tool requirements; it does not establish which organizational mode accounts for actual AI use. For fixed-qualification wages, the interaction is $+1.144$ log points times 100 ($t = 2.89$), consistent with relatively lower wage growth in more exposed occupations with lower bundle strength. New-posting volume does not show the predicted relative protection at higher strength: its interaction is -1.979 log points times 100 ($t = -1.06$).

Trend sensitivity differs across these outcomes. After subtracting an extrapolated linear pretrend, the residual-posting and GenAI-tool interactions remain negative ($t = -3.25$ and -3.32), whereas the exposed-task and mean-wage interactions are no longer distinguishable from zero ($t = 0.35$ and 0.71). The original continuous associations are therefore reproducible, but their robustness is uneven; the wage and exposed-task results do not on their own establish a shock-induced change.

Table 4: Continuous DDD: late-window contrasts and trend diagnostics

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

Notes: Late minus pre compares January–June 2026 with August 2023–March 2024. Units match Table 3; t is the window contrast divided by its clustered SE. Pretrend compares the second half of the pre window with the first. Detrended t tests the late coefficient against the extrapolated linear pretrend, including trend-estimation uncertainty. These checks are sensitivities, not proof of parallel trends.

Education and experience requirements. The proposed relative decline in hiring thresholds in exposed occupations with lower bundle strength is not established in this sample. The continuous late-minus-pre estimates are $+0.118$ percentage points for requiring

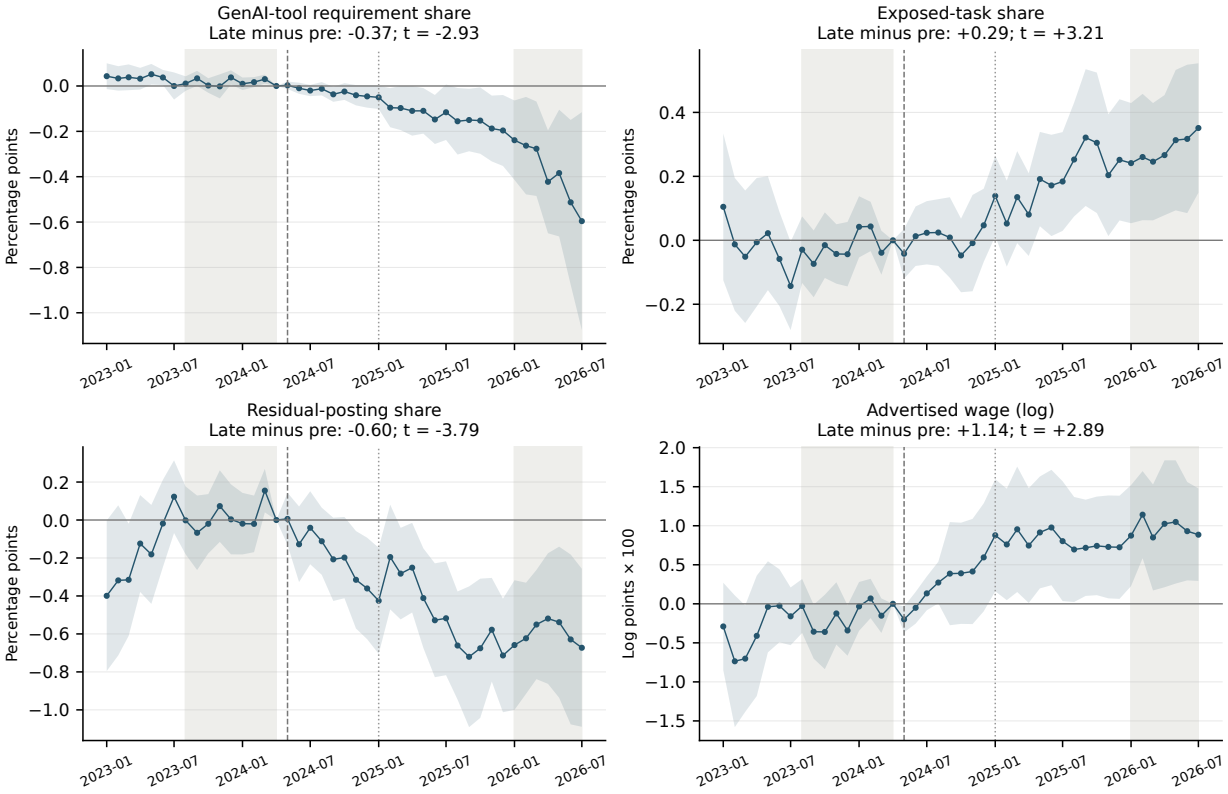


Figure 8: Continuous exposure-by-bundle event studies, China. Points show β_m from equation (50); bands are pointwise 95% confidence intervals. Shading marks the pre and late comparison windows. The dashed line marks May 2024 and the dotted line January 2025. April 2024 is normalized to zero.

a bachelor's degree or above ($t = 0.67$) and +0.232 points for requiring at least three years of experience ($t = 1.08$). The signs are compatible with the proposed ordering, but both estimates are imprecise. Static and detrended estimates also do not provide clear evidence of the predicted differential. Figure 9 reports these null results alongside the main outcomes.

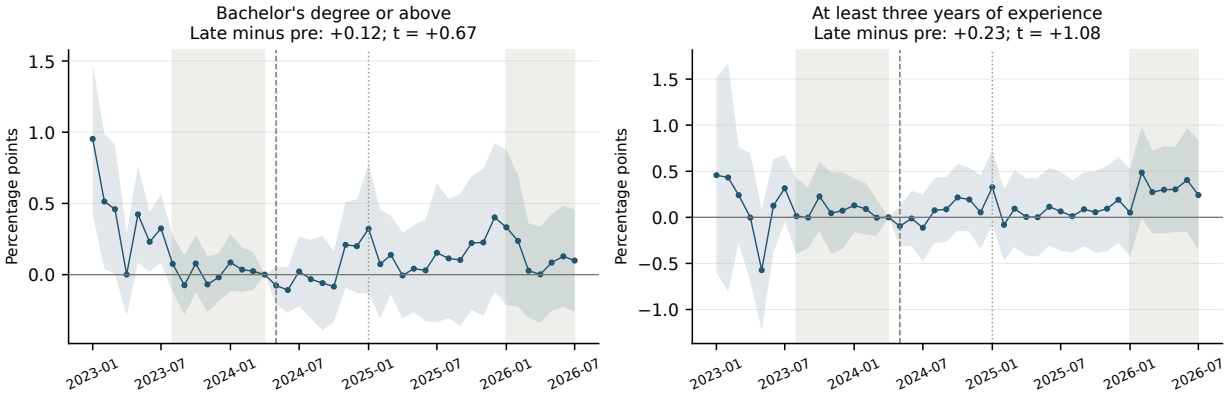


Figure 9: Education and experience requirements: continuous DDD event studies, China. Outcome shares condition on a recorded requirement. Plot conventions follow Figure 8.

Interpretation and specification uncertainty. The five-platform specification was selected during earlier exploration, not preregistered. The deduplicated search considered 16,416 combinations, including 576 five-platform combinations; the report used the highest-scoring five-platform choice. The earlier search-adjusted permutation exercise gave $p = 0.136$ for the five-platform search and $p = 0.106$ for the full search. Conventional coefficient-level intervals do not account for that selection or for testing multiple outcomes. We therefore treat these estimates as exploratory evidence of occupational restructuring, with stronger support for changes in advertised task mix than for a general causal claim about employment, qualification thresholds, or earnings.

5.6 Within-occupation wage gaps

We examine three interpretations of within-occupation wage gaps using the same deduplicated salary records and exposure-by-bundle design: dispersion of advertised wages, qualification premia, and the wage difference between mixed-task and residual postings. These are differences between advertised positions within a platform and occupation. They are not within-firm wage gaps, differences in realized worker earnings, or returns to a worker acquiring a qualification.

Definitions. The distributional outcomes are $\log(P90) - \log(P10)$ and $\log(P75) - \log(P25)$ of advertised salary midpoints, requiring at least 100 valid salaries per platform–occupation–month. To assess the role of qualification composition, we also calculate the baseline-weighted mean of P90–P10 spreads *within* education-by-experience cells. This sensitivity requires at least 30 baseline and 30 current salaries per cell, at least 100 salaries overall, and at least 50% baseline-weight coverage. It averages within-cell spreads; it is not a quantile of a reweighted marginal wage distribution.

The education premium is mean log pay in postings requiring at least a bachelor's degree minus mean log pay in postings with other recorded education requirements, within experience categories. Explicitly unrestricted requirements belong to the latter group; missing education is excluded. The experience premium is the difference between postings requiring at least three years and those requiring less than three years, within education categories; missing numeric experience is excluded. For task bundles, we compare postings containing *both* exposed and other recognized responsibility tasks with residual postings containing recognized tasks but no exposed task, within education-by-experience cells. Exposed-only and unrecognized-only postings are excluded from this contrast.

For each premium, we retain baseline comparison strata with at least five salaries on each side during October 2023–April 2024 and freeze weights using the two sides' combined baseline counts. A current stratum needs at least three salaries on each side; a monthly comparison requires at least 100 salaries across retained pairs and 50% baseline-weight coverage. Education and experience comparisons also retain the 40% platform-month field-reporting rule. The education categories are graduate degree, bachelor's, college, secondary, and other/missing; experience categories are missing or below one year, one to below three, three to below five, and five or more years. These are the qualification cells used in the mean-wage analysis. Employer composition is not additionally held fixed for these new wage-gap outcomes.

Results. Table 5 and Figure 10 apply the continuous interaction in equation (50) to all six outcomes. The late-minus-pre interaction is +0.581 log points times 100 for raw P90–P10 ($t = 1.52$), +0.320 for P75–P25 ($t = 1.48$), and −0.104 for within-qualification P90–P10 ($t = -0.18$). None is statistically distinguishable from zero at the conventional 5% level. The static raw-dispersion interactions are positive and significant ($t = 2.26$ and 2.06 for P90–P10 and P75–P25), but the within-qualification measure is not. Thus, the evidence for a general dispersion gradient is sensitive to the time contrast and qualification adjustment; it does not establish broad wage compression at lower bundle strength.

The experience premium provides a more consistent signal. Its continuous late-minus-

pre interaction is +0.542 log points times 100 ($t = 2.10$), and its static interaction is +0.554 ($t = 3.45$). The detrended contrast remains positive ($t = 2.03$). The pretrend contrast is $t = -1.63$: it is not rejected at 5%, but this does not establish parallel trends, particularly given the limited pre-period and exploratory design. The positive interaction indicates a relative increase in the exposure gradient of the experience premium as bundle strength rises. It does not by itself identify an absolute decline at a particular low-strength value, nor a change in the return to experience for the same worker.

The education-premium interaction is +0.211 ($t = 0.45$), and the mixed-task-minus-residual wage-gap interaction is −0.246 ($t = -1.12$); neither is significant. Together with the null hiring-requirement results, these estimates distinguish a possible experience-premium association from a general decline in entry requirements or wage gaps. All inference is conventional and unadjusted for multiple outcomes. Appendix Tables 7–8 report the continuous static estimates, sample sizes, and trend diagnostics; Appendix C preserves the separate rank-split sensitivity estimates.

6 Discussion and conclusion

We have built a model of endogenous task bundling. Our main argument proceeds in three steps. First, we show that in competitive equilibrium there is a wedge in human workers' revenue shares between bundled and unbundled jobs. A worker in a bundled job performs both tasks and thus earns a larger share of gross revenue than an unbundled worker who only performs Task 2.

Second, we show that unbundling produces a capacity shock. In bundled production, the worker splits her time across two tasks. When the job unbundles, she stops doing Task 1 and devotes everything to Task 2. Her task-2 output jumps discretely. Each surviving worker now produces more than she did before, not because her skill at Task 2 improves, but because she is released from doing Task 1.

Third, if demand for the service is inelastic, the flood of task-2 output from surviving workers depresses the product price. The price falls enough that the least-productive workers can no longer cover their opportunity cost. They leave the occupation. This nests the standard task-substitution mechanism of Autor, Levy, and Murnane (2003), Acemoglu and Autor (2011), and Acemoglu and Restrepo (2018, 2019), but only in the weak-bundle case. At fixed prices, by contrast, better AI expands the residual human role. Displacement arises only once the endogenous-price closure is imposed. The strong bundle mitigates the capacity shock, which helps protect jobs.

Our model provides a new lens to look at the impact of AI on labor outcomes at both the

Table 5: Within-occupation wage gaps: continuous DDD late-minus-pre contrasts

Wage-gap outcome	Continuous DDD	SE	t
P90–P10	0.581	0.382	1.52
Within-cell P90–P10	-0.104	0.565	-0.18
P75–P25	0.320	0.216	1.48
Education premium	0.211	0.467	0.45
Experience premium	0.542	0.258	2.10
Mixed-task minus residual	-0.246	0.219	-1.12

Notes: Estimates are late-minus-pre contrasts of β_m in equation (50). All coefficients and SEs are log points multiplied by 100. E and M remain continuous and standardized; SEs are clustered by O*NET code. Windows match Table 4.

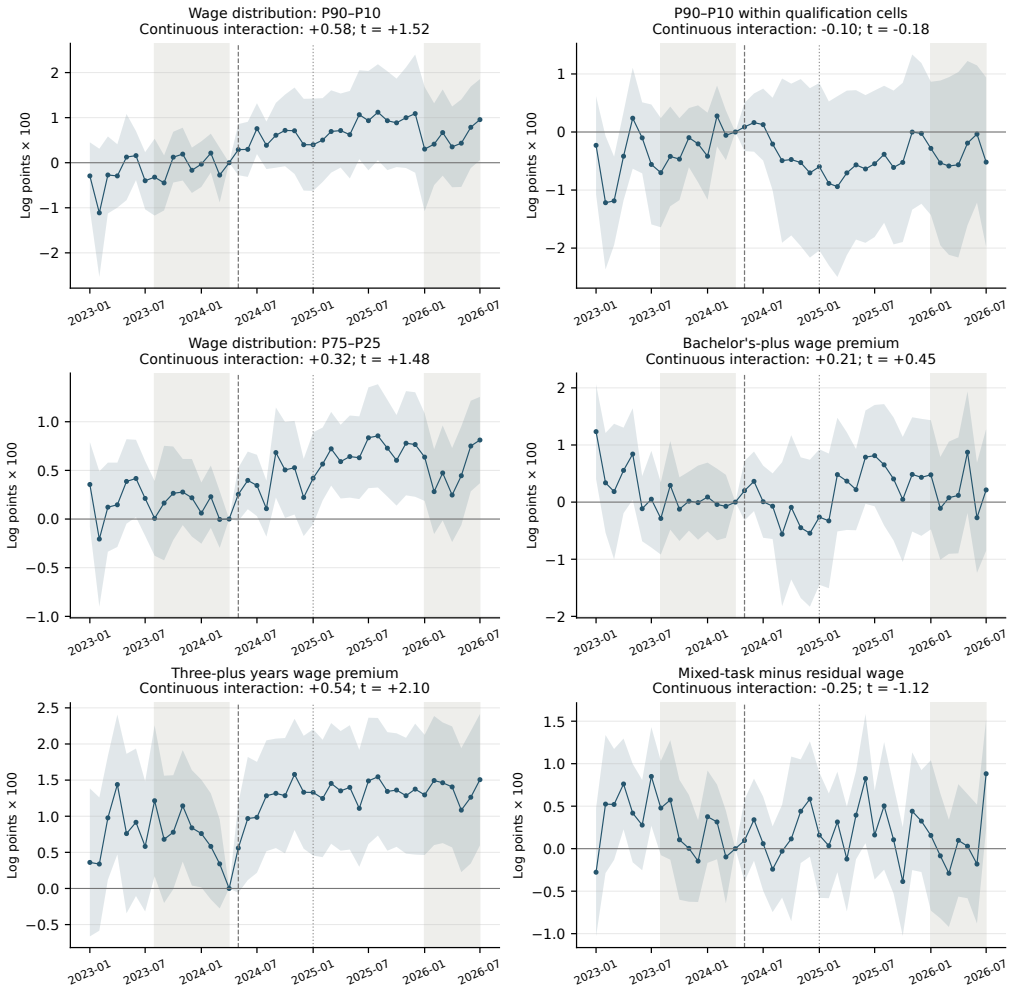


Figure 10: Within-occupation wage gaps: continuous exposure-by-bundle event studies. Points show β_m; bands are pointwise 95% intervals using O*NET-cluster degrees of freedom. Outcomes, sample sizes, and qualification controls differ across panels as defined in the text. Dates and shaded windows follow Figure 8.

micro and macro levels. At the micro level, Section 5 classifies occupations by the strength of their task bundles and finds that the responses of posted task content, residual roles, and posted wages to AI exposure differ with bundle strength. One key object is to measure the coordination cost between tasks, that is, the output loss due to breaking two given tasks apart. While coordination cost in our model is reduced-form, we believe that it is possible to study its main determinants. Empirically, c is distinct from low autonomous AI capability and from worker skill composition: AI performance on Task 1 could be the same in two occupations and skill distributions could also be the same, but they would have different outcomes if separating the tasks is different in the two occupations. The first determinant is the importance of shared context. For example, when the entity (human or machine) that classifies images is the one who also talks to the referring physician, splitting the two tasks means the classifier lacks clinical context and the communicator lacks detail. Second, liability may also raise coordination cost. When a single professional must be liable for the full service, separating tasks introduces a verification problem, which, depending on the matter, may prevent full automation of the task and keep the bundle together. Finally, in general, any cross-task spillovers tend to increase coordination cost: what the worker learns doing Task 1 improves how she does Task 2, and vice versa. These three determinants are the ones our O*NET index measures; a single component explains 58 percent of their joint variance across occupations.

At the macro level, one key objective is to study the impact of AI on the aggregate labor share. Bundling can offset the standard task-substitution channel. If most jobs are strong bundles, the aggregate labor share is more resilient than the task-exposure models predict. Mapping the model's revenue-share objects to the aggregate macro labor share would require additional assumptions about who owns the AI capital embodied in μ_H and μ_A.

Our work offers a first step to study the impact of AI on the shift from bundled production to unbundled production. One important extension is to embed the task choice within team production with various human specialists. This will help make sense of Dell'Acqua et al. (2025), which suggests that given the high coordination cost of human to human collaboration, a cybernetic teammate may beat a human one.

A Proofs

Proof of Lemma 1. Fix (s, m). Using (4), the bundled worker solves

$$\max_{t_1,t_2\geq 0} \left[\mu_H(m)^{1-\eta}(t_1 s_1)^{\eta}\right]^{\alpha}(t_2 s_2)^{1-\alpha} \quad \text{subject to} \quad t_1 + t_2 \leq 1.$$

Because the objective is strictly increasing in both t_1 and t_2, the time constraint (2) binds at the optimum. Hence $t_2 = 1 - t_1$, and the problem reduces to

$$\max_{t_1 \in [0,1]} \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta} s_2^{1-\alpha} t_1^{\alpha\eta} (1 - t_1)^{1-\alpha}.$$

The multiplicative term

$$C \equiv \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta} s_2^{1-\alpha} > 0$$

is constant with respect to t_1, so it suffices to maximize

$$f(t_1) \equiv t_1^{\alpha\eta} (1 - t_1)^{1-\alpha}.$$

Equivalently, maximize

$$\ln f(t_1) = \alpha\eta \ln t_1 + (1 - \alpha) \ln(1 - t_1).$$

The first-order condition is

$$\frac{d}{dt_1} \ln f(t_1) = \frac{\alpha\eta}{t_1} - \frac{1 - \alpha}{1 - t_1} = 0,$$

which gives

$$t_1^* = \frac{\alpha\eta}{\alpha\eta + 1 - \alpha}.$$

Using (7), this is

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

The solution is unique because

$$\frac{d^2}{dt_1^2} \ln f(t_1) = -\frac{\alpha\eta}{t_1^2} - \frac{1 - \alpha}{(1 - t_1)^2} < 0$$

for all $t_1 \in (0, 1)$.

Substituting t_1^* and t_2^* into (4) yields

$$y_H(s, m) = \mu_H(m)^{\alpha(1-\eta)} \left(\frac{\alpha\eta}{\kappa} s_1 \right)^{\alpha\eta} \left(\frac{1 - \alpha}{\kappa} s_2 \right)^{1-\alpha},$$

which can be written as

$$y_H(s, m) = \Omega_H \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta} s_2^{1-\alpha},$$

where

$$\Omega_H = \left(\frac{\alpha\eta}{\kappa}\right)^{\alpha\eta}\left(\frac{1-\alpha}{\kappa}\right)^{1-\alpha}.$$

This is exactly (8)–(9). □

Proof of Proposition 1. Fix (P, m, c). A worker of type $s = (s_1, s_2)$ weakly prefers unbundling to bundling if and only if

$$w_U(P, s_2, m; c) \geq w_H(P, s, m).$$

Using (11), (12), (8), and (5), this condition becomes

$$(1-\alpha)P(1-c)\mu_A(m)^{\alpha}s_2^{1-\alpha} \geq \kappa P\Omega_H\mu_H(m)^{\alpha(1-\eta)}s_1^{\alpha\eta}s_2^{1-\alpha}.$$

Since $P > 0$ and $s_2 > 0$ for all types in $S \subset \mathbb{R}^2_+$, both terms cancel, and we obtain

$$(1-\alpha)(1-c)\mu_A(m)^{\alpha} \geq \kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}s_1^{\alpha\eta}.$$

Rearranging,

$$s_1^{\alpha\eta} \leq \frac{(1-\alpha)(1-c)}{\kappa\Omega_H}\frac{\mu_A(m)^{\alpha}}{\mu_H(m)^{\alpha(1-\eta)}}.$$

Using (13), this is

$$s_1^{\alpha\eta} \leq \frac{(1-\alpha)(1-c)}{\kappa\Omega_H}\Gamma(m),$$

or equivalently

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

which is (15). Thus

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

This comparison is independent of s_2, so conditional on entering the focal occupation, workers with low task-1 skill weakly prefer unbundled production, while workers with high task-1 skill weakly prefer bundled production.

For the comparative statics, take logs of (15):

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

Differentiating with respect to m gives

$$\frac{\partial \ln s_1^*(m,c)}{\partial m} = \frac{1}{\alpha\eta}\frac{\Gamma'(m)}{\Gamma(m)} > 0$$

by Assumption 1. Hence

$$\frac{\partial s_1^*(m,c)}{\partial m} > 0.$$

Differentiating with respect to c gives

$$\frac{\partial \ln s_1^*(m,c)}{\partial c} = -\frac{1}{\alpha\eta(1-c)} < 0,$$

so

$$\frac{\partial s_1^*(m,c)}{\partial c} < 0.$$

□

Proof of Corollary 1. Start from the cutoff in Proposition 1,

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

We first relate the position of $s_1^*(m,c)$ to the support $[\underline{s}_1, \overline{s}_1]$. Using (15), the condition $s_1^*(m,c) \geq \overline{s}_1$ is equivalent to

$$\frac{(1-\alpha)(1-c)}{\kappa\Omega_H}\Gamma(m) \geq \overline{s}_1^{\alpha\eta},$$

that is,

$$(1-\alpha)(1-c)\mu_A(m)^\alpha \geq \kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}\overline{s}_1^{\alpha\eta}.$$

Solving for c gives

$$c \leq 1 - \frac{\kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}\overline{s}_1^{\alpha\eta}}{(1-\alpha)\mu_A(m)^\alpha} \equiv \underline{c}(m),$$

which is (17). Hence

$$c \leq \underline{c}(m) \quad \Longleftrightarrow \quad s_1^*(m,c) \geq \overline{s}_1.$$

Similarly, $s_1^*(m,c) \leq \underline{s}_1$ is equivalent to

$$(1-\alpha)(1-c)\mu_A(m)^\alpha \leq \kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}\underline{s}_1^{\alpha\eta},$$

or

$$c \geq 1 - \frac{\kappa \Omega_H \mu_H(m)^{\alpha(1-\eta)} \underline{s}_1^{\alpha\eta}}{(1-\alpha)\mu_A(m)^{\alpha}} \equiv \bar{c}(m),$$

which is (18). Hence

$$c \geq \bar{c}(m) \quad \Longleftrightarrow \quad s_1^*(m,c) \leq \underline{s}_1.$$

We now verify the ordering of the two cutoffs. By Assumption 2,

$$\kappa \Omega_H \mu_H(m)^{\alpha(1-\eta)} \overline{s}_1^{\alpha\eta} \leq (1-\alpha)\mu_A(m)^{\alpha},$$

so $\underline{c}(m) \geq 0$. Since $\underline{s}_1 \leq \overline{s}_1$,

$$\underline{s}_1^{\alpha\eta} \leq \overline{s}_1^{\alpha\eta},$$

which implies

$$\frac{\kappa \Omega_H \mu_H(m)^{\alpha(1-\eta)} \underline{s}_1^{\alpha\eta}}{(1-\alpha)\mu_A(m)^{\alpha}} \leq \frac{\kappa \Omega_H \mu_H(m)^{\alpha(1-\eta)} \overline{s}_1^{\alpha\eta}}{(1-\alpha)\mu_A(m)^{\alpha}}.$$

Therefore

$$\underline{c}(m) \leq \bar{c}(m).$$

Finally, because $\underline{s}_1 > 0$, the ratio subtracted from 1 in (18) is strictly positive, so $\bar{c}(m) < 1$. Combining these inequalities gives

$$0 \leq \underline{c}(m) \leq \bar{c}(m) < 1.$$

For part (i), if $c \leq \underline{c}(m)$, then $s_1^*(m,c) \geq \overline{s}_1$. Hence every worker type in the support satisfies $s_1 \leq s_1^*(m,c)$. By Proposition 1, every worker who enters the focal occupation weakly prefers unbundled production.

For part (ii), if $c \geq \bar{c}(m)$, then $s_1^*(m,c) \leq \underline{s}_1$. Hence every worker type in the support satisfies $s_1 \geq s_1^*(m,c)$. By Proposition 1, every worker who enters the focal occupation weakly prefers bundled production.

For part (iii), if $c \in (\underline{c}(m), \bar{c}(m))$, then

$$s_1^*(m,c) \in (\underline{s}_1, \overline{s}_1).$$

Hence there exist worker types with $s_1 < s_1^*(m,c)$ and worker types with $s_1 > s_1^*(m,c)$. By Proposition 1, the former weakly prefer unbundled production and the latter weakly prefer bundled production. Thus both organizational forms are privately optimal for some worker types. □

Proof of Lemma 3. Fix $(P, \bar{U}, m, c)$ and let $\tau = \bar{U}/P$ as in (19). We construct an alloca-

tion and then show that it is the unique equilibrium up to zero-measure tie sets.

Define the sets $\mathcal{U}(\tau, m, c)$ and $H(\tau, m, c)$ by (23) and (24), and let the outside set be

$$O(\tau, m, c) \equiv S \setminus \Big(\mathcal{U}(\tau, m, c) \cup H(\tau, m, c)\Big).$$

Because $s_1^*(m, c)$ is a number, $b_U(\tau, m, c)$ is a number, and $b_H(s_1; \tau, m)$ is continuous in s_1, all three sets are measurable. We show that assigning workers according to this partition is a competitive equilibrium.

Take any worker type $s = (s_1, s_2)$.

If $s_1 < s_1^*(m, c)$, then Proposition 1 implies

$$w_U(P, s_2, m; c) \geq w_H(P, s, m).$$

There are then two possibilities. If $s_2 \geq b_U(\tau, m, c)$, then by (22),

$$w_U(P, s_2, m; c) \geq \bar{U}.$$

Since $w_U \geq w_H$, unbundled production attains the maximum in (20), so $s \in \mathcal{U}(\tau, m, c)$. If instead $s_2 < b_U(\tau, m, c)$, then (22) implies

$$w_U(P, s_2, m; c) < \bar{U}.$$

Because $w_H \leq w_U$, bundled production also pays less than $\bar{U}$. Hence outside work attains the maximum in (20), so $s \in O(\tau, m, c)$.

If $s_1 > s_1^*(m, c)$, then Proposition 1 implies

$$w_H(P, s, m) > w_U(P, s_2, m; c).$$

Again there are two possibilities. If $s_2 \geq b_H(s_1; \tau, m)$, then by (21),

$$w_H(P, s, m) \geq \bar{U}.$$

Since $w_H > w_U$, bundled production attains the maximum in (20), so $s \in H(\tau, m, c)$. If instead $s_2 < b_H(s_1; \tau, m)$, then (21) implies

$$w_H(P, s, m) < \bar{U}.$$

Because $w_U < w_H$, unbundled production also pays less than $\bar{U}$. Hence outside work attains

the maximum in (20), so $s \in O(\tau, m, c)$.

Finally, if $s_1 = s_1^*(m, c)$, then Proposition 1 implies

$$w_H(P, s, m) = w_U(P, s_2, m; c).$$

If this common inside payoff is at least $\bar{U}$, either inside option is optimal; the selection rule in (24) assigns such types to bundled production. If the common inside payoff is below $\bar{U}$, the worker chooses outside work. Thus every worker type is assigned an option that attains the maximum in (20). This establishes existence.

For uniqueness, note that once $(P, \bar{U}, m, c)$ are fixed, the cutoff $s_1^*(m, c)$ and the participation thresholds $b_U(\tau, m, c)$ and $b_H(s_1; \tau, m)$ are uniquely determined by (15), (22), and (21). Hence every worker type with a strict ranking among the three options is assigned uniquely.

Possible multiplicity arises only on indifference sets:

$$s_1 = s_1^*(m, c), \qquad s_2 = b_U(\tau, m, c), \qquad s_2 = b_H(s_1; \tau, m).$$

The first is a vertical line in S, the second is a horizontal line, and the third is a continuous curve. Each is a one-dimensional subset of the two-dimensional set S, hence has measure zero. Therefore the competitive equilibrium with fixed product price is unique up to zero-measure tie sets. □

Proof of Proposition 2. We first derive the participation thresholds (21) and (22). A worker entering bundled production must earn at least the outside option, so $w_H(P, s, m) \geq \bar{U}$. Using the definition of bundled earnings from (11) and maximized output from (8), this condition is:

$$\kappa P \Omega_H \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta} s_2^{1-\alpha} \geq \bar{U}.$$

Rearranging for $s_2^{1-\alpha}$ and using the relative price definition $\tau \equiv \bar{U}/P$ from (19), we obtain:

$$s_2^{1-\alpha} \geq \frac{\tau}{\kappa \Omega_H \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta}}.$$

Raising both sides to the power of $1/(1-\alpha)$ yields exactly the bundled participation threshold $b_H(s_1; \tau, m)$ defined in (21).

Similarly, a worker entering unbundled production requires $w_U(P, s_2, m; c) \geq \bar{U}$. Using (12) and (5), this is:

$$(1-\alpha) P (1-c) \mu_A(m)^{\alpha} s_2^{1-\alpha} \geq \bar{U}.$$

Rearranging for $s_2^{1-\alpha}$ and substituting τ yields:

$$s_2^{1-\alpha} \geq \frac{\tau}{(1-\alpha)(1-c)\mu_A(m)^\alpha}.$$

Raising both sides to the power of $1/(1-\alpha)$ gives the unbundled participation threshold $b_U(\tau, m, c)$ defined in (22).

Now, we compare the two thresholds. Noting that the exponent $1/(1-\alpha)$ is strictly positive, we have

$$b_U(\tau, m, c) \leq b_H(s_1; \tau, m)$$

if and only if the bases satisfy the same inequality:

$$\frac{\tau}{(1-\alpha)(1-c)\mu_A(m)^\alpha} \leq \frac{\tau}{\kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}s_1^{\alpha\eta}}.$$

Since $\tau > 0$, we can divide both sides by τ and cross-multiply to obtain:

$$(1-\alpha)(1-c)\mu_A(m)^\alpha \geq \kappa\Omega_H\mu_H(m)^{\alpha(1-\eta)}s_1^{\alpha\eta}.$$

But this is exactly the inequality that characterizes the preference for unbundling over bundling in the proof of Proposition 1. Hence,

$$b_U(\tau, m, c) \leq b_H(s_1; \tau, m) \quad \Longleftrightarrow \quad s_1 \leq s_1^*(m, c),$$

with strict inequality if and only if $s_1 < s_1^*(m, c)$.

Therefore, for every worker type that weakly prefers unbundling, the participation cutoff for the residual human role is weakly lower than the participation cutoff for bundled production. In that sense, once task 1 is peeled away, entry into the occupation becomes weakly easier. □

Proof of Proposition 3. Suppose the fixed-price equilibrium is interior. Write, for brevity,

$$s_1^* \equiv s_1^*(m, c), \qquad b_U \equiv b_U(\tau, m, c).$$

By (23), the unbundled active set is

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

Hence unbundled employment is

$$N_U(\tau, m, c) = \int_{\underline{s}_1}^{s_1^*} \int_{b_U}^{\bar{s}_2} g(s_1, s_2)\, ds_2\, ds_1.$$

Because the equilibrium is interior, both limits vary in the interior of the support, so Leibniz's rule applies. Differentiating with respect to m gives

$$\frac{dN_U(\tau, m, c)}{dm} = \left(\int_{b_U}^{\bar{s}_2} g(s_1^*, s_2)\, ds_2 \right) \frac{ds_1^*(m, c)}{dm} - \left(\int_{\underline{s}_1}^{s_1^*} g(s_1, b_U)\, ds_1 \right) \frac{db_U(\tau, m, c)}{dm},$$

which is exactly (27).

Next, differentiate (22). Since τ, c, and α are fixed in the fixed-price benchmark,

$$\ln b_U(\tau, m, c) = \frac{1}{1-\alpha} \Big[\ln \tau - \ln(1-\alpha) - \ln(1-c) - \alpha \ln \mu_A(m) \Big].$$

Therefore,

$$\frac{d \ln b_U(\tau, m, c)}{dm} = -\frac{\alpha}{1-\alpha} \frac{d \ln \mu_A(m)}{dm},$$

which is (28). Since μ_A is strictly increasing in m by (3), it follows that

$$\frac{d \ln b_U(\tau, m, c)}{dm} < 0 \qquad \Longrightarrow \qquad \frac{db_U(\tau, m, c)}{dm} < 0.$$

Finally, by Proposition 1,

$$\frac{ds_1^*(m, c)}{dm} > 0.$$

Because g is strictly positive on S, the first term in (27) is strictly positive. Since $db_U/dm < 0$, the second term in (27) is also strictly positive. Hence

$$\frac{dN_U(\tau, m, c)}{dm} > 0.$$

Thus better machine capability expands unbundled employment in the fixed-price benchmark. □

Proof of Proposition 4. For fixed $(\bar{U}, m, c)$, let

$$Y(\tau, m, c) \equiv Y_H(\tau, m, c) + Y_U(\tau, m, c),$$

where Y_H and Y_U are defined in (25)–(26). By (31), an endogenous-price equilibrium is a

fixed point of

$$\tau = F(\tau; m, c) \equiv \bar{U}\, Y(\tau, m, c)^{1/\varepsilon}.$$

We first show that $Y(\tau, m, c)$ is continuous and weakly decreasing in τ. By (21) and (22), both participation thresholds $b_H(s_1; \tau, m)$ and $b_U(\tau, m, c)$ are continuous and strictly increasing in τ. Hence, by (23) and (24), if $\tau' > \tau$, then

$$\mathcal{U}(\tau', m, c) \subseteq \mathcal{U}(\tau, m, c), \qquad H(\tau', m, c) \subseteq H(\tau, m, c).$$

Therefore,

$$Y(\tau', m, c) \leq Y(\tau, m, c).$$

To prove continuity, let $\tau_n \to \tau$. Since the sorting cutoff $s_1^*(m, c)$ in (15) does not depend on τ, the only points at which the indicators of the active sets can fail to converge are points on the participation boundaries

$$s_2 = b_U(\tau, m, c) \qquad \text{or} \qquad s_2 = b_H(s_1; \tau, m).$$

These are one-dimensional subsets of the compact set S, hence have measure zero. By (8) and (5), the functions y_H and y_U are continuous and bounded on S, so dominated convergence implies

$$Y(\tau_n, m, c) \to Y(\tau, m, c).$$

Thus $Y(\tau, m, c)$ is continuous in τ, and therefore $F(\tau; m, c)$ is continuous and weakly decreasing as well.

For existence, note that when $\tau = 0$, the participation thresholds in (21) and (22) are both zero, so every worker chooses one of the two inside options and $Y(0, m, c) > 0$. Hence

$$F(0; m, c) = \bar{U}\, Y(0, m, c)^{1/\varepsilon} > 0.$$

Next define normalized inside earnings:

$$\omega_H(s; m) \equiv \frac{w_H(P, s, m)}{P} = \kappa y_H(s, m), \qquad \omega_U(s; m, c) \equiv \frac{w_U(P, s_2, m; c)}{P} = (1-\alpha) y_U(s_2, m; c),$$

where the equalities use (11) and (12). By (8) and (5), both functions are continuous on S, so compactness implies

$$M(m, c) \equiv \max_{s \in S} \max\{\omega_H(s; m), \omega_U(s; m, c)\} < \infty.$$

If $\tau > M(m,c)$, then both inside options pay less than the outside option for every worker, so no worker enters and $Y(\tau,m,c)=0$. Hence

$$F(\tau;m,c)=0<\tau \qquad \text{for all } \tau > M(m,c).$$

Therefore the continuous function

$$G(\tau)\equiv F(\tau;m,c)-\tau$$

satisfies $G(0)>0$ and $G(\tau)<0$ for all sufficiently large τ. By the intermediate value theorem, there exists at least one fixed point $\tau^*(m,c)$.

To show uniqueness, suppose $\tau_1<\tau_2$ are two fixed points. Since F is weakly decreasing,

$$\tau_2=F(\tau_2;m,c)\le F(\tau_1;m,c)=\tau_1,$$

a contradiction. Hence the fixed point $\tau^*(m,c)$ is unique. By (30), this also pins down a unique equilibrium price

$$P^*(m,c)=\frac{\bar{U}}{\tau^*(m,c)}.$$

Finally, this unique $\tau^*(m,c)$ determines the participation thresholds $b_U(\tau^*,m,c)$ and $b_H(s_1;\tau^*,m)$ through (22) and (21), and therefore determines the sets of unbundled workers $\mathcal{U}(\tau^*,m,c)$, bundled workers $H(\tau^*,m,c)$, and outsiders

$$S\setminus\Big(\mathcal{U}(\tau^*,m,c)\cup H(\tau^*,m,c)\Big)$$

through (23) and (24). This is the sense in which the induced allocation is unique: the partition of worker types across the three options is uniquely determined, except on indifference boundaries—namely $s_1=s_1^*(m,c)$, $s_2=b_U(\tau^*,m,c)$, and $s_2=b_H(s_1;\tau^*,m)$—which have measure zero. Thus the endogenous-price competitive equilibrium is unique up to zero-measure tie sets. □

Proof of Proposition 5. In pure unbundling, all active workers choose the residual human role. Hence, by (23),

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

so

$$N_U^*(m,c)=\int_{\underline{s}_1}^{\bar{s}_1}\int_{b_U^*(m,c)}^{\bar{s}_2} g(s_1,s_2)\,ds_2\,ds_1.$$

Because $b_U^*(m,c) \in (\underline{s}_2, \overline{s}_2)$, Leibniz's rule yields

$$\frac{dN_U^*(m,c)}{dm} = -\left(\int_{\underline{s}_1}^{\overline{s}_1} g(s_1, b_U^*(m,c))\, ds_1\right) \frac{db_U^*(m,c)}{dm},$$

which is exactly (35).

Since g is strictly positive on S, the term in parentheses is strictly positive. Hence

$$\frac{dN_U^*(m,c)}{dm} < 0 \iff \frac{db_U^*(m,c)}{dm} > 0.$$

Because $b_U^*(m,c) > 0$, the sign of $db_U^*(m,c)/dm$ is the same as the sign of $d\ln b_U^*(m,c)/dm$. Using (34), we obtain

$$\frac{db_U^*(m,c)}{dm} > 0 \iff \frac{1}{\varepsilon}\frac{d\ln Y^*(m,c)}{dm} > \alpha \frac{d\ln \mu_A(m)}{dm},$$

which is exactly (36). Therefore human employment falls if and only if (36) holds. □

Proof of Proposition 6. In the interior regime, the unbundled active set is given by (23), so

$$N_U^*(m,c) = \int_{\underline{s}_1}^{s_1^*(m,c)} \int_{b_U^*(m,c)}^{\overline{s}_2} g(s_1, s_2)\, ds_2\, ds_1.$$

Because $s_1^*(m,c) \in (\underline{s}_1, \overline{s}_1)$ and $b_U^*(m,c) \in (\underline{s}_2, \overline{s}_2)$, Leibniz's rule yields

$$\frac{dN_U^*(m,c)}{dm} = \left(\int_{b_U^*(m,c)}^{\overline{s}_2} g(s_1^*(m,c), s_2)\, ds_2\right) \frac{ds_1^*(m,c)}{dm} - \left(\int_{\underline{s}_1}^{s_1^*(m,c)} g(s_1, b_U^*(m,c))\, ds_1\right) \frac{db_U^*(m,c)}{dm},$$

which is exactly (37).

The first term is the sorting effect. By Proposition 1,

$$\frac{ds_1^*(m,c)}{dm} > 0,$$

so, since $g > 0$ on S, this term is strictly positive. It captures workers who switch from bundled to unbundled production as the sorting cutoff $s_1^*(m,c)$ in (15) shifts to the right.

The second term is the threshold effect. Its sign is determined by $db_U^*(m,c)/dm$. In the fixed-price benchmark, (28) implies $db_U(\tau,m,c)/dm < 0$, so the threshold effect is positive: better AI lowers the participation cutoff and draws in additional workers. Under endogenous prices, however, (34) shows that price compression can dominate the direct productivity gain, making $db_U^*(m,c)/dm > 0$; in that case the threshold effect becomes negative. This establishes the decomposition and sign discussion in the proposition. □

Proof of Corollary 2. The formulas (39) and (41) follow from Lemma 2, evaluated at the endogenous-price equilibrium price $P^*(m,c)$, together with (8) and (5). Thus

$$w_H^*(s,m,c) = \kappa P^*(m,c)\Omega_H \mu_H(m)^{\alpha(1-\eta)} s_1^{\alpha\eta} s_2^{1-\alpha},$$

and

$$w_U^*(s_2,m,c) = (1-\alpha)P^*(m,c)(1-c)\mu_A(m)^{\alpha} s_2^{1-\alpha}.$$

Taking logs and differentiating with respect to m gives

$$\frac{d\ln w_H^*(s,m,c)}{dm} = \frac{d\ln P^*(m,c)}{dm} + \alpha(1-\eta)\frac{d\ln \mu_H(m)}{dm},$$

and

$$\frac{d\ln w_U^*(s_2,m,c)}{dm} = \frac{d\ln P^*(m,c)}{dm} + \alpha\frac{d\ln \mu_A(m)}{dm}.$$

It remains to show directly that the price effect is negative.

Fix c, and let $m' > m$. Let $\tau^*(m,c)$ denote the unique equilibrium value of τ from Proposition 4. Consider output at the common relative price $\tau^*(m,c)$.

First take a type $s \in \mathcal{U}(\tau^*(m,c),m,c)$, as defined in (23). Then $s_1 < s_1^*(m,c)$. Since Proposition 1 implies that $s_1^*(m,c)$ is increasing in m, we also have $s_1 < s_1^*(m',c)$. Thus this type still weakly prefers unbundling at m'. Since μ_A is strictly increasing in m by (3),

$$y_U(s_2,m';c) > y_U(s_2,m;c).$$

Hence every type that was unbundled at m produces strictly more at m', when evaluated at the same $\tau^*(m,c)$.

Next take a type $s \in H(\tau^*(m,c),m,c)$, as defined in (24). Since μ_H is strictly increasing in m,

$$y_H(s,m') > y_H(s,m).$$

At m', such a type either remains bundled or switches to unbundling. If it remains bundled, its output is $y_H(s,m') > y_H(s,m)$. If it switches to unbundling, then by Lemma 2,

$$(1-\alpha)y_U(s_2,m';c) \geq \kappa y_H(s,m'),$$

so

$$y_U(s_2,m';c) \geq \frac{\kappa}{1-\alpha} y_H(s,m') > y_H(s,m'),$$

because $\kappa > 1-\alpha$. Hence every type that was bundled at m also produces strictly more at

m', when evaluated at the same $\tau^*(m,c)$.

Therefore every previously active type produces strictly more at m', while previously inactive types contribute weakly more than before. Since the equilibrium output $Y^*(m,c) = Y(\tau^*(m,c),m,c)$ is strictly positive, the set of active types has positive measure, and so aggregate output strictly rises:

$$Y(\tau^*(m,c),m',c) > Y(\tau^*(m,c),m,c) = Y^*(m,c).$$

Now define the fixed-point map

$$F(\tau;m,c) \equiv \bar{U}\, Y(\tau,m,c)^{1/\varepsilon},$$

as in (31). The inequality above implies

$$F(\tau^*(m,c);m',c) > F(\tau^*(m,c);m,c) = \tau^*(m,c).$$

By Proposition 4, for fixed (m',c), the map $F(\cdot;m',c)$ is weakly decreasing in τ and has a unique fixed point. Hence

$$\tau^*(m',c) > \tau^*(m,c).$$

Using (30),

$$P^*(m,c) = \frac{\bar{U}}{\tau^*(m,c)},$$

so

$$P^*(m',c) < P^*(m,c).$$

Thus the equilibrium product price is strictly decreasing in m. Equivalently, wherever the derivative exists,

$$\frac{d\ln P^*(m,c)}{dm} < 0.$$

Substituting this sign into the two log-derivative formulas above yields (40) and (42), with the price effect signed directly as negative. Hence fixed-type earnings fall whenever this negative price effect dominates the corresponding direct productivity effect. □

B Sample sizes and additional estimates

Table 6: China core-outcome sample sizes

Outcome	Occupations	O*NET clusters	Observations
GenAI-tool requirement share	829	245	87,308
Exposed-task share	829	245	87,014
Residual-posting share	829	245	87,108
Advertised wage (log)	881	255	102,186
Bachelor's degree or above	804	243	57,285
At least three years of experience	754	239	38,545
New postings (log)	769	240	69,520

Notes: Occupations and mapped O*NET clusters are counted in the eligible outcome panel; observations are the static regression's estimation count after removing singleton fixed effects. Monthly and static specifications use the same outcome panel.

Table 7: Within-occupation wage gaps: continuous static post-May-2024 estimates

Wage-gap outcome	Continuous DDD	SE	t
P90–P10	0.798	0.352	2.26
Within-cell P90–P10	-0.109	0.463	-0.24
P75–P25	0.378	0.183	2.06
Education premium	0.044	0.416	0.11
Experience premium	0.554	0.161	3.45
Mixed-task minus residual	-0.116	0.134	-0.86

Notes: Units and columns match Table 5. These estimates pool all months on each side of the May 2024 cutoff and therefore differ from the late-minus-pre event-study contrasts.

Table 8: Within-occupation wage gaps: continuous-model diagnostics and samples

Outcome	Pretrend t	Detrended t	Occupations	Clusters	Observations
P90–P10	0.28	-0.28	900	259	108,127
Within-cell P90–P10	1.85	-1.67	830	245	74,656
P75–P25	-0.33	0.54	900	259	108,127
Education premium	0.09	0.05	828	249	77,855
Experience premium	-1.63	2.03	785	242	59,782
Mixed-task minus residual	-0.96	1.17	806	245	75,759

Notes: Both diagnostic t statistics concern the continuous interaction. Windows are defined in the main text. Occupations and clusters are counted in the eligible panel; observations are the static model's estimation count after removing singleton fixed effects.

C Rank-based grouping as a separate sensitivity specification

For comparison with a different functional form, define $H_o = 1$ when PC1 is strictly above the median of the outcome's eligible occupations and $L_o = 1 - H_o$. Tied values remain together in the at-or-below-median group. These are sample-relative ranks, not validated weak- and strong-bundle regimes. The pooled event study is

$$y_{ut} = \alpha_u + \lambda_{pt} + \sum_{m \neq 2024\text{-}04} \mathbf{1}[t = m](\beta_m^H E_o H_o + \beta_m^L E_o L_o + \delta_m^H H_o) + \varepsilon_{ut}. \tag{51}$$

The panel, exposure standardization, fixed effects, weights, and clustering remain unchanged. Replacing continuous M by H changes both the interaction and the bundle-by-month controls; $\beta_m^H - \beta_m^L$ is therefore a different estimand from β_m in the main model. Its standard error uses the jointly estimated covariance.

Table 9 retains the split results for all four main outcomes and all six wage-gap outcomes. The exposed-task and residual-posting differences are not significant, while the mean-wage difference is positive and significant. That mean-wage contrast has a detectable pretrend ($t = 2.27$) and is not significant after detrending ($t = -0.67$). The raw P90–P10 and experience-premium split differences are also significant, but the former is not significant after detrending ($t = 1.27$), and the latter has a detectable pretrend ($t = -2.08$). These results neither validate the median as an economic boundary nor imply that a non-significant binary contrast refutes a continuous gradient.

For exposed-task and residual-posting shares, both models use 829 occupations and 245 O*NET clusters. After removing singleton fixed effects, their common estimation samples contain 87,014 and 87,108 observations, respectively. Re-estimating the continuous specification reproduces its original monthly coefficients to within 10^{-15} while holding y, E, and weights unchanged. The different binary results thus arise from the changed specification rather than a change in the outcome definition or input sample. In these panels, the median partitions occupations into 415 at or below and 414 above, while 39.0% of the index's total cross-occupation variation remains within the two groups. That variance share quantifies variation not represented by H; it is not an estimate of lost statistical power.

Table 9: Alternative rank-split model: late-minus-pre contrasts

Outcome	L	H	H − L	Difference SE	Difference t
Main outcomes					
GenAI-tool requirement share	1.289	0.863	-0.426	0.361	-1.18
Exposed-task share	-0.109	0.054	0.162	0.217	0.75
Residual-posting share	0.579	0.256	-0.323	0.510	-0.63
Advertised wage (log)	-3.314	-1.191	2.124	0.852	2.49
Within-occupation wage gaps					
P90–P10	0.526	2.141	1.614	0.749	2.15
Within-cell P90–P10	2.425	2.858	0.434	0.998	0.43
P75–P25	0.449	1.193	0.744	0.436	1.71
Education premium	-0.574	0.640	1.214	1.017	1.19
Experience premium	-0.409	0.810	1.219	0.517	2.36
Mixed-task minus residual	0.582	0.112	-0.471	0.588	-0.80

Notes: L denotes PC1 at or below the sample median; H denotes PC1 above it. The first two columns report within-group exposure gradients; H−L is their difference. SE and t refer to that difference. Shares are percentage points; wage and wage-gap coefficients are log points multiplied by 100. Units differ from the continuous interaction, which scales M by one standard deviation.

References

Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, volume 4B, chapter 12, pages 1043–1171. Elsevier, 2011.

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.

David Autor and Neil Thompson. Expertise. NBER Working Paper No. 33941, 2025.

David H. Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333, 2003.

Gary S. Becker and Kevin M. Murphy. The division of labor, coordination costs, and knowledge. *Quarterly Journal of Economics*, 107(4):1137–1160, 1992.

Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. Generative AI at work. NBER Working Paper No. 31161, 2023.

Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab working paper, revised August 2026, 2025.

Leland D. Crane and Paul E. Soto. AI and coder employment: Compiling the evidence. Finance and Economics Discussion Series 2026-018, Board of Governors of the Federal Reserve System, 2026.

Fabrizio Dell'Acqua, Charles Ayoubi, Hila Lifshitz, Raffaella Sadun, Ethan Mollick, Lilach Mollick, Yi Han, Jeff Goldman, Hari Nair, Stewart Taub, et al. The cybernetic teammate: A field experiment on generative ai reshaping teamwork and expertise. Technical report, National Bureau of Economic Research, 2025.

Ozge Demirci, Jonas Hannane, and Xinrong Zhu. Who is AI replacing? The impact of generative AI on online freelancing platforms. *Management Science*, 71(10):8097–8108, 2025.

Wouter Dessein and Tano Santos. Adaptive organizations. *Journal of Political Economy*, 114(5):956–995, 2006.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308, 2024.

Lukas B Freund and Lukas F Mann. Job transformation, specialization, and the labor market effects of ai. Technical report, CESifo Working Paper, 2025.

Luis Garicano. Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904, 2000.

Luis Garicano and Esteban Rossi-Hansberg. Organization and inequality in a knowledge economy. *The Quarterly journal of economics*, 121(4):1383–1435, 2006.

Christina Gathmann, Felix Grimm, and Erwin Winkler. AI, task changes in jobs, and worker reallocation. IZA Discussion Paper No. 17554, 2024.

Paweł Gmyrek, Janine Berg, and David Bescond. Generative AI and jobs: A global analysis of potential effects on job quantity and quality. ILO Working Paper 96, International Labour Organization, Geneva, 2023.

Paweł Gmyrek, Janine Berg, Karol Kamiński, Filip Konopczyński, Agnieszka Ładna, Bence Náfrádi, Konrad Rosłaniec, and Marek Troszyński. Generative AI and jobs: A refined global index of occupational exposure. ILO Working Paper 140, International Labour Organization, Geneva, 2025.

Menaka Hampole, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. Artificial intelligence and the labor market. NBER Working Paper No. 33509, 2025.

Seyed Mahdi Hosseini Maasoum and Guy Lichtinger. Generative AI and occupational entry barriers: The labor-supply channel of technological change. SSRN Working Paper No. 6059674, 2026.

Xiang Hui, Oren Reshef, and Luofeng Zhou. The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science*, 35(6):1977–1989, 2024.

Anders Humlum and Emilie Vestergaard. Still waters, rapid currents: Early labor market transformation under generative AI. NBER Working Paper No. 33777, revised March 2026. Earlier version circulated as "Large Language Models, Small Labor Market Effects", 2025.

Enrique Ide and Eduard Talamàs. Artificial intelligence in the knowledge economy. *Journal of Political Economy*, 133(12):3762–3800, 2025.

National Center for O*NET Development. O*NET database, release 31.0. https://www.onetcenter.org/db_releases.html, 2026. August 2026 release.

Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023.

Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. The impact of AI on developer productivity: Evidence from GitHub Copilot. *arXiv preprint arXiv:2302.06590*, 2023.

Sherwin Rosen. Substitution and division of labour. *Economica*, 45(179):235–250, 1978.

Sherwin Rosen. Specialization and human capital. *Journal of Labor Economics*, 1(1):43–49, 1983.