WEAK BUNDLE, STRONG BUNDLE: HOW AI REDRAWS JOB BOUNDARIES

Abstract

This paper studies how the effect of AI on an occupation depends not just on which tasks AI can perform but also on how costly it is to unbundle those tasks from the job. Much of the discussion of AI and labor markets starts from task exposure: if AI can perform more tasks in an occupation, that occupation should lose employment or earnings. This is incomplete because labor markets price jobs, not tasks. Jobs bundle tasks together, and the effect of AI depends on how costly it is to break the bundle. 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, activating the standard task-substitution channel once product demand is sufficiently inelastic. 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. Thus, bundling provides a force that protects jobs and workers' share of downstream revenue.

JEL Classification: J24, J23, O33, L23, D20, J31

Keywords: N/A

Luis Garicano - luis.garicano@gmail.com
London School of Economics and CEPR

Jin Li - jli1@hku.hk
Hong Kong University

Yanhui Wu - yanhuiwu@hku.hk
Hong Kong University and CEPR

Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries

Luis Garicano* Jin Li† Yanhui Wu‡

March 30, 2026

Abstract

This paper studies how the effect of AI on an occupation depends not just on which tasks AI can perform but also on how costly it is to unbundle those tasks from the job. Much of the discussion of AI and labor markets starts from task exposure: if AI can perform more tasks in an occupation, that occupation should lose employment or earnings. This is incomplete because labor markets price jobs, not tasks. Jobs bundle tasks together, and the effect of AI depends on how costly it is to break the bundle. 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, activating the standard task-substitution channel once product demand is sufficiently inelastic. 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. Thus, bundling provides a force that protects jobs and workers' share of downstream revenue.

1 Introduction

In 2016, Geoffrey Hinton declared that medical schools should stop training radiologists because AI would soon outperform them at reading scans. The remark captured a view that still shapes the discussion: once AI can do a task, the workers who do that task lose.[1]

*London School of Economics.

†University of Hong Kong.

‡University of Hong Kong.

[1]Hinton made the remark at the *Machine Learning and the Market for Intelligence* conference in Toronto, November 2016.

That logic underlies a growing empirical literature that measures the AI shock using task exposure (e.g., Eloundou et al., 2024). Occupations with more tasks exposed to AI receive a larger shock, and the inference is monotone: greater exposure means lower employment or lower earnings. But this across-occupation view leaves out an internal organizational margin: job design.

Reading scans is a task. Radiology is a job. A radiologist does not only classify images. She triages cases, communicates with other physicians, trains residents, takes responsibility for difficult calls, and signs a diagnosis that others are willing to act on. What the market buys is not a classification exercise but a bundled service.

We study the survival of the bundle as a result of the tradeoff between the benefit of specialization and the cost of coordination. 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. In a strong bundle, breaking the job destroys enough value that the job survives as a whole: AI assists but the human still sells the full service and retains a large share of revenue. In a weak bundle, the cost of splitting is small: AI replaces some tasks, the human role narrows, and the labor share falls. The question, then, is not only whether AI can perform a task, but whether performing that task separately destroys enough value 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 results that build on one another. First, workers keep a larger share of revenue in bundled production than in unbundled production, because human work contributes to both tasks in the bundle but only to Task 2 outside it. Second, AI pushes the frontier toward unbundled jobs. Third, unbundling produces a discrete capacity shock: workers who previously split time across two tasks now devote everything to Task 2. With

inelastic demand, the resulting flood of output depresses the product price enough to displace the least-productive workers. This nests the standard task-substitution mechanism of Acemoglu and Restrepo (2018, 2019). Fourth, strong bundles protect workers. AI still raises performance on Task 1, but the task remains inside the bundle, so the worker keeps the larger revenue share, and the capacity shock is mitigated.

Several recent patterns fit the model better than a simple "exposure leads to job loss" story. Early evidence suggests that generative AI changes the composition of work more than the level of employment.[2] That is what one expects when 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. The empirical question is not only whether employment falls, but whether workers continue to sell the whole service or are pushed toward a thinner slice. The gains from AI are also asymmetric. Studies of writing, customer service, and coding find larger short-run productivity gains for weaker or less experienced workers (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 also studies jobs as bundles of tasks. Autor and Thompson (2025) study how automation subtracts tasks from an occupation, and how that subtraction changes 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, and wage effects depend on how the new task mix interacts with workers' multidimensional skills. We share the emphasis on bundling and within-occupation heterogeneity but focus on a novel, prior question: the perimeter of the bundle itself is endogenous. AI can either improve performance inside the existing job or cause the job to split, with the outcome governed by coordination costs. This distinction determines whether the task-substitution mechanism operates at all. Empirically, they explain heterogeneous wage effects across workers inside transformed occupations; we explain heterogeneous labor-share 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

[2]Humlum and Vestergaard (2025) link chatbot adoption surveys to Danish administrative records and find occupational switching and task restructuring but precise null effects on earnings and hours. Gathmann et al. (2024) combine patent-based AI exposure measures with German survey and administrative data and find that AI shifts the task content of jobs with small displacement effects. At the level of predicted exposure, Gmyrek et al. (2023) and Gmyrek et al. (2025) find that generative AI is more likely to automate some tasks within occupations than to automate occupations outright, so that most jobs are transformed rather than eliminated.

showing that AI changes occupational entry barriers and effective labor supply by automating some tasks and making others easier. Our contribution is to show that those entry effects depend not only on changing expertise requirements, but also on whether the job bundle survives: once AI unbundles the job, workers no longer need to satisfy the skill requirements of the full bundle in order to enter the residual human role.

More broadly, our paper connects to the task-based automation literature (Acemoglu and Restrepo, 2018, 2019) and to the empirical literature that measures AI exposure at the occupation or task level. That approach maps where AI is likely to matter, but it often treats exposed tasks as independently reallocable. Our unit of analysis is the job bundle. The same improvement in AI on Task 1 can leave the job intact in a strong bundle and redraw it around a narrower human role in a weak bundle. In the weak-bundle case, our model nests the standard displacement mechanism: AI takes some tasks, the labor share falls, and inelastic demand can reduce employment. Bundling is the force that can shut that mechanism down.

The model also relates to multidimensional assignment. In this literature, standard models endow workers with a single ability index (Sattinger, 1993). Lindenlaub (2017) introduces multiple skill dimensions, and Ocampo (2023) endogenizes occupational boundaries in skill space. Our model adds an organizational margin: in bundled production the market values both skill dimensions jointly; in unbundled production it values only task-2 skill, so the job opens to workers who could not have entered the bundled occupation.

Finally, 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. Rosen (1983) studies instead 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) study the trade-off between ex ante coordination versus ex post adaptation. Garicano (2000) shows how communication allows to increase 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 reshapes the allocation of heterogeneous workers across hierarchical layers and the resulting wage distribution.

The paper proceeds as follows. Section 2 states the model, introduces worker time allocation, derives the downstream earnings shares, and establishes the sorting boundary between bundled and unbundled work. Section 3 introduces the outside option and the fixed-price participation margin. It shows that once Task 1 is peeled away, the residual human role becomes easier to access and expands at fixed prices. Section 4 closes the model with in-

elastic demand and shows how the time-release capacity shock generates displacement in weak-bundle occupations. Section 5 discusses interpretation and concludes. Appendix A contains proofs.

2 Model and organizational sorting

We begin with the organizational margin inside the focal occupation. This section abstracts from the outside option and asks when AI stays inside the job or peels Task 1 away from it. The key objects are the human revenue share inside each organizational form and the sorting cutoff between bundled and unbundled production.

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 on the horizon of interest. A worker has type $s = (s_1, s_2)$ drawn from compact support

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

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

Each worker is endowed with one unit of working time. In bundled production the worker must allocate time across the two tasks:

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

Frontier AI capability is denoted by m. 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 production. Throughout, μ_H and μ_A are continuously differentiable and strictly increasing in m.

There are two organizational forms.

Bundled production. 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 job remains bundled. The task-1 input is itself a Cobb–Douglas composite of machine services (weight $1-\eta$) and human time-weighted skill (weight η).

Unbundled production. Autonomous machine cognition 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

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

The coordination loss c is reduced form. It summarizes the value lost when two tasks currently held inside one job are instead supplied by separate agents. Shared context, accountability, communication frictions, spillovers between tasks, and measurement problems can all raise c.

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 $(1 - \alpha)P y_U$. 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.

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

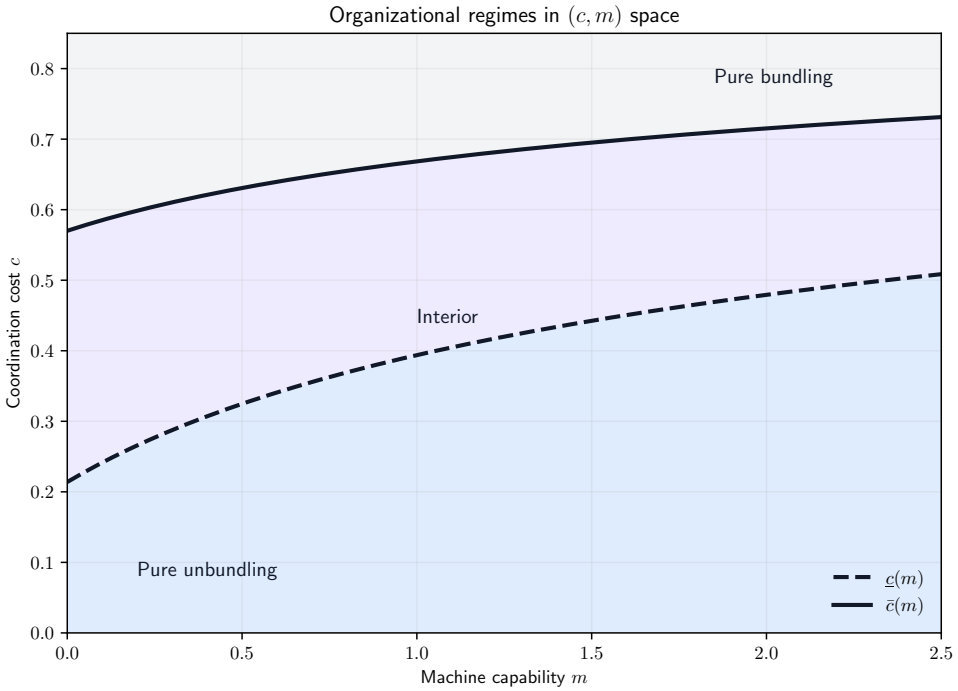


Figure 1: Organizational regimes in (c, m) space. The dashed curve is $\underline{c}(m)$ and the solid curve is $\bar{c}(m)$. Below $\underline{c}(m)$ the occupation is in pure unbundling; between the two curves it is in interior equilibrium; above $\bar{c}(m)$ it is in pure bundling. Benchmark: $S = [0.2, 1]^2$, 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.

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

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

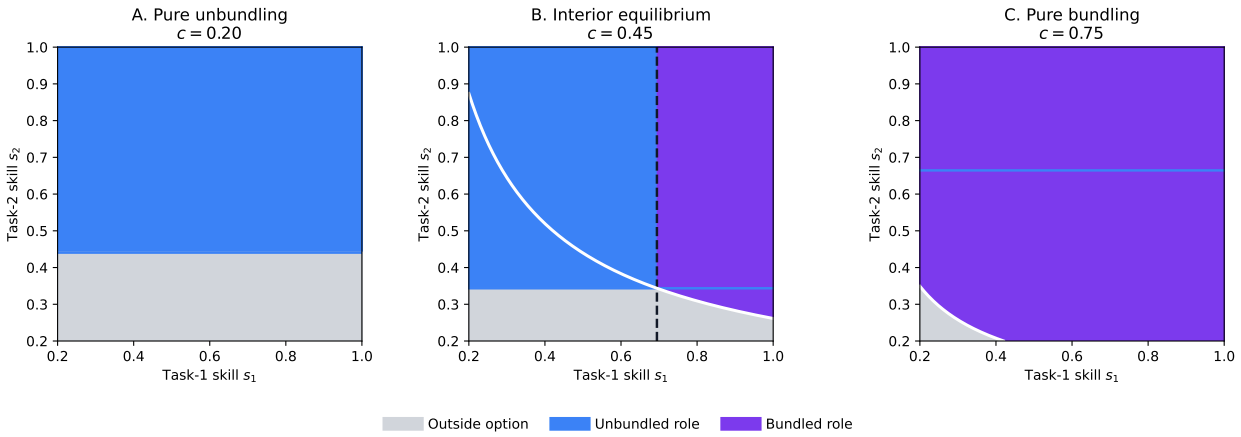


Figure 2: Occupational choice in skill space at $m = 0.8$ under the endogenous-price benchmark used in Figure 3. The three panels show representative coordination costs from pure unbundling, interior equilibrium, and pure bundling. Gray types stay outside the occupation, blue types supply the residual 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.

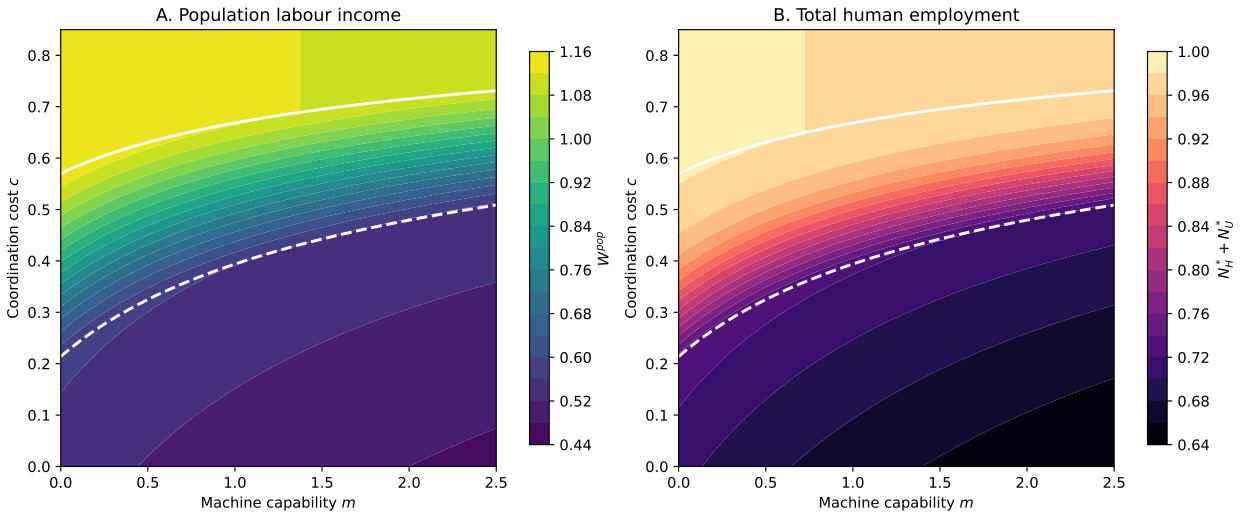


Figure 3: Population labor income and human employment in (c, m) space under inelastic demand. Panel A plots population labor income $W^{pop} = P^*[\kappa Y_H^* + (1-\alpha)Y_U^*]$. Panel B plots total human employment $N_H^* + N_U^*$. The white dashed and solid curves are the organizational cutoffs $\underline{c}(m)$ 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

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

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

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

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.[3] 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.

[3]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.

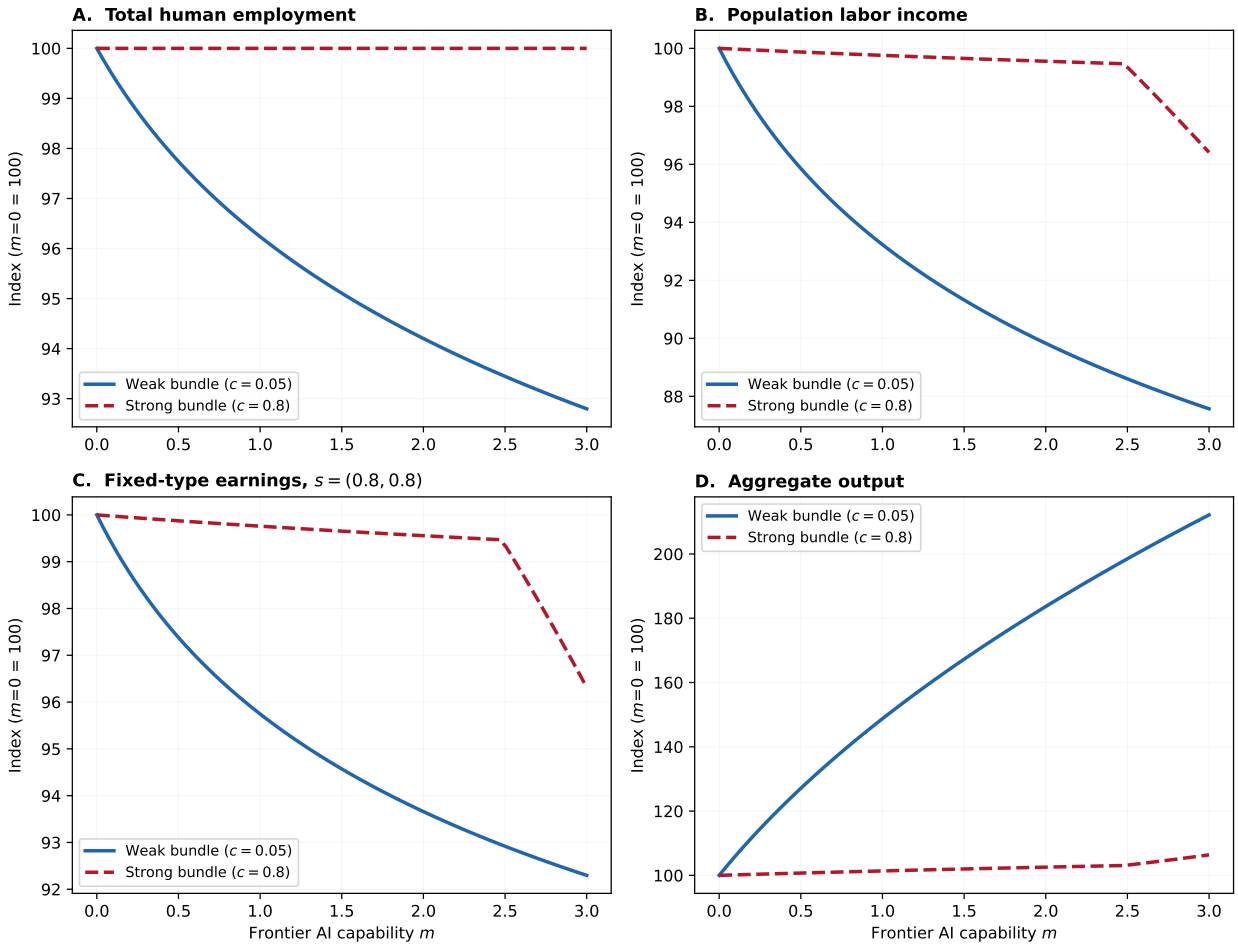


Figure 4: How bundling protects workers. Each panel compares a weak-bundle occupation ($c = 0.05$, solid blue) with a strong-bundle occupation ($c = 0.80$, dashed red) as frontier AI capability m rises. Panel A: total human employment. Panel B: population labor income W^{pop}. Panel C: chosen earnings 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 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 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 the labor outcomes at both the micro and macro levels. At the micro level, researchers can classify occupations by the strength of their task bundles and test whether labor-market responses to AI exposure differ systematically across weak-bundle and strong-bundle occupations. 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: two occupations could face similar measured AI performance on Task 1 and similar worker skill distributions, yet differ sharply in how much output is lost when diagnosis is separated from communication, drafting from approval, or classification from liability-bearing sign-off. 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.

At the macro level, one key objective is to study the impact of AI on the aggregate labor share. The standard task-substitution channel lowers it: AI takes a task, the human loses the associated revenue. Bundling can offset that channel. In a strongly bundled occupation, AI still improves task-1 performance, but the human retains both tasks and a larger revenue share. Whether the aggregate labor share rises or falls depends partly on the distribution of coordination costs between tasks within occupations. If most employment sits in strongly bundled jobs, the aggregate labor share is more resilient than the task-exposure models predict. The distribution of bundle strength across the occupational structure is therefore a first-order empirical question for understanding AI's macroeconomic impact. 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 team mate 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 allocation 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,

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

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),

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

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

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

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. □

References

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.

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.

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.

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.

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.

Anders Humlum and Emilie Vestergaard. Large language models, small labor market effects. NBER Working Paper No. 33777, 2025.

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

Ilse Lindenlaub. Sorting multidimensional types: Theory and application. *Review of Economic Studies*, 84(2):718–789, 2017.

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

Sergio Ocampo. A task-based theory of occupations with multidimensional heterogeneity. Working Paper, University of Western Ontario, 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.

Michael Sattinger. Assignment models of the distribution of earnings. *Journal of Economic Literature*, 31(2):831–880, 1993.