NBER WORKING PAPER SERIES

ENDOGENOUS TASK BUNDLING, SKILLS AND AUTOMATION

Joshua S. Gans

Working Paper 35211
http://www.nber.org/papers/w35211

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

I thank David Autor for helpful comments and Refine.ink, ChatGPT 5.5 and Claude Opus 4.7 for valuable research assistance. Responsibility for all errors remains my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w35211

Endogenous Task Bundling, Skills and Automation
Joshua S. Gans
NBER Working Paper No. 35211
May 2026
JEL No. D23, J23, J24, J31, O33

ABSTRACT

Empirical measures of AI's wage effect typically hold fixed the bundle of activities a worker is paid for at its pre-AI shape. We argue that this assumption hides much of the action. When automation breaks a job apart, firms decide how to recombine the surviving activities; whether they rebundle them into one broad role or split them into specialist roles changes which surviving skills the labour market actually rewards. A skill that played no role in the pre-AI wage can become the dominant component of the post-AI wage, while a skill that anchored the pre-AI wage can disappear from the schedule. We develop an assignment model in which the priced human bundle is endogenous, and we use it to show that a fixed-bundle wage regression can mis-sign the effect of AI exposure. In general, the omitted-redesign bias has no unconditional sign: it is the residual covariance between exposure and role-specific redesign terms. Under explicit sufficient conditions, exposure-correlated unbundling loads specialist comparative-advantage premia onto the exposure coefficient, while exposure-correlated rebundling loads a different, often opposite, omitted term. The sign must therefore be measured from local post-AI partition changes rather than assumed from exposure alone.

Joshua S. Gans
University of Toronto
Rotman School of Management
and NBER
joshua.gans@rotman.utoronto.ca

1 Introduction

A radiologist's day, before AI, runs together image classification and physician consultation. Quality assurance and model governance may sit in a separate technical role. When AI begins to read the images, the old radiologist role no longer exists in human form. The hospital faces a choice: it can recombine consultation with governance into one broad clinical role, or it can split them into two specialist roles—a senior consultant and a model-governance specialist. The two redesigns reward different workers. A radiologist with balanced consultation and governance skills is valuable under recombination. A radiologist with a sharp comparative advantage in consultation or in technical oversight is valuable under separation, but in different specialist roles. And a skill that was never priced in the original bundled role—governance for a radiologist whose paid work centred on diagnosis and consultation—can become the dominant component of the post-AI wage, but only if the hospital chooses the specialist partition and assigns that worker to the governance role.

We ask how automation changes wages when this redesign decision is part of the response. Two empirical research programs measure AI's wage effect by holding the priced bundle of activities fixed at its pre-AI shape. AI-exposure scores quantify which activities large language models can perform (Eloundou et al., 2024). Earnings functions estimated within a fixed occupation describe how the value of remaining expertise responds when some activities are automated (Autor and Thompson, 2025). Both designs assume that the labour market continues to price the same bundle of human activities after automation as before. The radiology vignette suggests this assumption can hide most of the response. The activities the market rewards jointly are not a primitive; they are an equilibrium object that firms can redraw when automation breaks an old role.

The economics of the redesign rests on a familiar trade-off. A bundled role pays for breadth: an above-average worker on each surviving activity earns the same as a strongly specialised worker, because what the firm buys is overall service quality from one person's time. A specialist role pays for comparative advantage: the worker who is far above average on one activity earns a premium, because the firm can pair them with a complementary specialist on the other. Skill dispersion in the workforce and the relative coordination efficiency of the two designs together determine which arrangement firms choose. Automation enters this trade-off through three doors. It removes the human return to the automated activity. It changes the productivity of whichever job design firms select for the survivors. And, because the surviving activities can be either rebundled or split, it changes which surviving skills the market is buying jointly. We refer to this third channel as *rebundling-as-revaluation*: a skill that sat outside a worker's pre-AI role can be revalued upward, and a skill that anchored the pre-AI wage can be revalued out of the schedule.

The fixed-bundle research design is silent on this third channel. It measures the loss of the automated activity and the within-bundle reweighting of the survivors, but not the cross-bundle reweighting that occurs when firms redraw the human boundary. When AI exposure is correlated with post-AI partitions and with the workers sorted into those partitions, the fixed-bundle estimate of AI's wage effect absorbs omitted role-specific redesign payoffs into the exposure coefficient. The sign of the resulting bias is not pinned down by exposure alone, or by a coarse unbundling-versus-rebundling label. It depends on the residual covariance between exposure and the omitted redesign term. If high-exposure occupations tend to split surviving tasks and the omitted specialist premium is positive for the affected workers, the fixed-bundle coefficient loads that comparative-advantage premium onto exposure. If high exposure instead correlates with rebundling, or with workers whose omitted redesign terms have the opposite sign, the bias can run the other way. A labour economist using a single fixed-bundle specification is, therefore, estimating a mixture of redesign effects whose sign depends on the local assignment response and skill distribution, not on exposure by itself.

We formalise this reasoning in an assignment model in which firms choose workers and job designs simultaneously. The model has three activities, one of which AI will eventually perform, and a continuum of workers heterogeneous in the skills they bring to each. Before AI, the activity that would be automated was bundled with one surviving activity in one human role, while the other surviving activity was performed in a separate role at a separate price. We work from this partial pre-AI configuration because it is the smallest environment in which the redesign question is non-trivial. The integrated case, in which all three activities sit in one broad role before AI, is well behaved but does not produce rebundling-as-revaluation, because under integration, every surviving skill was already in the pre-AI bundle. We do not assume the partial configuration is feasible: we study automation only in pre-AI economies that support it as a selected equilibrium under fixed primitives, so the wage comparison runs between two equilibria of the same economy.

The model produces a complete role-by-role wage decomposition. The sharpest case is a worker who began in the bundled role and is reassigned to specialist work on the activity that was previously priced separately—in the radiology metaphor, a governance-leaning radiologist whose technical oversight skill was not priced in the original diagnosis-and-consultation role. Three forces combine for this worker. They lose the return to the automated activity. They gain the return to a previously unpriced surviving skill. And the signed imbalance between their two surviving skills, which was irrelevant to their pre-AI wage, now matters because the post-AI specialist role prices the skill on the previously separate activity. The coefficient on signed imbalance is large and negative: wages fall with a tilt toward the old surviving task and rise with a tilt toward the newly priced specialist task. A fixed-bundle regression cannot recover this term because the relevant skill was not in the pre-AI bundle and the relevant role transition was not in the pre-AI choice set. The other five role transitions have analogous decompositions, recorded in the body of the paper.

Three contributions follow. First, the model is a tractable assignment environment in which the priced human bundle is endogenous, the pre-AI baseline is supported by primitives, and the post-AI partition is determined by the interaction of skill dispersion and relative coordination efficiency. Second, the role-by-role wage decomposition gives a complete worker-level accounting of automation's wage effect and isolates the bundle-redesign channel that fixed-bundle methods miss. Third, the omitted-variable bias on a fixed-bundle AI-exposure coefficient has no sign restriction. Its sign depends on the local covariance between exposure and the full omitted redesign term, a claim that points naturally to occupational data on pre-AI co-tasking, post-AI co-tasking changes and specialist entry.

The paper connects to several strands of work. The task-based automation literature (Acemoglu and Restrepo, 2018, 2019) studies how automation displaces and reinstates labour through activity-level changes; AI-exposure measures (Eloundou et al., 2024) identify which activities are technically exposed; and Autor and Thompson (2025) study how AI changes the value of expertise within a fixed occupational bundle. Our contribution is to ask what these exposure-based research designs miss when the priced bundle is itself endogenous. A second strand studies the division of labour and the organisational cost of splitting work (Becker and Murphy, 1992; Garicano et al., 2026). Garicano et al. (2026) bring this logic to AI in a two-activity environment that rules out re-bundling of the surviving activities by construction. We retain the organisational intuition in the smallest environment in which the post-AI human partition is non-trivial. The Cobb–Douglas activity technology gives a transparent form of quality complementarity in the spirit of Kremer (1993); the time-spreading cost of asking one worker to perform several activities follows the time-allocation logic in Gans and Goldfarb (2026); and the multidimensional sorting machinery is related to Lindenlaub (2017) and Ocampo (2022). Demirer et al. (2026) study which steps in an ordered workflow machines perform; we ask which surviving human skills are priced jointly after the automated activity has been removed.

The rest of the paper proceeds as follows. Section 2 develops the environment and the rebundling mechanism. Section 3 defines production-mode assignment equilibrium and the symmetric wage menu. Section 4 derives skill-balance sorting and the return to comparative advantage. Section 5 develops the partial pre-AI baseline and the worker-level wage decomposition. Section 6 develops the empirical implications: the structural form of the wage equation, the omitted-variable bias on the exposure coefficient, and the measurement and identifying assumptions behind a role-by-role empirical specification. Section 7 concludes. Proofs and the integrated-baseline extension are collected in the appendices.

2 Environment and the Rebundling Mechanism

The environment begins before automation. There are three tasks, all initially performed by humans. AI is then introduced and performs task 1. Tasks 2 and 3 remain human. The focal pre-AI baseline is the partial partition $\{1,2\}|\{3\}$: tasks 1 and 2 are bundled in one role, while task 3 is separately priced. The post-AI firm chooses whether to buy the surviving human tasks, 2 and 3, from one bundled worker or from two specialists. The pre-AI assignment problem determines whether the partial baseline is supportable and supplies the wage schedule against which the AI-induced redesign is evaluated.

For convenience, Table 1 collects the main notation used throughout the paper.

2.1 Tasks, time and skill

A unit of service requires three tasks, indexed $1, 2, 3$. Final service quality is Cobb-Douglas,

$$q_1^{\alpha_1} q_2^{\alpha_2} q_3^{\alpha_3}, \qquad \alpha_j > 0, \quad \sum_{j=1}^{3} \alpha_j = 1. \tag{1}$$

The Cobb-Douglas form gives a transparent notion of task complementarity. Weak performance on one task lowers the value of the others. The symmetric benchmark sets $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$.

A human worker with skill $s_j > 0$ who devotes time $t_j \geq 0$ to task j produces task quality $q_j = t_j s_j$. A worker assigned to a bundle of human tasks has one unit of time to allocate across the tasks in that bundle; a worker assigned to a singleton specialist role devotes the full unit of time to that role. Let the full pre-AI skill vector be $\tilde{s} = (s_1, s_2, s_3) \in \tilde{S} \subset \mathbb{R}^3_{++}$ and write log skills as $z = \ln s_1$, $x = \ln s_2$, $y = \ln s_3$. Let F denote the distribution of (z, x, y) on a compact set $\tilde{X} \subset \mathbb{R}^3$. The post-AI sorting problem depends on the marginal distribution G of surviving log skills (x, y) on a compact set $X \subset \mathbb{R}^2$. The joint distribution

Table 1: Notation guide

Notation	Meaning	Notation	Meaning
s_j	Human skill in task j	(z, x, y)	Log skills $(\ln s_1, \ln s_2, \ln s_3)$
F, G	Pre-AI joint distribution of (z, x, y) and post-AI marginal distribution of (x, y)	a, d	Average surviving skill and surviving-task comparative advantage in the symmetric benchmark
ρ	Production mode, that is, a partition of tasks into worker roles	$\mathcal{B}_{12\|3}$	Maintained partial pre-AI baseline: $\{12\}\|\{3\}$
$\mathfrak{T}^0_{12\|3}, \mathcal{G}_{12\|3}$	Fixed class of partial-baseline-admissible primitive-equilibrium pairs and their surviving-skill marginals	P, C, U	Partial pre-AI mode, post-AI bundled mode, and post-AI unbundled mode
α_j	Task weight in Cobb–Douglas service quality	$\kappa_B, \kappa, \lambda$	Bundle weight, total surviving-task weight, and task-2 share in surviving weighted skill
$\Omega(B), \Omega_{23}$	Time-spreading factor for bundle B and for the bundled surviving pair	$\Phi^0_\rho, \Phi_C, \Phi_U$	Pre-AI and post-AI coordination/productivity shifters by mode
$M_A(m_A)$	Automated-task productivity factor after AI adoption	A^0_ρ, A_C, A_U	Reduced-form productivity terms for pre-AI modes and post-AI bundled/unbundled technologies
f^0_ρ, f_C, f_U	Pre-AI mode output and post-AI bundled/unbundled output technologies	$p, P(Y), \bar{u}$	Product price, inverse demand, and worker outside option
μ_C, μ_U, η	Assignment measures for bundled workers, specialist teams, and idle workers	η_s, h	Specialist scarcity wedge and post-AI skill-balance threshold

F remains essential for pre/post wage comparisons because the automated-task skill z may be correlated with average surviving skill and comparative advantage.

For any non-empty bundle $B \subseteq \{1, 2, 3\}$, define

$$\kappa_B = \sum_{j \in B} \alpha_j, \qquad \Omega(B) = \prod_{j \in B} \left(\frac{\alpha_j}{\kappa_B} \right)^{\alpha_j}. \tag{2}$$

The factor $\Omega(B)$ is the loss from spreading one worker's time across the tasks in B. For singleton bundles, $\Omega(\{j\}) = 1$. The case used repeatedly below is the surviving-pair bundle, $\Omega_{23} = \Omega(\{2, 3\}) = (\alpha_2/(\alpha_2 + \alpha_3))^{\alpha_2} (\alpha_3/(\alpha_2 + \alpha_3))^{\alpha_3}$. Optimal time allocation within a bundle sets

$$t_j = \frac{\alpha_j}{\kappa_B}, \qquad j \in B, \tag{3}$$

so time shares are pinned down by task weights rather than by the worker's skill vector.

This keeps the assignment geometry sharp and makes the time-spreading loss a technological feature of bundling rather than a worker-specific choice variable. We write $\kappa = \alpha_2 + \alpha_3 = 1 - \alpha_1$, $\lambda = \alpha_2/\kappa$ and $1 - \lambda = \alpha_3/\kappa$, so the weighted average surviving log skill is $\lambda x + (1-\lambda)y$, which collapses to $(x+y)/2$ in the symmetric benchmark.

2.2 Pre-AI job design and post-AI production modes

Before automation, all three tasks are human. A human production mode is a partition $\rho = \{B_1, \ldots, B_{K(\rho)}\}$ of $\{1,2,3\}$, where each bundle B_ℓ is assigned to one worker role. The feasible pre-AI modes are

$$\{123\}, \quad \{12\}|\{3\}, \quad \{13\}|\{2\}, \quad \{23\}|\{1\}, \quad \{1\}|\{2\}|\{3\}.$$

Let $r = (z, x, y)$ and write r_j for the log skill in task j. Given workers $(r^1, \ldots, r^{K(\rho)})$, pre-AI output in mode ρ is

$$f^0_\rho(r^1, \ldots, r^{K(\rho)}) = \Phi^0_\rho \prod_{\ell=1}^{K(\rho)} \left[\Omega(B_\ell) \exp\left\{ \sum_{j \in B_\ell} \alpha_j r^\ell_j \right\} \right]. \tag{4}$$

The factor $\Phi^0_\rho > 0$ captures pre-AI coordination efficiency for the partition ρ: communication across workers, internal organisation, monitoring, and any non-time cost or benefit of grouping tasks together. The pre-AI active job design is the support of the assignment measures over these modes. The focal baseline requires that this support include the partial mode $\{1,2\}|\{3\}$ and exclude the fully integrated mode as the baseline comparison. Whether that restriction is satisfied is determined by the full primitive vector: the distribution F, the time-spreading factors $\Omega(B)$, the coordination factors Φ^0_ρ, demand and the outside option.

In the partial pre-AI mode $P = \{12\}|\{3\}$, one worker performs tasks 1 and 2 and another worker performs task 3. In the symmetric benchmark,

$$f^0_P((z,x), y') = A^0_{12|3} \exp\left\{ \frac{z + x + y'}{3} \right\}, \qquad A^0_{12|3} = \Phi^0_{12|3} \Omega(\{1,2\}). \tag{5}$$

This partial case is the main pre-AI baseline. In the symmetric benchmark its output is

$$f^0_I(z,x,y) = A^0_I \exp\left\{ \frac{z + x + y}{3} \right\}, \qquad A^0_I = \Phi^0_I \Omega(\{1,2,3\}). \tag{6}$$

The remaining pre-AI modes are also not focal baselines, but they remain relevant because they compete for workers and help determine whether the partial baseline can be supported.[1]

After automation, AI performs task 1 with quality $\chi(m_A) > 0$, where m_A indexes AI capability, and the common automated-task productivity factor is $M_A(m_A) = \chi(m_A)^{\alpha_1}$. Tasks 2 and 3 remain human.

The post-AI bundled mode, denoted C, assigns tasks 2 and 3 to one worker. Let $\Phi_C(m_A) > 0$ be the coordination factor for this mode; it captures the efficiency of integrating the AI-performed task with a bundled human role, together with any mode-specific organisational cost not already captured by time spreading. Output from a bundled worker of type (x, y) is

$$f_C(x, y; m_A) = A_C(m_A)\exp\{\alpha_2 x + \alpha_3 y\}, \qquad A_C(m_A) = M_A(m_A)\Phi_C(m_A)\Omega_{23}, \tag{7}$$

which collapses in the symmetric benchmark to $f_C(x, y; m_A) = A_C(m_A)\exp\{(x + y)/3\}$.

The post-AI unbundled mode, denoted U, assigns one worker to task 2 and another to task 3. There is no time-spreading loss because each specialist devotes the full unit of time to one role. Let $\Phi_U(m_A) > 0$ be the coordination factor; it captures human–human handoffs, AI–human interfaces, verification costs, and the organisational overhead or benefit of a more modular production process. Output from a task-2 specialist with log skill x and a task-3 specialist with log skill y' is

$$f_U(x, y'; m_A) = A_U(m_A)\exp\{\alpha_2 x + \alpha_3 y'\}, \qquad A_U(m_A) = M_A(m_A)\Phi_U(m_A), \tag{8}$$

which is $A_U(m_A)\exp\{(x + y')/3\}$ in the symmetric benchmark.

The unbundled technology is a two-worker technology. It is, therefore, misleading to compare f_U with f_C as though the only question were whether two specialists can produce more than one generalist; they often can. The market comparison asks whether the output of the specialist team is large enough to pay for two workers, each of whom has alternative uses in the assignment economy. Take two identical balanced workers in the symmetric benchmark, each with $x = y = a$. If both are used as bundled generalists, they produce two service units with total output $2A_C e^{2a/3}$; if the same two workers are instead used as one specialist team, they produce one service unit with output $A_U e^{2a/3}$. For balanced workers, unbundling dominates only when $A_U > 2A_C$, not merely when $A_U > A_C$.

Two productivity ratios summarise the comparison: the relative coordination effi-

[1] The integrated mode $I = \{123\}$ remains a feasible competing mode in the pre-AI assignment problem. The integrated mode is discussed as an extension at the end of Section 5.4; it is not the baseline maintained for the main transition results.

ciency of unbundling, $\Lambda(m_A) \equiv \Phi_U(m_A)/\Phi_C(m_A)$, and the direct technological trade-off, $\Gamma(m_A) \equiv A_U(m_A)/A_C(m_A) = \Lambda(m_A)/\Omega_{23}$. The common automated-task factor $M_A(m_A)$ cancels from Γ. Holding $\Lambda(m_A)$ fixed, an improvement in the quality of the automated task raises the productivity of both post-AI modes but does not move the role-choice boundary between bundled and specialist work among workers facing the same menu. It can still change aggregate output, product prices, wages and, when outside options bind, the composition of participating workers. Directional redesign of the human job boundary requires AI to change relative coordination efficiency, feasible production modes, adoption, or which tasks remain human.

Final output sells at inverse demand $P(Y)$, where Y is aggregate output, and product-market surplus is $R(Y) = \int_0^Y P(z)\,\mathrm{d}z$. Workers have outside option $\bar{u} \geq 0$. The main closed-form sorting results use the full-participation case $\bar{u} = 0$, or equivalently, analyse the role-choice problem conditional on participation. Positive outside options introduce a participation cutoff because all job values scale with the product price; role-choice comparisons among active jobs are unchanged by the product price, but observed mode shares can change if the set of participating workers changes.

2.3 The role-choice threshold and why an assignment equilibrium is needed

Before stating the formal equilibrium concept, it is useful to preview the post-AI role-choice logic in geometric terms. In the symmetric benchmark, a worker with surviving log skills (x, y) chooses between bundled work, which pays for the average $(x+y)/3$ scaled by the bundled multiplier A_C, and specialist work, which pays the larger of two values proportional to $e^{2x/3}$ and $e^{2y/3}$ scaled by half the specialist multiplier $A_U/2$. Specialist work weakly dominates bundled work when

$$|x-y| \geq h \equiv 3\ln\left(\frac{2A_C}{A_U}\right).$$

The factor of 2 inside the logarithm is the opportunity cost of the second worker on a specialist team: a specialist team uses two workers to produce one service unit, so the specialist comparative-advantage premium must be large enough to compensate two outside options before the labour market selects unbundling. A purely technological comparison would set the boundary at $A_C = A_U$; the labour-market boundary is at $A_U = 2A_C$. The intermediate range $A_C < A_U < 2A_C$ is where engineering efficiency and labour-market sorting disagree. Workers with balanced surviving skills are priced for breadth; workers with strong comparative advantage in one surviving task are priced as specialists. Proposition 2 below makes this precise.

The boundary is endogenous, not imposed. The model does not assume separate wages for task-2 and task-3 skills, which would presume unbundling, or a single wage for a combined task-2/task-3 job, which would presume bundling. A multidimensional worker is indivisible: a worker of type (x, y) can be used once, as a bundled generalist, as a task-2 specialist, as a task-3 specialist, or not at all. Likewise, an unbundled service unit requires a pair of workers, and the feasibility of that team depends on the alternative uses of both. These are assignment constraints, not factor-demand equations. The production-mode formulation treats each feasible job design as a candidate production plan, and the active support of production modes is selected by market prices rather than imposed as an occupational primitive.

3 Production-Mode Equilibrium

A production-mode assignment equilibrium is a price–allocation–payoff triple in which firms choose bundled or specialist production, workers choose roles, and the product market clears, with the active set of job designs selected endogenously by no-profit conditions. Under standard regularity (compact type space, finite worker measure, continuous and weakly decreasing inverse demand), the planner problem has a solution, the product market clears, and the optimum is supported by a continuous worker payoff function $u \in C(X)$ satisfying the no-profit inequalities

$$u(x, y) \geq \bar{u}, \tag{9}$$

$$u(x, y) \geq p f_C(x, y; m_A), \tag{10}$$

$$u(x, y) + u(x', y') \geq p f_U(x, y'; m_A) \tag{11}$$

with equality on the support of idle workers, active bundled workers, and active specialist teams respectively. Details, including the formal regularity assumptions, the planner formulation, and the proof that aggregate output is continuous and weakly increasing in the product price under the symmetric menu, are collected in Appendix B.1. The same existence and decentralisation argument extends to an arbitrary finite task set with any feasible collection of human production modes; Appendix C records this general statement and reads rebundling and unbundling as changes in the support of the active production-mode measures. Supporting payoffs need not be unique; the symmetric core selection introduced below is one transparent member of the set, and the remainder of the paper uses it throughout.

3.1 The symmetric benchmark and the wage menu

The closed-form results use the symmetric benchmark:

$$\alpha_1 = \alpha_2 = \alpha_3 = \frac{1}{3} \quad \text{and} \quad G(B) = G(B^T) \text{ for every Borel } B \subset X, \tag{12}$$

where $B^T = \{(y, x) : (x, y) \in B\}$. Exchangeability means that the population is symmetric across the two surviving human tasks. It does not mean that individual workers are balanced.

The symmetric job values are

$$W_C(x, y; p) = pA_C \exp\left\{\frac{x+y}{3}\right\}, \tag{13}$$

$$W_2(x; p) = p\frac{A_U}{2} \exp\left\{\frac{2x}{3}\right\}, \tag{14}$$

$$W_3(y; p) = p\frac{A_U}{2} \exp\left\{\frac{2y}{3}\right\}. \tag{15}$$

The corresponding price-free values are

$$R_C(x, y) = A_C \exp\left\{\frac{x+y}{3}\right\}, \tag{16}$$

$$R_2(x) = \frac{A_U}{2} \exp\left\{\frac{2x}{3}\right\}, \tag{17}$$

$$R_3(y) = \frac{A_U}{2} \exp\left\{\frac{2y}{3}\right\}. \tag{18}$$

Define the worker payoff menu

$$u_p(x, y) = \max\{\bar{u}, W_C(x, y; p), W_2(x; p), W_3(y; p)\}. \tag{19}$$

Workers choose any job attaining the maximum. If W_2 is chosen, the worker enters the task-2 specialist pool; if W_3 is chosen, the worker enters the task-3 specialist pool.[2]

[2] On tie sets, the assignment is represented by a measurable tie-breaking kernel that is invariant under transposition: the mass assigned from type (x, y) to the task-2 role is mirrored by the mass assigned from type (y, x) to the task-3 role, and similarly for the other roles. If tie sets have zero G-measure, the choice of tie-breaking rule is irrelevant.

Proposition 1 (Supporting symmetric equilibrium menu)**.** *Fix $p > 0$ and suppose (12) holds. Assign workers to jobs that attain the maximum in (19) using a transposition-invariant measurable tie-breaking kernel. Match task-2 and task-3 specialists by equal quantiles of their role-relevant log skills. Then no firm can earn a positive profit using either production mode. Active bundled firms and active specialist teams earn zero profit.*

The proofs of all propositions are in Appendix B. The proposition gives a convenient symmetric wage menu, not a uniqueness result. The exact envelope below—its kinks, slopes and role-by-role coefficients—comes from that symmetric core selection. The broader economics is more robust: bundled jobs reward average surviving skill, specialist jobs reward comparative advantage, and scarcity raises pay in the scarce specialist role. Thus, whenever a closed-form wage coefficient is reported below, the coefficient should be read as a coefficient under this selected symmetric payoff menu. AI quality leaves the role-choice threshold unchanged only when it scales both post-AI modes through the common factor $M_A(m_A)$. If it also changes Φ_U/Φ_C, feasible modes, adoption, or which tasks remain human, then it can change job design through those channels. With positive outside options, it can also change observed mode shares through participation even when the threshold for participating workers is unchanged.

4 Skill-Balance Sorting and the Return to Comparative Advantage

This section characterises post-AI sorting in the full-participation symmetric benchmark. It derives the threshold that separates bundled from specialist jobs and the implied wage schedule. The results are stated for an arbitrary surviving-skill distribution G; Section 5 restricts G to distributions that arise from partial-baseline-admissible pre-AI equilibria.

4.1 The job-design threshold

Compare the bundled value and the task-2 specialist value. From (13) and (14), task-2 specialisation weakly dominates bundled work if and only if

$$p\frac{A_U}{2}e^{2x/3} \geq pA_C e^{(x+y)/3}. \tag{20}$$

Cancelling p and rearranging gives the role-choice threshold

$$x - y \geq h, \qquad h \equiv 3 \ln\Bigg(\underbrace{2}_{\substack{\text{second-}\\ \text{worker cost}}} \cdot \underbrace{\frac{\Omega_{23}}{\Lambda(m_A)}}_{\substack{\text{bundling cost}\\ \text{vs. specialist}\\ \text{coordination}}} \Bigg). \tag{21}$$

Symmetrically, task-3 specialisation dominates bundled work when $y - x \geq h$.

The threshold has a direct interpretation. Unbundling uses two workers, so a worker must have a sufficiently strong comparative advantage before the selected specialist wage can dominate the bundled wage. The factor of 2 is the opportunity cost of the second worker in the unbundled team. A rise in the relative coordination efficiency of unbundling, $\Lambda(m_A)$, lowers h. A more severe time-spreading loss from bundling, represented by a lower Ω_{23}, also lowers h. Holding $\Lambda(m_A)$ fixed, an improvement in the quality of the automated task raises A_C and A_U in the same proportion and leaves h unchanged.

The factor of 2 inside the logarithm distinguishes the labour-market threshold from the purely technological comparison between bundled and unbundled production. A technological ranking would compare A_C with A_U directly. The labour-market threshold also requires the specialist team to compensate two workers, each with alternative uses in the assignment economy. Hence specialist production may be technologically more productive per service unit, in the sense that $A_U > A_C$, while balanced workers are still assigned to bundled roles whenever $A_U < 2A_C$. The intermediate region $A_C < A_U < 2A_C$ is precisely where engineering efficiency and labour-market sorting disagree.

4.2 Equilibrium regions in skill space

Let

$$D_{\max} = \sup\{|x - y| : (x, y) \in \operatorname{supp} G\} \tag{22}$$

be the largest log-skill difference in the support of the worker distribution.

Proposition 2 (Skill-balance sorting)**.** *Suppose $\bar{u} = 0$, (12) holds, and role-choice tie sets have G-measure zero. Let h be defined by (21).*

(i) If $h \leq 0$, all workers are assigned to specialist jobs up to null tie sets.

(ii) If $0 < h < D_{\max}$, workers with $|x - y| < h$ are assigned to bundled generalist jobs; workers with $x - y > h$ are assigned to task-2 specialist jobs; workers with $y - x > h$ are assigned to task-3 specialist jobs.

(iii) If $h \geq D_{\max}$, all workers are assigned to bundled generalist jobs up to null tie sets.

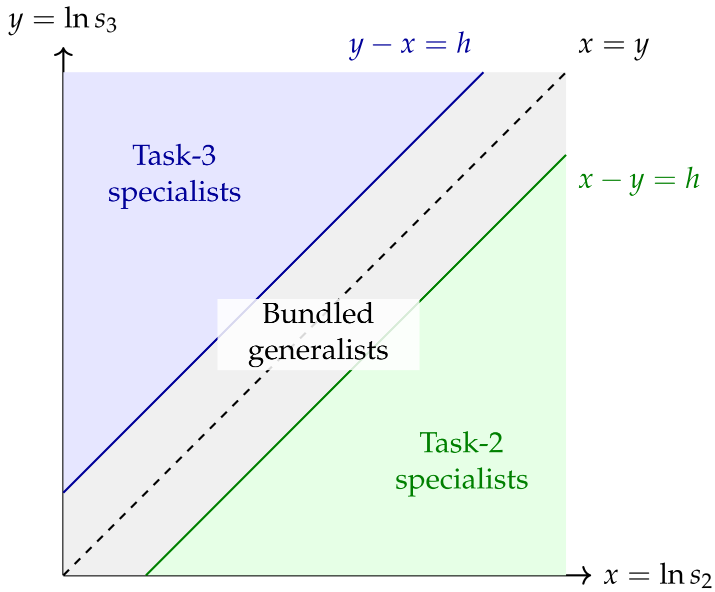

The specialist regions expand as the threshold $h = 3\ln(2A_C/A_U)$ falls.

Figure 1: Skill-balance sorting in the intermediate case. Workers near the diagonal have balanced surviving skills and become bundled generalists. Workers far below the diagonal are relatively strong at task 2 and become task-2 specialists. Workers far above the diagonal become task-3 specialists.

Thus, the post-AI job design is governed by the interaction of technology and the surviving skill distribution. The same coordination technology can generate pure bundling, coexistence, or pure specialisation depending on the support of comparative advantage. For a pre/post prediction, however, the relevant support is not arbitrary; it must be the marginal support inherited from a pre-AI economy in which the partial baseline $\{1,2\}|\{3\}$ was actually active.

Let

$$M_U(h) = G\big(\{(x,y) : |x-y| > h\}\big) \tag{23}$$

be the mass of workers assigned to specialist jobs in the full-participation symmetric equilibrium. This is a worker-mass measure: each unbundled service unit uses two specialist workers.

Proposition 3 (Specialist employment)**.** *Suppose Proposition 2 applies. Specialist worker mass is weakly decreasing in the threshold h. Since*

$$h = 3\ln(2\Omega_{23}) - 3\ln\Lambda(m_A),$$

specialist worker mass is weakly increasing in the relative coordination efficiency of unbundling,

$\Lambda(m_A)$, and weakly increasing in the severity of the time-spreading cost. If $\partial \ln \Lambda(m_A)/\partial m_A > 0$, so that AI capability makes unbundled production relatively easier to coordinate, specialist worker mass is weakly increasing in m_A. If AI capability shifts Φ_C and Φ_U proportionally, then $\Lambda(m_A)$ is unchanged and AI quality does not move the post-AI sorting threshold.

Specialist worker mass rises with skill dispersion: when $d = (x-y)/2$ is uniform on $[-\sigma, \sigma]$ and independent of $a = (x+y)/2$,

$$M_U(h) = 1 - \frac{h}{2\sigma}, \qquad 0 < h < 2\sigma, \tag{24}$$

so $\partial M_U/\partial\sigma > 0$ at any fixed interior h. The derivation appears in the proof.

4.3 The selected wage envelope

Let $\ell(x,y) = \ln u_p(x,y)$ denote log worker payoff when $\bar{u} = 0$, and write $a = (x+y)/2$ and $d = (x-y)/2$. Combining (13)–(15) gives the selected wage envelope as the upper envelope of a flat bundled schedule and a V-shaped specialist schedule:

$$\ell(a,d) = \ln p + \frac{2a}{3} + \max\Big\{ \underbrace{\ln A_C}_{\text{bundled}}, \underbrace{\ln\left(\tfrac{A_U}{2}\right) + \tfrac{2|d|}{3}}_{\text{best specialist}} \Big\}. \tag{25}$$

The comparison depends on $|d|$, not the sign of d, because either task can be the worker's comparative advantage. The kink and the slope are properties of the symmetric payoff selection; the more robust prediction is that specialist production introduces a return to role-specific comparative advantage while bundled production prices average surviving skill alone. Defining the imbalance threshold

$$d^* \equiv \frac{h}{2} = \frac{3}{2}\ln\left(\frac{2A_C}{A_U}\right), \tag{26}$$

(25) can be rewritten in the kinked form

$$\ell(a,d) = \underbrace{\ln p + \ln A_C}_{\text{bundled level}} + \underbrace{\tfrac{2a}{3}}_{\text{breadth premium}} + \underbrace{\max\{0, \tfrac{2}{3}(|d| - d^*)\}}_{\text{comparative-advantage premium}}. \tag{27}$$

The third term is the contribution of specialist sorting: it is zero in the bundled region and rises linearly in $|d|$ once the worker's comparative advantage is large enough to justify the second worker on a specialist team.

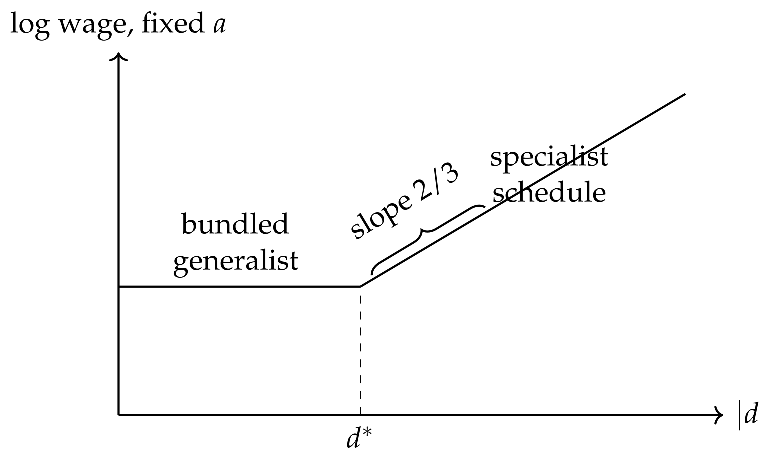

Figure 2: The selected return to skill imbalance when $d^* > 0$. The wage envelope is flat in $|d|$ in the bundled region and increasing in $|d|$ once specialist work is selected.

Proposition 4 (Return to skill imbalance under the symmetric payoff selection)**.** *In the full-participation symmetric equilibrium, holding average surviving log skill a fixed, selected log wages are flat in $|d|$ for $|d| < d^*$ and increase with slope $2/3$ in $|d|$ for $|d| > d^*$ whenever $d^* > 0$. If $d^* \leq 0$, specialist work dominates bundled work for all non-tie types, and selected log wages increase with slope $2/3$ in $|d|$ throughout the support. In either case, workers assigned to specialist jobs receive a positive selected return to role-specific comparative advantage.*

A fixed-bundle prediction assumes that the surviving human tasks remain jointly priced after automation. In the symmetric benchmark, the same-price fixed-bundle prediction is

$$\ell^{FB}(a) = \ln p + \ln A_C + \frac{2a}{3}. \tag{28}$$

It depends on average surviving skill but not on skill imbalance.

Corollary 1 (Same-price fixed-bundle prediction error)**.** *Under the symmetric payoff selection, the true-minus-fixed-bundle log wage error at a common product price is*

$$\ell(a,d) - \ell^{FB}(a) = \max\left\{0, \frac{2}{3}(|d| - d^*)\right\}. \tag{29}$$

A fixed-bundle earnings function omits the gains to comparative advantage among workers assigned to specialist jobs.

This is a same-price envelope comparison, not a full equilibrium counterfactual. A true fixed-bundle counterfactual generally has a different aggregate output and therefore a different product price. The robust point is narrower: a fixed-bundle earnings function omits the skill-imbalance margin created by post-AI specialist production. It will tend to miss workers whose wage gains come not from being good at the automated task, but from being unusually good at one of the tasks that remain human.

5 The Partial Pre-AI Baseline and Rebundling-as-Revaluation

The post-AI sorting theorem describes how an economy with surviving-skill distribution G assigns workers to bundled and specialist roles. The pre/post question is narrower. The maintained starting point is the partial pre-AI mode $\{1,2\}|\{3\}$: task 1, later automated, is bundled with surviving task 2, while surviving task 3 is separately priced. The post-AI economy then either rebundles task 2 with task 3, or unbundles the old task-1/task-2 role and leaves the surviving tasks in specialist roles.

The partial baseline is the right starting point because it is the minimal pre-AI configuration in which automation of task 1 can simultaneously do three things: remove the human return to the automated-task skill z; change the role of task 2 (from being bundled with the automated task to being either re-bundled with task 3 or priced separately); and make a previously unpriced surviving skill newly valuable. This last channel is the distinctive economic mechanism. A worker initially in the $\{1,2\}$ role was never paid for task-3 skill y before AI; under post-AI specialist work, y may become the dominant component of their wage. A worker initially in the $\{3\}$ role was never paid for task-2 skill x before AI; under post-AI rebundling or task-2 specialisation, x may become newly priced. The integrated baseline $\{1,2,3\}$ does not deliver this rebundling-as-revaluation effect, because under integration every surviving skill was already in the pre-AI bundle. The integrated case is recorded as an extension in Appendix A.3.

5.1 Partial-baseline-admissible primitives

Assumption 1 (Partial pre-AI baseline)**.** A primitive-equilibrium pair (θ^0, e^0) is partial-baseline-admissible if e^0 is a selected pre-AI equilibrium for the primitive vector $\theta^0 = (F, \{\Phi^0_\rho\}_{\rho\in\mathcal{P}^0}, P, \bar{u})$, in the sense of Proposition A.1, and the selected mode masses $m^0_\rho = \nu^0_\rho(\mathbb{R})$ satisfy

$$m^0_{12|3} > 0, \qquad m^0_{123} = 0. \tag{30}$$

The first condition says that the initial selected equilibrium involves a role bundling task 1 with task 2 and a separate role for task 3. The second condition excludes the fully integrated mode as the maintained baseline. Other non-focal modes may be active in the same pre-AI economy. The wage comparisons below are for workers assigned to the focal partial mode; a pure partial-baseline economy is the special case in which no other pre-AI mode has positive mass.

The restriction is imposed on the full primitive-equilibrium pair, not on the surviving-skill marginal alone. This distinction matters because the selected pre-AI assignment may not be unique. A definition that first fixes an arbitrary marginal G and then asks whether there exist some coordination factors $\{\Phi^0_\rho\}$ and some selected equilibrium capable of rationalising the desired baseline would be too weak: the coordination factors and the selected pre-AI support are part of the observed or calibrated pre-AI object, not free instruments to be chosen after the post-AI marginal has been specified. The fixed admissible class $\mathfrak{T}^0_{12|3}$ of primitive-equilibrium pairs satisfying Assumption 1, the surviving-skill image $\mathcal{G}_{12|3} = \{\pi_{23\#}F : ((F,\cdot),e^0) \in \mathfrak{T}^0_{12|3}\}$, and the closed-form selected pre-AI menu are recorded in Appendix A. The restriction is not empty: Proposition 5 constructs a compact atomless economy with an exchangeable surviving-skill marginal in which the partial mode is active and the integrated mode is inactive.

Proposition 5 (Non-emptiness of the partial-baseline-admissible class)**.** *There exist compactly supported primitives and a selected pre-AI equilibrium satisfying Assumption 1. The construction can be chosen with symmetric task weights, zero outside option, a pure partial pre-AI support* $\mathcal{M}^0 = \{\{1,2\}|\{3\}\}$*, and an exchangeable surviving-skill marginal* $\pi_{23\#}F$*.*

The admissibility restriction is therefore substantive rather than empty: it rules in a non-vacuous set of pre-AI primitive-equilibrium pairs and rules out arbitrary surviving-skill marginals that cannot be generated by such a pre-AI support.

5.2 Restricted post-AI continuation set

After AI, task 1 is automated and the relevant worker distribution is the marginal distribution of surviving skills,

$$G = \pi_{23\#}F. \tag{31}$$

For the transition analysis, G is restricted to $\mathcal{G}_{12|3}$. In the symmetric benchmark the post-AI threshold is

$$h = 3\ln\left(\frac{2\Omega_{23}}{\Lambda(m_A)}\right), \qquad \Lambda(m_A) = \frac{\Phi_U(m_A)}{\Phi_C(m_A)}. \tag{32}$$

Let

$$T_2(h) = G\big(\{(x,y) : x - y > h\}\big), \tag{33}$$
$$T_3(h) = G\big(\{(x,y) : y - x > h\}\big), \tag{34}$$
$$M_C(h) = G\big(\{(x,y) : |x - y| < h\}\big) \tag{35}$$

be the two specialist-tail masses and the central bundled mass. In the symmetric benchmark, the specialist mode U is active when both specialist tails have positive mass, and the bundled mode C is active when $M_C(h) > 0$.

5.3 Pre/post job-design responses on the restricted domain

Definition 1 (Partial-baseline job-design responses)**.** Under Assumption 1:

(i) *Rebundling* occurs when $m^0_{12|3} > 0$ and the post-AI bundled mode $C = \{2,3\}$ is active. Task 3 was separately priced before AI but is bought together with task 2 after task 1 is automated.

(ii) *Unbundling* occurs when $m^0_{12|3} > 0$ and the post-AI specialist mode $U = \{2\}|\{3\}$ is active. The old $\{1,2\}$ role is broken by automation, and task 2 is priced as a surviving specialist role rather than being rebundled with task 3.

The terminology is relational. In the maintained partial baseline, rebundling means that task 2 moves from its old bundle with the automated task into a new human bundle with task 3. Unbundling means that automation breaks the old task-1/task-2 role and leaves task 2 and task 3 separately priced after AI.[3]

Proposition 6 (Restricted pre/post job-design map)**.** *Suppose the full-participation symmetric benchmark applies, role-choice tie sets have zero mass, and the pre-AI primitive-equilibrium pair is partial-baseline-admissible. Let $G = \pi_{23\#}F \in \mathcal{G}_{12|3}$, and let h and $D_{\max}$ be defined by (32) and (22).*

(i) If $h \leq 0$, then the only active post-AI mode is U. The transition is unbundling.

(ii) If $0 < h < D_{\max}$, then U is active if and only if $T_2(h) > 0$ and $T_3(h) > 0$, while C is active if and only if $M_C(h) > 0$. The transition includes unbundling exactly when U is active and rebundling exactly when C is active.

(iii) If $h \geq D_{\max}$, then the only active post-AI mode is C. The transition is rebundling.

[3] In the integrated extension at the end of Section 5.4, the same word unbundling has the more familiar meaning that tasks 2 and 3 were jointly priced before AI and separately priced after AI.

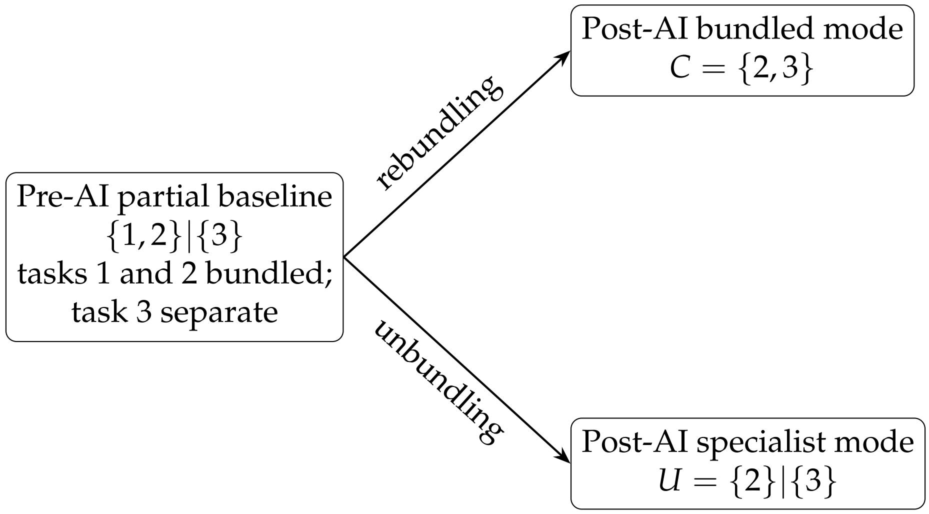

Figure 3: Job-design responses from the partial pre-AI baseline. The maintained starting point is $\{1,2\}|\{3\}$; post-AI bundling is therefore rebundling, while post-AI specialisation is the unbundled outcome.

The proposition is the restricted-domain version of the post-AI sorting theorem. The unrestricted theorem asks what would happen for an arbitrary surviving-skill distribution. The transition theorem asks what can happen for surviving-skill distributions inherited from pre-AI economies in which task 1 and task 2 were bundled while task 3 was separate. See Figure 3.

5.4 Pre/post wage impacts

The worker-level effect of automation is the difference between the post-AI payoff and the pre-AI payoff,

$$\Delta(z,x,y) = \ln u^1(x,y) - \ln u^0(z,x,y), \tag{36}$$

where u^0 is determined by a partial-baseline-admissible pre-AI assignment equilibrium and u^1 by the post-AI assignment equilibrium. Under the selected pre-AI menu in Proposition A.1, a worker assigned before AI to block B in mode ρ has selected log payoff

$$\ell^0_{\rho,B}(r) = \ln p^0 + \ln A^0_\rho + \ln \kappa_B + \bar{r}_B + \xi_{\rho,B}. \tag{37}$$

If the same worker is assigned after AI to the bundled surviving-task mode C, then

$$\Delta_{C|\rho,B}(r) = \ln\left(\frac{p^1}{p^0}\right) + \ln A^1_C - \left(\ln A^0_\rho + \ln \kappa_B\right) + \frac{x+y}{3} - \bar{r}_B - \xi_{\rho,B}. \tag{38}$$

If the worker is assigned after AI to the task-2 specialist role, then

$$\Delta_{2|\rho,B}(r) = \ln\left(\frac{p^1}{p^0}\right) + \ln\left(\frac{A_U^1}{2}\right) - \left(\ln A_\rho^0 + \ln \kappa_B\right) + \frac{2x}{3} - \bar{r}_B - \xi_{\rho,B}. \tag{39}$$

If the worker is assigned after AI to the task-3 specialist role, then

$$\Delta_{3|\rho,B}(r) = \ln\left(\frac{p^1}{p^0}\right) + \ln\left(\frac{A_U^1}{2}\right) - \left(\ln A_\rho^0 + \ln \kappa_B\right) + \frac{2y}{3} - \bar{r}_B - \xi_{\rho,B}. \tag{40}$$

Thus, the general incidence formula subtracts the role-specific pre-AI skill index that the worker was paid for. The automated-task skill matters because it was part of the focal pre-AI $\{1,2\}$ role; the surviving-task imbalance matters according to the post-AI job selected.

In the symmetric partial baseline, a worker in the bundled task-1/task-2 role is paid for the average of their automated-task skill and their task-2 skill, while a worker in the separate task-3 role is paid for task-3 skill alone:

$$\ell^0_{12|3,12}(z,x) = \underbrace{\ln p^0 + \ln A^0_{12|3} + \ln\tfrac{2}{3} - \tfrac{\delta_{12|3}}{3}}_{\text{constant for the role}} + \underbrace{\tfrac{z+x}{2}}_{\substack{\text{average skill on the}\\ \text{bundled tasks 1 and 2}}}, \tag{41}$$

$$\ell^0_{12|3,3}(y) = \underbrace{\ln p^0 + \ln A^0_{12|3} + \ln\tfrac{1}{3} + \tfrac{2\delta_{12|3}}{3}}_{\text{constant for the role}} + \underbrace{y}_{\substack{\text{task-3 skill,}\\ \text{the only priced skill}}}. \tag{42}$$

Here $\delta_{12|3}$ is the selected pre-AI matching wedge for the partial mode, defined in Proposition A.1. Subtracting these pre-AI payoffs from the post-AI selected schedules (13)–(15) gives the wage-impact decomposition below. Throughout, we write $a = (x+y)/2$ and $d = (x-y)/2$ for average surviving skill and signed surviving-task imbalance.

Proposition 7 (Wage-impact decomposition from a partial pre-AI baseline)**.** *Suppose the symmetric full-participation benchmark applies and the worker was initially assigned to the partial pre-AI mode $\{1,2\}|\{3\}$. Let $A_C^1 = A_C(m_A)$ and $A_U^1 = A_U(m_A)$.*

Each wage-impact expression below has the same four-part structure: a product-price term, a productivity term comparing post-AI to pre-AI organisational multipliers, the loss of the human return to the automated task (zero for workers who were never paid for it pre-AI), and a redesign-channel term in the surviving-skill coordinates (a,d).

For a worker initially in the $\{1,2\}$ role, post-AI assignment to C, to the task-2 specialist role, *and to the task-3 specialist role gives*

$$\Delta_{C|12} = \underbrace{\ln\left(\frac{p^1}{p^0}\right)}_{price} + \underbrace{\ln\left(\frac{A_C^1}{(2/3)A_{12|3}^0}\right) + \frac{\delta_{12|3}}{3}}_{productivity} \underbrace{-\frac{z}{2}}_{\substack{\textit{lost auto-}\\ \textit{task return}}} + \underbrace{\frac{a}{6} - \frac{d}{2}}_{\textit{redesign channel}}, \tag{43}$$

$$\Delta_{2|12} = \underbrace{\ln\left(\frac{p^1}{p^0}\right)}_{price} + \underbrace{\ln\left(\frac{A_U^1/2}{(2/3)A_{12|3}^0}\right) + \frac{\delta_{12|3}}{3}}_{productivity} \underbrace{-\frac{z}{2}}_{\substack{\textit{lost auto-}\\ \textit{task return}}} + \underbrace{\frac{a+d}{6}}_{\textit{redesign channel}}, \tag{44}$$

$$\Delta_{3|12} = \underbrace{\ln\left(\frac{p^1}{p^0}\right)}_{price} + \underbrace{\ln\left(\frac{A_U^1/2}{(2/3)A_{12|3}^0}\right) + \frac{\delta_{12|3}}{3}}_{productivity} \underbrace{-\frac{z}{2}}_{\substack{\textit{lost auto-}\\ \textit{task return}}} + \underbrace{\frac{a}{6} - \frac{7d}{6}}_{\textit{redesign channel}}. \tag{45}$$

For a worker initially in the $\{3\}$ *role, post-AI assignment to C, to the task-2 specialist role, and to the task-3 specialist role gives*

$$\Delta_{C|3} = \underbrace{\ln\left(\frac{p^1}{p^0}\right)}_{price} + \underbrace{\ln\left(\frac{A_C^1}{(1/3)A_{12|3}^0}\right) - \frac{2\delta_{12|3}}{3}}_{productivity} + \underbrace{0}_{\substack{\textit{no auto-}\\ \textit{task return}}} + \underbrace{-\frac{a}{3} + d}_{\textit{redesign channel}}, \tag{46}$$

$$\Delta_{2|3} = \underbrace{\ln\left(\frac{p^1}{p^0}\right)}_{price} + \underbrace{\ln\left(\frac{A_U^1/2}{(1/3)A_{12|3}^0}\right) - \frac{2\delta_{12|3}}{3}}_{productivity} + \underbrace{0}_{\substack{\textit{no auto-}\\ \textit{task return}}} + \underbrace{-\frac{a}{3} + \frac{5d}{3}}_{\textit{redesign channel}}, \tag{47}$$

$$\Delta_{3|3} = \underbrace{\ln\left(\frac{p^1}{p^0}\right)}_{price} + \underbrace{\ln\left(\frac{A_U^1/2}{(1/3)A_{12|3}^0}\right) - \frac{2\delta_{12|3}}{3}}_{productivity} + \underbrace{0}_{\substack{\textit{no auto-}\\ \textit{task return}}} + \underbrace{-\frac{a}{3} + \frac{d}{3}}_{\textit{redesign channel}}. \tag{48}$$

The proposition is easiest to read through the three incidence channels summarised in Table 2: automated-task skill z, average surviving skill a, and signed surviving-task imbalance d. The loss of z is a pure automation effect for workers who began in $\{1,2\}$, because task 1 was priced before AI and disappears afterwards. Rebundling rewards high a, while specialist assignment rewards d in the direction of the active specialist role. What is distinctive about the partial baseline is the pre-AI benchmark being subtracted. Workers who began in $\{1,2\}$ were previously paid for $(z+x)/2$, not for y, so moving into task-3 specialist work can make $d < 0$ newly valuable. Workers who began in $\{3\}$ were previously paid for y, not for (z,x), so their post-AI gains depend on whether task 2 becomes newly rewarded through rebundling or specialist pricing.

Among the six cases, one is particularly stark. A worker who began in the $\{1,2\}$ role and is post-AI assigned to task-3 specialist work has wage-impact coefficient $-7/6$ on signed surviving-task imbalance d (equation (45)). Three forces combine. The worker loses the human return to the automated task, contributing $-z/2$. Post-AI specialist pricing rewards task-3 skill y, contributing $+2y/3$, even though the worker was never paid for y under the partial baseline. The signed imbalance $d = (x - y)/2$, which was irrelevant to the pre-AI wage, now appears with a coefficient that combines the post-AI task-3 specialist return ($-2d/3$) with the subtraction of the pre-AI bundled return to task-2 skill ($-d/2$ from the $x/2$ component of $(z + x)/2$); the sum is $-7d/6$. In the symmetric sorting rule this transition is selected for workers with $d < -d^*$, so the model assigns relatively task-3-strong workers to task-3 specialist work. No fixed-bundle earnings function can recover this, because the relevant skill (y) was not in the pre-AI bundle and the relevant role transition was not in the pre-AI choice set. The same logic produces the coefficient $5/3$ on d for a worker who began in the $\{3\}$ role and becomes a task-2 specialist (equation (47)): in that case, the previously unpriced skill x becomes the dominant component of the wage. Relative scarcity of one specialist side tilts the specialist log-payoff terms in the obvious direction; see Appendix B.3 for the constant-wedge family that supports this comparative static.

Table 2: Worker skill profiles and wage incidence from the partial baseline

Worker skill profile	Favoured when	Exposed to loss when
High automated-task skill z among workers initially in $\{1, 2\}$	Task 1 was part of the supported pre-AI role	Task 1 is automated and the human return to z drops out of the post-AI wage schedule
High average surviving skill a	Surviving tasks are rebundled after AI	Post-AI specialist pricing rewards role-specific skill more than breadth across surviving tasks
High signed imbalance toward task 2, $d > 0$	Task-2 specialist roles are active, or task 2 is newly valued for workers initially in the $\{3\}$ role	Rebundling rewards balanced breadth rather than specialist task-2 advantage
High signed imbalance toward task 3, $d < 0$	Workers initially in the $\{1, 2\}$ role move into task-3 specialist work, making previously unpriced task-3 skill valuable	Task-3 skill remains outside the worker's post-AI role
Imbalance tilted toward a scarce surviving role	The corresponding specialist role receives a positive scarcity wedge	Imbalance is tilted toward the relatively abundant role

6 Empirical Implications: Bias from Fixed Task Bundles

Autor and Thompson (2025) study how automation changes wages and employment by changing the expertise content of the tasks that remain in an occupation. Their empirical object is a fixed-occupation, fixed-bundle object: once a task is removed, the remaining tasks are still priced within the same occupational bundle. The assignment model here gives a different estimand when automation also changes the job boundary. The fixed-bundle estimate remains a valid estimate of a fixed-bundle effect; it is not, by itself, an estimate of the total worker-level effect when the priced human bundle changes.

The distinction is sharp in the partial baseline. Before AI, task 1 is bundled with task 2, while task 3 is separately priced. A fixed-bundle design treats automation of task 1 as changing the expertise requirement of the original role. In the model, however, automation may change the human task partition from

$$\{1,2\}|\{3\} \quad \text{to either} \quad C = \{2,3\} \quad \text{or} \quad U = \{2\}|\{3\}.$$

The post-AI wage schedule may therefore price average surviving skill, signed comparative advantage, or a skill that was outside the worker's original pre-AI role. A fixed-bundle regression omits that redesign component.

Let ΔXPT_o^{AT} and $\Delta TASK_o^{AT}$ denote occupation-level measures of expertise change and task-quantity change. Let $B^0_{12|3,o}$ indicate that occupation o satisfies the partial pre-AI baseline: task 1 and task 2 were jointly demanded before AI, while task 3 was separately priced. Let $R^0_{12,i}$ and $R^0_{3,i}$ indicate worker i's pre-AI role inside that partial baseline. Let $I^1_{C,io}$, $I^1_{2,io}$ and $I^1_{3,io}$ indicate the worker's post-AI assignment to the rebundled role, task-2 specialist role and task-3 specialist role. In the symmetric full-participation benchmark these indicators are generated by the threshold rule

$$I^1_{2,io} = \mathbf{1}\{d_i > d_o^*\}, \qquad I^1_{3,io} = \mathbf{1}\{d_i < -d_o^*\}, \qquad I^1_{C,io} = \mathbf{1}\{|d_i| < d_o^*\},$$

where $d_o^* = h_o/2$. In empirical work the indicators can instead be measured from observed post-AI jobs or vacancies. For worker skills write

$$a_i = \frac{x_{i2} + x_{i3}}{2}, \qquad d_i = \frac{x_{i2} - x_{i3}}{2},$$

where a_i is average surviving skill and d_i is signed surviving-task imbalance.

A fixed-bundle wage regression has the form

$$\Delta \ln w_{io} = \beta_X \Delta XPT_o^{AT} + \beta_N \Delta TASK_o^{AT} + X_i'\theta + Z_o'\delta + e_{io}. \tag{49}$$

In the model, the structural wage change for workers in the partial baseline is piecewise affine in (z_i, a_i, d_i) and in the initial and post-AI roles. It can be written as

$$\Delta \ln w_{io} = B^0_{12|3,o} \sum_{r\in\{12,3\}} \sum_{k\in\{C,2,3\}} R^0_{r,i} I^1_{k,io} \left[K_{rk,o} + b_{rk,z} z_i + b_{rk,a} a_i + b_{rk,d} d_i\right] + \varepsilon_{io}, \tag{50}$$

where the constants $K_{rk,o}$ collect product-price effects, productivity terms and selected matching wedges. The skill coefficients implied by the symmetric partial-baseline decomposition are shown in Table 3.

Table 3: Role-specific skill coefficients in the partial-baseline wage change

Initial and post-AI role	Coefficient on z_i	Coefficient on a_i	Coefficient on d_i
$C \mid 12$	$-1/2$	$1/6$	$-1/2$
$2 \mid 12$	$-1/2$	$1/6$	$1/6$
$3 \mid 12$	$-1/2$	$1/6$	$-7/6$
$C \mid 3$	0	$-1/3$	1
$2 \mid 3$	0	$-1/3$	$5/3$
$3 \mid 3$	0	$-1/3$	$1/3$

The omitted redesign term is the part of (50) that is not spanned by the fixed-bundle variables in (49). Write the population equation as

$$\Delta \ln w_{io} = \beta_X \Delta XPT^{AT}_o + \beta_N \Delta TASK^{AT}_o + X_i'\theta + Z_o'\delta + \mathcal{R}_{io} + \varepsilon_{io}. \tag{51}$$

Let Q_{io} collect the included controls other than ΔXPT^{AT}_o, and let tildes denote population residuals after linear projection on Q_{io}. Provided exposure has positive residual variance and is uncorrelated with the structural error, the probability limit of the fixed-bundle coefficient decomposes into the structural object and an omitted-redesign bias:

$$\operatorname{plim} \hat{\beta}^{AT}_X = \underbrace{\beta_X}_{\text{structural exposure effect}} + \underbrace{\frac{\operatorname{Cov}\left(\widetilde{\Delta XPT^{AT}_o}, \widetilde{\mathcal{R}}_{io}\right)}{\operatorname{Var}\left(\widetilde{\Delta XPT^{AT}_o}\right)}}_{\text{omitted-redesign bias}}. \tag{52}$$

The sign of the bias is not pinned down by exposure alone. Equation (52) signs the bias only after specifying how exposure covaries with the entire omitted redesign term. Directional statements therefore require additional restrictions on the post-AI partition, the affected worker groups, the distribution of signed skill imbalance, product-price changes and participation.

A model-faithful empirical augmentation adds a redesign block to the fixed-bundle specification:

$$\Delta \ln w_{io} = \underbrace{\alpha + \beta_X \Delta XPT_o^{AT} + \beta_N \Delta TASK_o^{AT} + X_i'\theta + Z_o'\delta}_{\text{fixed-bundle terms}}$$
$$+ \underbrace{A_o B^0_{12|3,o} \sum_{\substack{r\in\{12,3\}\\ k\in\{C,2,3\}}} R^0_{r,i} I^1_{k,io} \left[\gamma_{rk} + \gamma_{rk,z} z_i + \gamma_{rk,a} a_i + \gamma_{rk,d} d_i\right]}_{\text{role-by-role redesign block}} + \varepsilon_{io}, \qquad (53)$$

where A_o is AI exposure or adoption. The restricted version imposes the coefficient pattern in Table 3; the flexible version estimates the interactions and tests whether they match the model. This specification uses the signed imbalance d_i and role indicators. Replacing those terms by a single $|d_i|$ regressor is only valid for the selected post-AI wage envelope when the worker's pre-AI role is not being subtracted. It is not valid for the partial-baseline wage changes, because moving from $\{1,2\}$ to task-2 specialisation and moving from $\{1,2\}$ to task-3 specialisation have different signed-d_i coefficients.

A fixed-bundle post-AI schedule at a common product price omits the specialist premium $\max\{0, \frac{2}{3}(|d_i| - d_o^*)\}$, which is positive for workers who become specialists and zero for workers who remain in the bundled region under the selected post-AI envelope. The partial-baseline wage change subtracts the worker's pre-AI role-specific wage, so the omitted term in (52) is more general than this non-negative premium. Whether the fixed-bundle estimate overstates or understates the wage effect for any given group depends on the role transition, product-price changes, participation changes, displacement and the correlation between exposure and the post-AI partition.

A stylised magnitude calculation. A simple numerical illustration shows that the bias in (52) can be large enough, under the maintained covariance structure, to flip the sign of the AI-exposure coefficient. Take the symmetric benchmark, set $d_o^* = 0.1$ so the role-choice threshold is interior, and consider a stylised population in which surviving-task imbalance d_i is symmetrically distributed about zero with standard deviation $\sigma_d = 0.15$. Within $B^0_{12|3,o}$ occupations, suppose the pre-AI role split is even between $\{1,2\}$ and $\{3\}$, and that workers initially in $\{1,2\}$ who satisfy $|d_i| > d_o^*$ are reassigned to whichever specialist role matches the sign of d_i, while the remainder remain bundled. Using the coefficients in Table 3, the role-specific contribution to the omitted redesign term $\mathcal{R}_{io}$ is dominated by the $-7d/6$ and $5d/3$ entries on the off-diagonal cells. Suppose AI exposure ΔXPT_o^{AT} is correlated with the share of $\{1,2\}$ workers reassigned to task-3 specialist work, π_o, with correlation $\rho_{X\pi}$, and that the residualised variance of exposure is $\mathrm{Var}(\widetilde{\Delta XPT_o^{AT}}) = 1$ after standardisation. Plugging into (52) gives the bias

$$\text{bias} \approx -\tfrac{7}{6}\,\rho_{X\pi}\,\sigma_\pi\ \mathrm{E}[d_i \mid d_i < -d_o^*],$$

where σ_π is the cross-occupation standard deviation of π_o and the conditional mean of d_i in the task-3 specialist tail is approximately -0.20 under the assumed Gaussian shape. With $\sigma_\pi = 0.10$, the bias is approximately $0.023\,\rho_{X\pi}$. Table 4 shows the implied bias for three illustrative correlations.

For an exposure coefficient in the range of -0.005 to -0.020 log points per unit standardised exposure, biases of this magnitude can flip the sign at moderate correlation, and do so in the displayed scenario at $\rho_{X\pi} = 0.8$. The illustration uses only equations (50)–(52) and the coefficients in Table 3; it is not a calibration and it does not establish a general sign. Its purpose is to show that the omitted-redesign term can be economically material under plausible correlations.

Measurement and identifying assumptions. A role-by-role implementation of (53) requires measures of pre-AI task bundling, post-AI task bundling and worker skills. Pre-AI co-tasking can proxy for the partial-baseline indicator $B^0_{12|3,o}$ when tasks 1 and 2 are jointly demanded within an occupation and tasks 2 and 3 are not. Pre-AI sub-occupation or job-description data can proxy for $R^0_{12,i}$ and $R^0_{3,i}$: workers whose original job description bundled task 1 with task 2 are assigned to $R^0_{12,i} = 1$, and workers in the separately priced task-3 role are assigned to $R^0_{3,i} = 1$. Post-AI co-tasking and specialist-title entry can proxy for $I^1_{C,io}$, $I^1_{2,io}$ and $I^1_{3,io}$: $I^1_{C,io} = 1$ when post-AI vacancies and titles bundle tasks 2 and 3, and $I^1_{2,io} = 1$ or $I^1_{3,io} = 1$ when vacancies and titles separate them. The cross-product $R^0_{r,i} I^1_{k,io}$ is then measured from the interaction of pre-AI co-tasking patterns with post-AI co-tasking changes. The role-specific skill coefficients $\gamma_{rk,z}$, $\gamma_{rk,a}$ and $\gamma_{rk,d}$ are estimated from within-cell variation in z_i, a_i and d_i across workers; the restricted version of (53) replaces the parameters by the model-implied entries of Table 3 and tests the joint restriction.

Table 4: Illustrative magnitude of the omitted-variable bias on $\hat{\beta}_X^{AT}$

Correlation $\rho_{X\pi}$ between exposure and task-3 specialist reassignment	Bias on $\hat{\beta}_X^{AT}$
0.2	$\approx$ 0.005 log points
0.5	$\approx$ 0.012 log points
0.8	$\approx$ 0.019 log points

This empirical design is not identified by vocabulary alone. It requires that co-tasking measures recover priced task bundles rather than mere textual co-occurrence; that post-AI titles and vacancies describe actual assignments rather than relabelling; that proxies for z_i, a_i and d_i measure the relevant worker skills; that post-AI role assignment is not driven by unobserved wage shocks after conditioning on the included controls, or is otherwise instrumented; and that occupational reclassification does not mechanically create the specialist-entry signal. Under these assumptions, (53) is a model-based estimating equation for the redesign channel; without them, it is a descriptive decomposition of observed co-tasking and wage changes.

Signing the bias under sufficient conditions. The current literature does not pin down the sign of the bias in (52) from exposure alone. A sign conclusion requires sufficient conditions on the omitted term. One such case is an occupation group in which high exposure predicts specialist reassignment, the specialist premium is positive for the affected workers, product-price and participation effects are either controlled for or orthogonal to exposure, and the relevant skill-imbalance distribution loads positive omitted redesign terms onto high-exposure occupations. In that case, a fixed-bundle regression attributes the omitted comparative-advantage premium to AI exposure itself and overstates the wage effect of exposure. A second case is an occupation group in which high exposure predicts rebundling and the omitted role-specific redesign term is negative for the affected workers after the same controls. In that case, the bias runs the other way and the fixed-bundle regression understates the wage effect of exposure. Outside such restrictions, (52) has no sign theorem. The testable object is therefore the local covariance between exposure, post-AI partition and the signed distribution of surviving-skill imbalance, occupation by occupation.

The same logic applies to employment. Falling measured expertise requirements can expand employment in a fixed-bundle occupation, but rebundling may raise the breadth of surviving skill required after the automated task is removed. Unbundling may expand entry into one specialist role while shrinking another. Occupation-level employment changes can therefore mix task removal, job redesign and occupational reclassification.

When only occupation-level data are available, the worker-level interactions in (53) should be replaced by role shares and conditional moments, such as $\Pr(R^0 = r, I^1 = k \mid o)$, $\mathrm{E}[z_i \mid R^0 = r, I^1 = k, o]$, $\mathrm{E}[a_i \mid R^0 = r, I^1 = k, o]$ and $\mathrm{E}[d_i \mid R^0 = r, I^1 = k, o]$.

The empirical implication is narrow but important. Fixed-bundle estimates are informative about fixed-bundle wage effects. They become incomplete estimates of automation's effect on workers when automation changes which surviving skills are bundled into jobs.

The omitted-redesign term is most consequential in occupations where AI exposure is correlated with post-AI co-tasking changes, specialist-role entry, or the signed distribution of surviving-skill imbalance.

7 Conclusion

Endogenous bundling makes AI exposure an incomplete predictor of wage effects. The partial pre-AI baseline $\{1,2\}|\{3\}$ is the minimal environment in which this matters: automation of task 1 removes a human return, repositions task 2 in either a rebundled or specialist role, and can make a previously unpriced surviving skill the dominant component of a worker's post-AI wage. The wage decomposition is role-specific. A worker who began bundled and is reassigned as a task-3 specialist has wage-impact coefficient $-7/6$ on signed surviving-task imbalance under the symmetric selected menu. That coefficient reflects both the post-AI price of task-3 skill and the subtraction of the pre-AI return to task-2 skill, and it cannot be recovered by a fixed-bundle earnings function because the relevant skill was not in the pre-AI bundle and the relevant role transition was not in the pre-AI choice set.

The empirical agenda is to estimate (53) using pre-AI co-tasking proxies for the partial-baseline indicator and pre-AI roles, and post-AI co-tasking changes and specialist-title entry as proxies for the post-AI assignment indicators. The role-specific coefficient pattern in Table 3 is the model's testable prediction under the symmetric selected menu. The sign of the omitted-variable bias on a fixed-bundle AI-exposure coefficient is not determined by exposure alone. It depends on the covariance between exposure and the full omitted redesign term, including post-AI partition, worker role transitions, signed skill imbalance, product prices and participation. The data should therefore test the local relationship between exposure and post-AI partition rather than assume a single sign.

APPENDICES

A Pre-AI Equilibrium and Baseline Support

This appendix collects the closed-form pre-AI equilibrium that supports Assumption 1, the formal definition of the partial-baseline-admissible class and its surviving-skill image, and the integrated-baseline counterpart of the wage-impact decomposition.

A.1 Closed-form selected pre-AI equilibrium

Let $N = \{1, 2, 3\}$ and let $\mathcal{P}^0 = \Pi(N)$ be the five feasible pre-AI partitions. For every partition $\rho \in \mathcal{P}^0$, define the pre-AI organisational multiplier

$$A^0_\rho \equiv \Phi^0_\rho \prod_{B \in \rho} \Omega(B). \tag{54}$$

For a block $B \subseteq N$, let

$$\bar{r}_B \equiv \frac{1}{\kappa_B} \sum_{j \in B} \alpha_j r_j, \qquad \kappa_B = \sum_{j \in B} \alpha_j. \tag{55}$$

Thus $\bar{r}_B$ is the task-weighted average log skill relevant for a worker assigned to block B. The output of a team using mode ρ can be written as

$$f^0_\rho(r^B : B \in \rho) = A^0_\rho \exp\left\{ \sum_{B \in \rho} \kappa_B \bar{r}^B_B \right\}. \tag{56}$$

A pre-AI role-price wedge is a collection $\xi = \{\xi_{\rho,B} : \rho \in \mathcal{P}^0, B \in \rho\}$ satisfying the within-mode normalisation

$$\sum_{B \in \rho} \kappa_B \xi_{\rho,B} = 0 \qquad \text{for every } \rho \in \mathcal{P}^0. \tag{57}$$

Given product price $p > 0$, define the value of assigning worker $r = (z, x, y)$ to block B in mode ρ by

$$W^{0,\xi}_{\rho,B}(r; p) = p A^0_\rho \kappa_B \exp\{\bar{r}_B + \xi_{\rho,B}\}. \tag{58}$$

The selected pre-AI payoff menu is

$$u^{0,\xi}_p(r) = \max\left\{ \bar{u}, \max_{\rho \in \mathcal{P}^0} \max_{B \in \rho} W^{0,\xi}_{\rho,B}(r; p) \right\}. \tag{59}$$

Let

$$\mathcal{R}^{0,\xi}_{\rho,B}(p) = \{r \in \tilde{X} : W^{0,\xi}_{\rho,B}(r;p) = u^{0,\xi}_p(r)\} \tag{60}$$

be the region of types for which role (ρ, B) attains the selected menu. If $\bar{u} = 0$, these role-choice regions do not depend on p.

Proposition A.1 (Closed-form selected pre-AI equilibrium)**.** *Fix $p > 0$ and a wedge vector ξ satisfying (57). Suppose there exist finite non-negative measures $\gamma_{\rho,B}$ on $\tilde{X}$, an idle-worker measure ζ, and finite score measures ν_ρ on $\mathbb{R}$ such that:*

(i) $\gamma_{\rho,B}$ is supported on $\mathcal{R}^{0,\xi}_{\rho,B}(p)$ for every ρ and $B \in \rho$, while ζ is supported on $\{r : u^{0,\xi}_p(r) = \bar{u}\}$;

(ii) resources are exhausted,

$$\zeta + \sum_{\rho \in \mathcal{P}^0} \sum_{B \in \rho} \gamma_{\rho,B} = F; \tag{61}$$

(iii) for each mode ρ, the adjusted role-score pushforwards are common across all blocks,

$$(\bar{r}_B + \xi_{\rho,B})_\# \gamma_{\rho,B} = \nu_\rho \qquad \text{for every } B \in \rho; \tag{62}$$

(iv) the product price clears the product market,

$$p = P(Y^0), \qquad Y^0 = \sum_{\rho \in \mathcal{P}^0} \int A^0_\rho e^q \, \mathrm{d}\nu_\rho(q). \tag{63}$$

Then $u^{0,\xi}_p$ and the induced score-assortative matchings decentralise a pre-AI production-mode equilibrium. The pre-AI active support is

$$\mathcal{M}^0 = \{\rho \in \mathcal{P}^0 : \nu_\rho(\mathbb{R}) > 0\}. \tag{64}$$

Conversely, any equilibrium supported by the selected menu (59) has active modes characterised by (64), and every active team in mode ρ matches workers so that $\bar{r}_B + \xi_{\rho,B}$ is common across all $B \in \rho$.

The theorem gives an explicit selected-menu support test for each of the five pre-AI partitions. The integrated mode $\{1,2,3\}$ is active exactly when the corresponding score measure ν_{123} has positive mass. A partial-bundling mode such as $\{1,2\}|\{3\}$ is active exactly when there is positive common adjusted-score mass between workers assigned to the $\{1,2\}$ role and workers assigned to the $\{3\}$ role. Full separation $\{1\}|\{2\}|\{3\}$ is active exactly when the three singleton role pools have positive common adjusted-score mass.

The role-choice cells are finite intersections of half-spaces. For any two candidate roles (ρ, B) and (ρ', B'), the boundary on which they give the same selected role value is $\bar{r}_B - \bar{r}_{B'} = \ln(A^0_{\rho'}\kappa_{B'}/A^0_{\rho}\kappa_B) + \xi_{\rho',B'} - \xi_{\rho,B}$. Thus F enters the pre-AI support through the mass and score distributions of these cells, while $\{\Phi^0_\rho\}$ enters through the constants A^0_ρ.

In the symmetric task-weight case $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$, write $r = (z, x, y)$ and abbreviate the five mode multipliers by A^0_{123}, $A^0_{12|3}$, $A^0_{13|2}$, $A^0_{23|1}$, and $A^0_{1|2|3}$. The selected role values can be written as

$$W^0_{123}(z,x,y;p) = pA^0_{123}\exp\left\{\frac{z+x+y}{3}\right\}, \tag{65}$$

$$W^0_{12|3,12}(z,x;p) = pA^0_{12|3}\frac{2}{3}\exp\left\{\frac{z+x}{2} - \frac{\delta_{12|3}}{3}\right\}, \tag{66}$$

$$W^0_{12|3,3}(y;p) = pA^0_{12|3}\frac{1}{3}\exp\left\{y + \frac{2\delta_{12|3}}{3}\right\},$$

$$W^0_{13|2,13}(z,y;p) = pA^0_{13|2}\frac{2}{3}\exp\left\{\frac{z+y}{2} - \frac{\delta_{13|2}}{3}\right\}, \tag{67}$$

$$W^0_{13|2,2}(x;p) = pA^0_{13|2}\frac{1}{3}\exp\left\{x + \frac{2\delta_{13|2}}{3}\right\},$$

$$W^0_{23|1,23}(x,y;p) = pA^0_{23|1}\frac{2}{3}\exp\left\{\frac{x+y}{2} - \frac{\delta_{23|1}}{3}\right\}, \tag{68}$$

$$W^0_{23|1,1}(z;p) = pA^0_{23|1}\frac{1}{3}\exp\left\{z + \frac{2\delta_{23|1}}{3}\right\},$$

$$W^0_{1|2|3,j}(r_j;p) = pA^0_{1|2|3}\frac{1}{3}\exp\{r_j + \xi^S_j\}, \qquad j = 1,2,3. \tag{69}$$

where $\xi^S_1 + \xi^S_2 + \xi^S_3 = 0$. The partial-mode wedges are normalised so that the zero-profit matching conditions are

$$\frac{z+x}{2} - y' = \delta_{12|3}, \tag{70}$$

$$\frac{z+y}{2} - x' = \delta_{13|2}, \tag{71}$$

$$\frac{x+y}{2} - z' = \delta_{23|1}, \tag{72}$$

for active partial-bundling teams, and the full-separation matching condition is

$$z + \xi^S_1 = x' + \xi^S_2 = y'' + \xi^S_3. \tag{73}$$

Equations (65)–(69) are the complete pre-AI wage menu for all possible partitions under this selected equilibrium.

The partial baseline restriction used in the main text is a restriction on the support generated by this theorem. The maintained baseline is $\mathcal{B}_{12|3} = \{\{1,2\}|\{3\}\}$, and under the selected menu a fixed pre-AI primitive-equilibrium pair supports it when $\nu_{12|3}(\mathbb{R}) > 0$ and $\nu_{123}(\mathbb{R}) = 0$: the partial mode is active and the integrated mode is excluded. The theorem characterises all five partitions because the non-focal partitions are competing production modes whose role values determine whether the partial baseline can be sustained in equilibrium. The restriction is imposed after the pre-AI primitives and the selected pre-AI support have been fixed; it is not obtained by varying the coordination factors or the equilibrium selection to rationalise an arbitrary post-AI surviving-skill marginal.

A.2 Partial-baseline-admissible class and its surviving-skill image

Collect the pre-AI primitives in $\theta^0 = (F, \{\Phi^0_\rho\}_{\rho\in\mathcal{P}^0}, P, \bar{u})$ and a selected pre-AI equilibrium in $e^0 = (p^0, \xi^0, \{\gamma^0_{\rho,B}\}_{\rho,B}, \zeta^0, \{\nu^0_\rho\}_{\rho\in\mathcal{P}^0})$. The selected mass of teams using mode ρ is $m^0_\rho = \nu^0_\rho(\mathbb{R})$. Let $\mathfrak{T}^0_{12|3}$ denote a fixed class of primitive-equilibrium pairs (θ^0, e^0) satisfying Assumption 1, that is, $m^0_{12|3} > 0$ and $m^0_{123} = 0$. In a single-economy transition, $\mathfrak{T}^0_{12|3}$ is the singleton containing the actual pre-AI primitive vector together with the selected pre-AI equilibrium used for the wage comparison. The maintained baseline set is $\mathcal{B}_{12|3} = \{\{1,2\}|\{3\}\}$. The restricted post-AI continuation set is the image of this fixed admissible class under automation of task 1:

$$\mathcal{G}_{12|3}(\mathfrak{T}^0_{12|3}) = \{\pi_{23\#}F : ((F, \{\Phi^0_\rho\}, P, \bar{u}), e^0) \in \mathfrak{T}^0_{12|3}\}, \tag{74}$$

where $\pi_{23}(z,x,y) = (x,y)$. When the class $\mathfrak{T}^0_{12|3}$ is clear, write $\mathcal{G}_{12|3}$ for this set. The set is the image of a fixed admissible class of pre-AI primitive-equilibrium pairs under task-1 automation; it is not the set of all G that could be made admissible by re-choosing pre-AI organisational parameters or by selecting a different pre-AI support after the fact.

A.3 Integrated pre-AI baseline

The maintained transition analysis begins from $\{1,2\}|\{3\}$. If a fixed pre-AI primitive-equilibrium class instead supports the integrated mode $I = \{1,2,3\}$, the same post-AI modes have different meanings: $C = \{2,3\}$ is bundle contraction, while $U = \{2\}|\{3\}$ is unbundling of tasks that were jointly priced before AI. The corresponding restricted continuation set is

$$\mathcal{G}_{123}(\mathfrak{T}^0_{123}) = \{\pi_{23\#}F : ((F, \{\Phi^0_\rho\}, P, \bar{u}), e^0) \in \mathfrak{T}^0_{123}\}, \tag{75}$$

for a fixed class $\mathfrak{T}^0_{123}$ of pre-AI primitive-equilibrium pairs with $m^0_{123} > 0$. As in the partial case, this class is fixed before projection.

In the symmetric integrated benchmark,

$$\ell^0_I(z,x,y) = \ln p^0 + \ln A^0_I + \frac{z+x+y}{3}, \qquad A^0_I = \Phi^0_I \Omega(\{1,2,3\}). \tag{76}$$

Subtracting this pre-AI payoff from the post-AI selected schedules gives

$$\Delta_C = \ln\left(\frac{p^1}{p^0}\right) + \ln\left(\frac{A^1_C}{A^0_I}\right) - \frac{z}{3}, \tag{77}$$

$$\Delta_U = \ln\left(\frac{p^1}{p^0}\right) + \ln\left(\frac{A^1_U/2}{A^0_I}\right) - \frac{z}{3} + \frac{2|d|}{3}, \tag{78}$$

where $A^1_C = A_C(m_A)$, $A^1_U = A_U(m_A)$ and $d = (x-y)/2$. The integrated extension subtracts the pre-AI return to the full index $(z+x+y)/3$; the maintained partial-baseline formulas in Proposition 7 subtract the selected pre-AI role index for the worker's actual baseline role. In this baseline the word "unbundling" has its more familiar meaning of separating tasks that were previously jointly priced, while the partial baseline reads the same word as separating the surviving pair after automation has broken the old $\{1,2\}$ bundle.

B Proofs

B.1 Existence, decentralisation and product-market clearing

This subsection collects the formal regularity assumptions, the planner formulation, and the existence and decentralisation results that support the no-profit inequalities (9)–(11) used throughout the main text. Let μ_C be a finite measure over workers assigned to bundled production. Let μ_U be a finite measure over ordered pairs $((x,y),(x',y')) \in X \times X$, where the first worker is assigned to task 2 and the second to task 3. Its marginals are denoted μ^2_U and μ^3_U. Let η be the measure of idle workers. Feasibility is

$$\mu_C + \mu^2_U + \mu^3_U + \eta = G \tag{79}$$

as measures on X, with all measures non-negative. Aggregate output is $Y(\mu_C, \mu_U) = \int f_C(x,y;m_A)\, \mathrm{d}\mu_C(x,y) + \int f_U(x,y';m_A)\, \mathrm{d}\mu_U((x,y),(x',y'))$. Only the task-relevant coordinates x and y' enter the unbundled team's output, but the full worker type matters because it determines alternative job opportunities. The planner maximises total surplus, including idle workers' outside-option payoffs:

$$R\big(Y(\mu_C,\mu_U)\big)+\bar{u}\,\eta(X) \tag{80}$$

subject to (79).

Assumption B.1 (Regularity)**.** The type space X is compact. The measure G is finite. The functions f_C and f_U are continuous and non-negative. The inverse demand function P is continuous, strictly positive and weakly decreasing on the feasible output interval.

Proposition B.1 (Existence and decentralisation)**.** *Under Assumption B.1, the planner problem (80) has a solution. Since P is weakly decreasing, R is concave. If (μ_C^*,μ_U^*,η^*) solves (80) and $Y^*=Y(\mu_C^*,\mu_U^*)$, then the product price $p=P(Y^*)$ and a continuous worker payoff function $u\in C(X)$ satisfy (9)–(11) for all worker types, with equalities η^*-, μ_C^*-, and μ_U^*-almost everywhere on the corresponding active supports.*

The inequalities are the competitive discipline. A one-worker bundled firm cannot earn a positive profit by hiring any worker. A two-worker specialist firm cannot earn a positive profit by pairing any task-2 worker with any task-3 worker. An idle worker earns the outside option. The payoff function u is a supporting wage schedule, or core payoff in the assignment economy. It need not be unique; the symmetric core selection in Section 3.1 is one transparent member of the set.

When $\bar{u}=0$, all workers participate, role choice is independent of p, and the equilibrium product price is $p^*=P(Y^0)$, where Y^0 is the output of the menu-induced assignment. With $\bar{u}>0$, define the price-free job-revenue index $V(x,y)=\max\{R_C(x,y),R_2(x),R_3(y)\}$, where R_C,R_2,R_3 are the price-free job values (16)–(18). A worker participates at price p if and only if $pV(x,y)\geq\bar{u}$, and role choice among participating workers is independent of p.

Assumption B.2 (No mass on participation and tie boundaries)**.** The measure G is atomless, the role-choice tie sets have G-measure zero, and for every $r>0$ the level set $\{(x,y):V(x,y)=r\}$ has G-measure zero.

Proposition B.2 (Product-market equilibrium under the symmetric menu)**.** *Suppose Assumptions B.1 and B.2 hold, suppose the symmetry condition (12) holds, and suppose the symmetric wage menu defined by (19) is used. If $\bar{u}=0$, the equilibrium product price is $p^*=P(Y^0)$. If $\bar{u}>0$, aggregate output $Y(p)$ induced by the menu is continuous and weakly increasing in p, and* *there is a unique price p^* satisfying*

$$p^* = P\big(Y(p^*)\big). \tag{81}$$

Higher product prices draw in additional workers by lowering the participation cutoff $\bar{u}/p$. They do not reverse the ordering of active jobs, because the same product price multiplies all job revenues. The proofs of Propositions B.1 and B.2 are below.

B.2 Proofs

Proof of optimal time allocation, (3). A worker assigned to bundle B solves

$$\max_{\{t_j\}_{j\in B}} \prod_{j\in B} (t_j s_j)^{\alpha_j} \quad \text{subject to} \quad \sum_{j\in B} t_j = 1, \quad t_j \geq 0.$$

Since the skills are fixed for the worker and strictly positive, this is equivalent to maximising $\prod_{j\in B} t_j^{\alpha_j}$. Because every $\alpha_j > 0$, the level objective is zero on any boundary point with $t_j = 0$ and strictly positive at any feasible allocation with all $t_j > 0$, so the optimum is interior. On the interior, take logs and maximise $\sum_{j\in B} \alpha_j \ln t_j$ subject to $\sum_{j\in B} t_j = 1$. The Lagrangian first-order conditions are $\alpha_j/t_j = \nu$ for all $j \in B$, where ν is the multiplier on the time constraint. Summing over $j \in B$ gives $\nu = \kappa_B$, so $t_j = \alpha_j/\kappa_B$. Substitution gives (2) and, for $B = \{2, 3\}$, the surviving-pair factor Ω_{23} used in the main text.

Lemma B.1 (Compact measure linear-programme duality)**.** *Let X be compact and let G be a finite Borel measure on X. For each element i of a finite index set I, let K_i be compact, let $\tau_{i\ell} : K_i \to X$ be continuous for $\ell = 1, \ldots, k_i$, and let $c_i \in C(K_i)$. Consider the primal problem*

$$V(G) = \sup_{\{\mu_i \geq 0\}_{i\in I}} \sum_{i\in I} \int_{K_i} c_i \,\mathrm{d}\mu_i \quad \textit{subject to} \quad \sum_{i\in I} \sum_{\ell=1}^{k_i} (\tau_{i\ell})_\# \mu_i = G. \tag{82}$$

If the feasible set is non-empty, then the primal has a maximiser. Its dual is

$$\inf_{u\in C(X)} \int_X u \,\mathrm{d}G \quad \textit{subject to} \quad \sum_{\ell=1}^{k_i} u(\tau_{i\ell}(k)) \geq c_i(k) \quad \textit{for all } k \in K_i,\ i \in I. \tag{83}$$

The dual has a minimiser, the primal and dual values are equal, and complementary slackness holds: if $\{\mu_i^*\}$ and u^* are optimal, then $\sum_{\ell=1}^{k_i} u^*(\tau_{i\ell}(k)) = c_i(k)$ μ_i^*-almost everywhere, for each i.

Proof. The feasible set is weakly compact. Feasibility gives $G(X) = \sum_{i\in I} k_i \mu_i(K_i)$, so the feasible measures are uniformly bounded in total mass. Since each K_i is compact, the positive-measure balls with these mass bounds are weakly compact. The pushforward maps are continuous under weak convergence, so the resource constraint is closed. The objective is continuous because each c_i is continuous and bounded. Hence the primal has a maximiser and a finite value.

Weak duality is immediate: if u satisfies the dual constraints, then every feasible $\{\mu_i\}$ satisfies $\sum_i \int c_i \, \mathrm{d}\mu_i \leq \sum_i \sum_{\ell=1}^{k_i} \int u(\tau_{i\ell}(k)) \, \mathrm{d}\mu_i(k) = \int u \, \mathrm{d}G$, so the dual value is at least the primal value.

For the reverse inequality and dual attainment, use the separating-hyperplane proof of compact linear-programming duality. Let $\mathcal{M}(X)$ be the signed finite Radon measures on X with the weak topology $\sigma(\mathcal{M}(X), C(X))$, and define

$$\mathcal{C} = \left\{ \left(\sum_i \sum_{\ell=1}^{k_i} (\tau_{i\ell})_\# \mu_i, \, r \right) : \mu_i \geq 0, \, r \leq \sum_i \int c_i \, \mathrm{d}\mu_i \right\} \subset \mathcal{M}(X) \times \mathbb{R}.$$

The set $\mathcal{C}$ is convex and closed. To see closedness, take a convergent sequence $(m_n, r_n) \to (m, r)$ in $\mathcal{C}$ and choose representing measures $\mu_{i,n}$. Since $m_n(X) = \sum_i k_i \mu_{i,n}(K_i)$ and $m_n(X) \to m(X)$, the representing measures have uniformly bounded masses. By compactness, pass to weakly convergent subsequences $\mu_{i,n} \Rightarrow \mu_i$. Continuity of pushforwards and of the objective gives $(m, r) \in \mathcal{C}$.

Let V be the primal value. The point (G, V) lies on the upper boundary of $\mathcal{C}$: $(G, V) \in \mathcal{C}$ by primal attainment, while $(G, V + t) \notin \mathcal{C}$ for every $t > 0$. The supporting-hyperplane theorem for closed convex subsets of locally convex spaces gives a non-zero continuous linear functional supporting $\mathcal{C}$ at (G, V) and separating it from the vertical ray above (G, V). Because the topology on $\mathcal{M}(X)$ is induced by $C(X)$, this functional has the form $(u, -\beta)$ with $u \in C(X)$ and $\beta > 0$. Normalise by $\beta = 1$. The support inequality is then $\langle u, m \rangle - r \geq 0$ for all $(m, r) \in \mathcal{C}$, with $\int u \, \mathrm{d}G - V = 0$.

Apply the first inequality to a measure concentrated in a single mode i. If $\sum_{\ell=1}^{k_i} u(\tau_{i\ell}(k)) < c_i(k)$ for some $k \in K_i$, continuity gives a neighbourhood on which the same strict inequality holds; a small positive measure on that neighbourhood would generate an element of $\mathcal{C}$ violating the support inequality. Hence $\sum_{\ell=1}^{k_i} u(\tau_{i\ell}(k)) \geq c_i(k)$ for all $k \in K_i$, so u is dual feasible. The equality $\int u \, \mathrm{d}G = V$ shows that the dual value is no greater than the primal value. Together with weak duality, the values are equal and u is a dual minimiser.

Complementary slackness follows from equality in weak duality. For any primal optimiser $\{\mu_i^*\}$, the non-negative slack $\sum_{\ell=1}^{k_i} u(\tau_{i\ell}(k)) - c_i(k)$ integrates to zero against each μ_i^*, and therefore is zero μ_i^*-almost everywhere. □

Proof of Proposition B.1. Let $\mathcal{M}(X)$ and $\mathcal{M}(X \times X)$ denote the spaces of finite Borel measures on the compact spaces X and $X \times X$, endowed with the weak topology. Feasibility implies $\mu_C(X) + \mu_U^2(X) + \mu_U^3(X) + \eta(X) = G(X) < \infty$, so the feasible set is uniformly bounded in total mass and tight. By Prokhorov's theorem, every sequence of feasible allocations has a weakly convergent subsequence. The marginal maps from $\mathcal{M}(X \times X)$ to $\mathcal{M}(X)$ are continuous under weak convergence on compact spaces, and (79) is closed because it is equivalent to $\int \varphi \, \mathrm{d}\mu_C + \int \varphi \, \mathrm{d}\mu_U^2 + \int \varphi \, \mathrm{d}\mu_U^3 + \int \varphi \, \mathrm{d}\eta = \int \varphi \, \mathrm{d}G$ for every continuous test function φ on X. Thus the feasible set is compact.

The functions f_C and f_U are continuous on compact spaces and therefore bounded and continuous. Hence $(\mu_C, \mu_U) \mapsto Y(\mu_C, \mu_U)$ is continuous under weak convergence. The feasible output interval is compact, so R is continuous on it. The objective in (80) is continuous. A maximiser exists by Weierstrass' theorem.

Let $(\mu_C^*, \mu_U^*, \eta^*)$ be an optimum, let $Y^* = Y(\mu_C^*, \mu_U^*)$, and set $p = P(Y^*)$. Since P is weakly decreasing, $R(Y) = \int_0^Y P(z) \, \mathrm{d}z$ is concave and differentiable with derivative $P(Y)$. To show that the optimum is supported by the price p, take any feasible allocation (μ_C, μ_U, η) and form the feasible path $(\mu_C^t, \mu_U^t, \eta^t) = (1-t)(\mu_C^*, \mu_U^*, \eta^*) + t(\mu_C, \mu_U, \eta)$, $t \in [0, 1]$. The function $\psi(t) = R(Y(\mu_C^t, \mu_U^t)) + \bar{u}\, \eta^t(X)$ has a maximum at $t = 0$. Because output and idle-worker mass are affine along this path, the right derivative at zero is $\psi'_+(0) = p\{Y(\mu_C, \mu_U) - Y^*\} + \bar{u}\{\eta(X) - \eta^*(X)\}$. Optimality implies $\psi'_+(0) \le 0$. Therefore $(\mu_C^*, \mu_U^*, \eta^*)$ maximises the linear programme

$$\max_{\mu_C, \mu_U, \eta} \int p f_C \, \mathrm{d}\mu_C + \int p f_U \, \mathrm{d}\mu_U + \int \bar{u} \, \mathrm{d}\eta \quad \text{subject to (79).} \tag{84}$$

Its dual over $C(X)$, namely $\min_{u \in C(X)} \int u \, \mathrm{d}G$ subject to (9)–(11), is the dual in Lemma B.1 applied to three compact modes: idle workers (team space X, reward $\bar{u}$, identity role map); bundled production (team space X, reward $p f_C$, identity role map); and specialist production (team space $X \times X$, reward $p f_U$, two coordinate projections as role maps). Hence there is an optimal $u \in C(X)$ and the primal and dual values coincide.

Complementary slackness follows from equality of the primal and dual values. The non-negative slack $u - \bar{u}$ integrates to zero against η^*, the non-negative slack $u - p f_C$ integrates to zero against μ_C^*, and the non-negative slack $u(x, y) + u(x', y') - p f_U(x, y'; m_A)$ integrates to zero against μ_U^*. Therefore the outside-option, bundled, and specialist constraints bind on the corresponding active supports almost everywhere.

Proof of Proposition B.2. If $\bar{u} = 0$, all workers participate and all job values are multiplied by the common factor p. Role choice and matching are therefore independent of p. The assignment determines a fixed finite output Y^0, and the market-clearing price is $p^* =$

$P(Y^0)$.

Let $\bar{u} > 0$. Under the symmetric menu, role choice is determined by the price-free values R_C, R_2, R_3, and participation is determined by the nested sets $\{(x,y) : V(x,y) \geq \bar{u}/p\}$. As p rises, $\bar{u}/p$ falls, so the participating set expands. Because role-choice regions are independent of p, the bundled contribution is the integral of the positive continuous function f_C over an expanding family of measurable sets.

The specialist contribution can be written using either specialist side. Under exchangeability and transposition-invariant tie-breaking, the task-2 and task-3 specialist pools have equal mass and identical distributions of role-relevant log skill after imposing any participation cutoff based on V. Equal-quantile matching therefore pairs a task-2 specialist of role-relevant skill z with a task-3 specialist of role-relevant skill z. The aggregate specialist output is $\int_{U_2(p)} A_U e^{2x/3} \, \mathrm{d}G(x,y)$, where $U_2(p)$ is the task-2 specialist region after the participation cutoff. This is again the integral of a positive continuous function over an expanding family of measurable sets. Hence $Y(p)$ is weakly increasing.

Continuity follows from dominated convergence. The integrands are bounded because X is compact. By Assumption B.2, the indicator of the participation set changes only on a G-null boundary when $p_n \to p$. Thus $Y(p_n) \to Y(p)$.

Define $T(p) = P(Y(p))$ and $H(p) = p - T(p)$. Since Y is continuous and weakly increasing and P is continuous and weakly decreasing, T is continuous and weakly decreasing. Hence H is continuous and strictly increasing. At $p = 0$, no worker participates and $H(0) = -P(0) < 0$. Since feasible output is bounded, $P(Y(p))$ is bounded above on the feasible output interval, so $H(p) > 0$ for sufficiently large p. The intermediate value theorem gives a root of (81), and strict monotonicity of H gives uniqueness.

Proof of Proposition 1. By definition, $u_p(x,y) \geq W_C(x,y;p) = pf_C(x,y;m_A)$ for every worker type, so no bundled firm can make positive profit. For any candidate specialist pair with a task-2 worker of log skill x and a task-3 worker of log skill y', $u_p(x,y) + u_p(x',y') \geq W_2(x;p) + W_3(y';p) = (pA_U/2)e^{2x/3} + (pA_U/2)e^{2y'/3} \geq pA_U e^{(x+y')/3} = pf_U(x,y';m_A)$, where the second inequality is arithmetic-geometric mean. So no specialist firm can earn positive profit either.

For active bundled workers, the assignment rule selects the bundled job only when W_C attains the maximum, so $u_p = W_C = pf_C$ and profit is zero. For active specialists, exchangeability and transposition-invariant tie-breaking imply that the task-2 and task-3 specialist pools have equal mass and identical distributions of the role-relevant index. Equal-quantile matching therefore pairs a task-2 specialist with role-relevant log skill z with a task-3 specialist with role-relevant log skill z, and for such a team $W_2(z;p) + W_3(z;p) = pA_U e^{2z/3} = pf_U(z,z;m_A)$, so active specialist firms also earn zero profit.

Proof of Proposition 2. From (14) and (15), $W_2(x;p) \geq W_3(y;p)$ if and only if $x \geq y$. Comparing task-2 specialisation and bundled work, $W_2(x;p) \geq W_C(x,y;p) \iff (A_U/2)e^{2x/3} \geq A_C e^{(x+y)/3} \iff x - y \geq h$. Similarly, $W_3(y;p) \geq W_C(x,y;p)$ if and only if $y - x \geq h$.

If $h \leq 0$, then for every worker either $x \geq y$ or $y \geq x$. In the first case, $x - y \geq 0 \geq h$, so task-2 specialisation weakly dominates bundled work and task-3 specialisation. In the second case, task-3 specialisation weakly dominates bundled work and task-2 specialisation. Tie sets have zero mass, so workers are assigned to specialist jobs almost everywhere.

If $0 < h < D_{\max}$, the inequalities above imply that task-2 specialisation uniquely maximises the menu when $x - y > h$, task-3 specialisation uniquely maximises it when $y - x > h$, and bundled work uniquely maximises it when $|x - y| < h$. Tie sets have measure zero by assumption.

If $h \geq D_{\max}$, neither $x - y > h$ nor $y - x > h$ can occur on the support, so the bundled value weakly dominates both specialist values everywhere on the support, with equality only on possible tie sets.

Proof of Proposition 3. By Proposition 2, the specialist set is $\{(x,y) : |x - y| > h\}$ up to null tie sets. If $h_2 > h_1$, then $\{|x - y| > h_2\} \subseteq \{|x - y| > h_1\}$, so $M_U(h)$ is weakly decreasing in h. Since $h = 3\ln(2\Omega_{23}) - 3\ln\Lambda(m_A)$, h falls when $\Lambda(m_A)$ rises and rises when Ω_{23} rises. A smaller Ω_{23} means a more severe time-spreading loss, so specialist worker mass rises as time-spreading becomes more costly. If $\partial \ln\Lambda(m_A)/\partial m_A > 0$, then h falls with m_A and specialist worker mass weakly rises. If Φ_C and Φ_U shift proportionally, Λ is unchanged and the sorting threshold is unchanged.

Derivation of (24). If d is uniform on $[-\sigma, \sigma]$, then $|x - y| = 2|d|$, so for $0 < h < 2\sigma$, $M_U(h) = \Pr(2|d| > h) = 1 - h/(2\sigma)$. The derivative with respect to σ is $h/(2\sigma^2) > 0$ in the interior range. If $h \leq 0$, all non-tie workers specialise; if $h \geq 2\sigma$, no worker is in a specialist tail.

Proof of Proposition 4. Using $x = a + d$ and $y = a - d$, the bundled log value is $\ln W_C = \ln p + \ln A_C + (x+y)/3 = \ln p + \ln A_C + 2a/3$, and the larger specialist log value is $\max\{\ln W_2, \ln W_3\} = \ln p + \ln(A_U/2) + (2/3)\max\{x,y\} = \ln p + \ln(A_U/2) + (2/3)(a + |d|)$. Taking the maximum of these gives (25). The specialist expression exceeds the bundled expression if and only if $(2/3)|d| \geq \ln(2A_C/A_U)$, which is $|d| \geq d^*$ by (26). If $d^* > 0$, the derivative with respect to $|d|$ is zero below d^* and $2/3$ above d^*, away from the kink. If $d^* \leq 0$, then $|d| \geq d^*$ for all workers (with only possible ties at $d^* = 0$ and $d = 0$), so the specialist expression is selected almost everywhere and the slope is $2/3$.

Proof of Corollary 1. Subtract (28) from (27). The common terms $\ln p + \ln A_C + 2a/3$ cancel, leaving (29).

Proof of Proposition A.1. For any mode ρ and candidate tuple $(r^B)_{B\in\rho}$, the weighted arithmetic-geometric mean inequality gives $\sum_{B\in\rho}\kappa_B\exp\{\bar r^B_B+\xi_{\rho,B}\}\geq\exp\{\sum_{B\in\rho}\kappa_B(\bar r^B_B+\xi_{\rho,B})\}$. Using the normalisation (57), this is $\sum_{B\in\rho}W^{0,\xi}_{\rho,B}(r^B;p)\geq pA^0_\rho\exp\{\sum_{B\in\rho}\kappa_B\bar r^B_B\}=pf^0_\rho(r^B:B\in\rho)$. Since $u^{0,\xi}_p(r)\geq W^{0,\xi}_{\rho,B}(r;p)$ for every candidate role (ρ,B), no firm can earn positive profit in any pre-AI mode.

It remains to construct the teams whose role marginals are the measures $\gamma_{\rho,B}$. Fix a mode ρ and write $s_{\rho,B}(r)=\bar r_B+\xi_{\rho,B}$. Condition (62) says that $s_{\rho,B\#}\gamma_{\rho,B}=\nu_\rho$ for every block $B\in\rho$. Since the type space is compact and hence a standard Borel space, disintegrate each role measure with respect to the common score, $\gamma_{\rho,B}(\mathrm{d}r)=\int K_{\rho,B}(q,\mathrm{d}r)\,\nu_\rho(\mathrm{d}q)$, where $K_{\rho,B}(q,\cdot)$ is a probability kernel supported on $\{r:s_{\rho,B}(r)=q\}$ for ν_ρ-almost every q. Define the team measure for mode ρ by $\mu_\rho(\mathrm{d}(r^B)_{B\in\rho})=\int\prod_{B\in\rho}K_{\rho,B}(q,\mathrm{d}r^B)\,\nu_\rho(\mathrm{d}q)$. Its Bth marginal is $\gamma_{\rho,B}$, and every team drawn from μ_ρ has a common adjusted score q across all blocks. This is the score-assortative matching induced by (62).

For such an active team, the equality condition in weighted arithmetic-geometric mean is satisfied. Because each assigned worker lies in a role-choice region, $u^{0,\xi}_p=W^{0,\xi}_{\rho,B}$ on the support of $\gamma_{\rho,B}$. Hence every active team earns zero profit. On a team with common adjusted score q, output is $A^0_\rho\exp\{\sum_{B\in\rho}\kappa_B(q-\xi_{\rho,B})\}=A^0_\rho e^q$, again by (57). Integrating over the common score measure ν_ρ gives (63). The price equation in (63) clears the product market, while (61) exhausts the worker measure. Thus the constructed allocation and payoff menu decentralise a pre-AI production-mode equilibrium.

Conversely, within the selected menu (59), zero profit for an active team in mode ρ requires equality in the weighted arithmetic-geometric mean inequality above. Equality holds only when $\bar r_B+\xi_{\rho,B}$ is common across all active roles in that team. Therefore the role marginals of any active selected-menu equilibrium have common adjusted-score pushforwards, and a mode is active exactly when that common score measure has positive mass.

Proof of Proposition 5. Use the symmetric task weights $\alpha_1=\alpha_2=\alpha_3=1/3$, set $\bar u=0$, and take inverse demand to be the constant function $P(Y)=1$ on the feasible output interval. Fix $s>0$ and choose $L>s+\ln 4$. Let Q be the uniform probability measure on $[-s,s]$. For each score $q\in[-s,s]$, define two worker-type curves $r^{12}(q)=(q,q,-L)$ and $r^3(q)=(-L,-L,q)$. Let $\gamma_{12|3,12}$ be the pushforward of Q by r^{12}, $\gamma_{12|3,3}$ the pushforward of Q by r^3, all other role measures and the idle measure zero, and $F=\gamma_{12|3,12}+\gamma_{12|3,3}$. The support of F is compact and F is atomless. Its surviving-skill marginal is exchangeable because it puts equal mass on $(x,y)=(q,-L)$ and on its transpose $(-L,q)$.

Choose the selected wedge for the partial mode to be $\delta_{12|3} = 0$, choose all other selected wedges to be zero, and set the reduced partial multiplier $A^0_{12|3} = 1$. For $r^{12}(q)$, the selected partial-role values are $W^0_{12|3,12}(r^{12}(q); 1) = \frac{2}{3}e^q$ and $W^0_{12|3,3}(r^{12}(q); 1) = \frac{1}{3}e^{-L}$; for $r^3(q)$ they are $W^0_{12|3,3}(r^3(q); 1) = \frac{1}{3}e^q$ and $W^0_{12|3,12}(r^3(q); 1) = \frac{2}{3}e^{-L}$. Because $L > s + \ln 4$, the focal-role value strictly dominates on each curve. Choose the reduced multipliers of all non-focal modes, including the integrated mode $\{123\}$, equal to a common $\varepsilon > 0$ small enough that every non-focal role value is strictly below the relevant focal value on the compact support of F. Such an ε exists by compactness and strict positivity of the focal values. The required coordination factors are obtained from $\Phi^0_\rho = A^0_\rho / \prod_{B\in\rho} \Omega(B) > 0$.

The adjusted score of the active $\{1,2\}$ worker is $(z + x)/2 = q$, and the adjusted score of the active $\{3\}$ worker is $y = q$. Hence the common score measure for the partial mode is $\nu_{12|3} = Q$, while $\nu_\rho = 0$ for every other mode. The score-clearing and resource conditions in Proposition A.1 are satisfied. Aggregate pre-AI output is $Y^0 = \int e^q \, dQ(q)$, so the constant inverse demand gives $p = P(Y^0) = 1$. Proposition A.1 decentralises a selected pre-AI equilibrium with $m^0_{12|3} = \nu_{12|3}(\mathbb{R}) = 1 > 0$ and $m^0_{123} = \nu_{123}(\mathbb{R}) = 0$, so the support is the pure partial baseline.

Proof of Proposition 6. The result follows from Proposition 2 after imposing Assumption 1. If $h \le 0$, all workers are assigned to specialist jobs up to null tie sets, so $\mathcal{M}^1 = \{U\}$. Since the pre-AI focal mode is $\{1,2\}|\{3\}$, post-AI $U = \{2\}|\{3\}$ is the unbundled outcome of the surviving tasks.

If $0 < h < D_{\max}$, Proposition 2 assigns workers with $|x - y| < h$ to the bundled mode and workers in the two tails to specialist roles. Therefore C is active exactly when $M_C(h) > 0$. The specialist mode is active exactly when both specialist role pools can be populated, equivalently $T_2(h) > 0$ and $T_3(h) > 0$. Under exchangeability these two masses are equal, but the statement records both requirements. Combining these post-AI active modes with the partial pre-AI baseline gives the two cases in Definition 1: C is rebundling and U is unbundling.

If $h \ge D_{\max}$, neither specialist tail has positive mass on the support up to null tie sets, so every active worker is assigned to the bundled mode and $\mathcal{M}^1 = \{C\}$. Since task 3 was separate before AI, post-AI $C = \{2,3\}$ is rebundling.

Proof of Proposition 7. For a worker initially in the $\{1,2\}$ role, the pre-AI log payoff is (41). Post-AI bundled work gives $\ell^1_C(x, y) = \ln p^1 + \ln A^1_C + (x + y)/3$. Using $x = a + d$ and $y = a - d$, subtracting (41) gives (43). The task-2 specialist log payoff is $\ell^1_2(x) = \ln p^1 + \ln(A^1_U/2) + 2x/3$, and subtracting (41) gives (44). The task-3 specialist log payoff is $\ell^1_3(y) = \ln p^1 + \ln(A^1_U/2) + 2y/3$, and subtracting (41) gives (45). For a worker initially in the $\{3\}$ role, the pre-AI log payoff is (42). Subtracting it from $\ell_C^1(x,y)$, $\ell_2^1(x)$ and $\ell_3^1(y)$ in turn and using $(x+y)/3 - y = -a/3 + d$ gives (46), (47) and (48) respectively.

B.3 Constant-wedge specialist support

This subsection records the constant-wedge specialist role-price family used in passing in Section 3.1. The constant-wedge family is exact when the two specialist pools can be matched by a constant shift in role-relevant skill, and it gives a transparent comparative-static representation of relative specialist scarcity. In a general asymmetric economy, equilibrium payoffs are the unrestricted dual potentials in (9)–(11); a single scalar wedge need not describe every possible asymmetry.

Let $p > 0$ be a product price and let $\eta_s \in \mathbb{R}$ be a specialist scarcity wedge. The notation η_s is used for the wage wedge and should not be confused with the idle-worker measure η. Define specialist role values

$$W_2^{\eta_s}(x;p) = pA_U\lambda\exp\{\kappa x + (1-\lambda)\eta_s\}, \tag{85}$$
$$W_3^{\eta_s}(y;p) = pA_U(1-\lambda)\exp\{\kappa y - \lambda\eta_s\}. \tag{86}$$

The wedge shifts surplus toward one specialist side while preserving the no-profit frontier: a larger η_s raises the task-2 specialist schedule and lowers the task-3 specialist schedule. The bundled value remains $W_C(x,y;p) = pA_C\exp\{\alpha_2 x + \alpha_3 y\}$.

Lemma B.2 (Specialist no-profit inequality)**.** *For every $\eta_s \in \mathbb{R}$ and every pair of worker types,*

$$W_2^{\eta_s}(x;p) + W_3^{\eta_s}(y';p) \geq pf_U(x,y';m_A). \tag{87}$$

Equality holds if and only if $y' = x + \eta_s/\kappa$.

Define the constant-wedge menu

$$u_p^{\eta_s}(x,y) = \max\{\bar{u}, W_C(x,y;p), W_2^{\eta_s}(x;p), W_3^{\eta_s}(y;p)\}. \tag{88}$$

Proposition B.3 (Constant-wedge specialist support)**.** *Fix $p > 0$ and $\eta_s \in \mathbb{R}$. Assign workers to jobs that attain the maximum in (88). Let $S_2^{\eta_s}$ and $S_3^{\eta_s}$ denote the task-2 and task-3 specialist pools generated by this assignment, including any measurable tie-breaking rule. Suppose there is a specialist matching measure μ_U with marginals equal to the assigned specialist measures and satisfying*

$$y' = x + \frac{\eta_s}{\kappa} \qquad \mu_U\text{-almost everywhere.} \tag{89}$$

Then no bundled or specialist firm can earn positive profit under $u_p^{\eta_s}$, and every active specialist team earns zero profit. Conversely, any active specialist team that earns zero profit under the constant-wedge menu must satisfy (89).

Economically, the specialist scarcity wedge measures which specialist side is harder to replace. In the constant-shift case, the two specialist pools have the same shape, except that task-3 specialists are uniformly stronger or weaker by η_s/κ in the skill relevant for specialist production. Then equal-quantile matching pairs the right workers and satisfies (89). When the asymmetry is more complicated than a common shift, however, one number cannot capture all of the scarcity differences across the two specialist pools. In that case, η_s is best viewed as a useful comparative-static summary, while the full dual payoff function is needed to describe the equilibrium.

Job choice still depends on relative skill. Workers with a large enough advantage in task 2 choose task-2 specialist jobs, workers with a large enough advantage in task 3 choose task-3 specialist jobs, and workers in between choose bundled jobs. Specifically, task-2 specialisation weakly dominates bundled work when

$$x - y \geq h_2(\eta_s) \equiv \frac{\ln\{A_C/(A_U\lambda)\} - (1-\lambda)\eta_s}{\alpha_3}, \tag{90}$$

while task-3 specialisation weakly dominates bundled work when

$$y - x \geq h_3(\eta_s) \equiv \frac{\ln\{A_C/[A_U(1-\lambda)]\} + \lambda\eta_s}{\alpha_2}. \tag{91}$$

The boundary between the two specialist roles is

$$x - y = b(\eta_s) \equiv -\frac{\eta_s + \ln\{\lambda/(1-\lambda)\}}{\kappa}. \tag{92}$$

so workers choose task-2 specialisation over task-3 specialisation when $x - y \geq b(\eta_s)$, and task 3 otherwise. A positive wedge means task-2 specialists are relatively scarce. It shifts both pay and sorting toward task 2: the task-2 specialist schedule rises, the task-3 schedule falls, the task-2 region expands, and the task-3 region shrinks. Two workers with the same overall imbalance can therefore have different wage effects depending on whether their stronger skill is the scarce one or the abundant one.

Example B.1 (Relative scarcity of task-2 skill)**.** Consider the symmetric task-weight case $\alpha_2 = \alpha_3 = 1/3$, but allow $\eta_s \neq 0$. Writing $a = (x+y)/2$ and $d = (x-y)/2$, the two specialist log schedules are

$$\ln W_2^{\eta_s}(x;p) = \ln p + \ln\left(\frac{A_U}{2}\right) + \frac{2(a+d)}{3} + \frac{\eta_s}{2},$$

and

$$\ln W_3^{\eta_s}(y;p) = \ln p + \ln\left(\frac{A_U}{2}\right) + \frac{2(a-d)}{3} - \frac{\eta_s}{2}.$$

Task-2 specialisation dominates bundled work when

$$d \geq d^* - \frac{3\eta_s}{4},$$

and task-3 specialisation dominates bundled work when

$$d \leq -d^* - \frac{3\eta_s}{4}.$$

Thus, when $\eta_s > 0$, the task-2 specialist region expands and the task-3 specialist region contracts. If $\eta_s < 0$, the pattern reverses.

Proof of Lemma B.2. Let $A = \exp\{\kappa x + (1-\lambda)\eta_s\}$ and $B = \exp\{\kappa y' - \lambda\eta_s\}$. The weighted arithmetic-geometric mean inequality gives $\lambda A + (1-\lambda)B \geq A^\lambda B^{1-\lambda}$. Multiplying by pA_U yields $W_2^{\eta_s}(x;p) + W_3^{\eta_s}(y';p) \geq pA_U \exp\{\lambda\kappa x + (1-\lambda)\kappa y'\}$. Since $\lambda\kappa = \alpha_2$ and $(1-\lambda)\kappa = \alpha_3$, the right-hand side is $pf_U(x, y'; m_A)$. Equality in weighted AM-GM holds if and only if $A = B$, which is equivalent to $y' = x + \eta_s/\kappa$.

Proof of Proposition B.3. Since $u_p^{\eta_s}(x,y) \geq W_C(x,y;p) = pf_C(x,y;m_A)$ for every worker type, no bundled firm can earn positive profit. For any candidate specialist pair, $u_p^{\eta_s}(x,y) \geq W_2^{\eta_s}(x;p)$ and $u_p^{\eta_s}(x',y') \geq W_3^{\eta_s}(y';p)$, so by Lemma B.2 the sum is at least $pf_U(x,y';m_A)$, ruling out positive specialist profit. If a pair in the active matching satisfies (89), the equality condition in Lemma B.2 holds. Since both specialists are assigned to roles attaining the maximum in (88), their payoffs equal the corresponding specialist role values, and the specialist team's profit is zero. Conversely, zero profit for an active specialist team requires equality in the no-profit inequality, which is exactly (89).

Derivation of (90)–(92). Task-2 specialisation weakly dominates bundled work if and only if $pA_U\lambda e^{\kappa x + (1-\lambda)\eta_s} \geq pA_C e^{\alpha_2 x + \alpha_3 y}$. Taking logs and using $\kappa - \alpha_2 = \alpha_3$ gives $\alpha_3(x-y) + \ln(A_U\lambda/A_C) + (1-\lambda)\eta_s \geq 0$, which is (90). The task-3 comparison is analogous and gives (91). Finally, $W_2^{\eta_s} \geq W_3^{\eta_s}$ if and only if $\kappa(x-y) + \ln\{\lambda/(1-\lambda)\} + \eta_s \geq 0$, which is (92).

C General Finite Production-Mode Assignment

This appendix gives the finite-mode version of the assignment economy. The three-task Cobb–Douglas case used in the main text is the natural environment for the rebundling-as-revaluation mechanism, but the existence and decentralisation logic does not rely on three tasks. The general statement here makes that explicit and shows that rebundling and unbundling are best read as changes in the support of the active production-mode measures after restricting pre-AI modes to the surviving task set.

Let n be the number of tasks and let $\mathcal{N} = \{1, \ldots, n\}$ be the full task set. AI performs the subset $A \subset \mathcal{N}$, leaving human task set $H = \mathcal{N} \setminus A$. Throughout this appendix, worker types are surviving log-skill vectors $r = (r_j)_{j \in H}$ in a compact set $X \subset \mathbb{R}^{|H|}$, with distribution G. A human production mode is a partition $\rho = \{B_1, \ldots, B_{K(\rho)}\} \in \Pi(H)$, where each bundle B_ℓ is assigned to one worker role. For a worker of log type r assigned to bundle B, write $\sigma_B(r) = \exp\{\sum_{j \in B} \alpha_j r_j\}$. Optimal within-bundle time allocation gives $\Omega(B) = \prod_{j \in B} (\alpha_j / \kappa_B)^{\alpha_j}$ and $\kappa_B = \sum_{j \in B} \alpha_j$. Let $\Phi(\rho, A, m_A)$ be the coordination factor induced by separating human tasks from one another and from AI tasks, and let $M(A, m_A) = \prod_{j \in A} \chi_j(m_A)^{\alpha_j}$. A team using mode ρ and workers $(r^1, \ldots, r^{K(\rho)})$ then produces

$$f_\rho(r^1, \ldots, r^{K(\rho)}; m_A) = \Phi(\rho, A, m_A) M(A, m_A) \prod_{\ell=1}^{K(\rho)} \left[\Omega(B_\ell) \sigma_{B_\ell}(r^\ell)\right]. \tag{93}$$

For every mode ρ, let μ_ρ be a finite measure over $X^{K(\rho)}$, with μ_ρ^ℓ its ℓth marginal. Feasibility with idle measure η is

$$\sum_{\rho \in \Pi(H)} \sum_{\ell=1}^{K(\rho)} \mu_\rho^\ell + \eta = G, \tag{94}$$

aggregate output is $Y(\{\mu_\rho\}) = \sum_{\rho \in \Pi(H)} \int f_\rho(r^1, \ldots, r^{K(\rho)}; m_A) \, \mathrm{d}\mu_\rho$, and the planner maximises

$$R(Y(\{\mu_\rho\})) + \bar{u} \, \eta(X) \tag{95}$$

subject to (94).

Proposition C.1 (Finite-mode existence and decentralisation). *Suppose X is compact, G is finite, every f_ρ is continuous and non-negative, and P is continuous, strictly positive and weakly decreasing on the feasible output interval. Then (95) has a solution. If R is concave, any solution can be decentralised by a product price $p = P(Y^*)$ and a continuous worker payoff function*

$u \in C(X)$ *satisfying*

$$u(r) \geq \bar{u}, \tag{96}$$

$$\sum_{\ell=1}^{K(\rho)} u(r^{\ell}) \geq p f_{\rho}(r^1, \ldots, r^{K(\rho)}; m_A) \quad \textit{for every } \rho \in \Pi(H). \tag{97}$$

The inequalities bind on the support of idle workers and active teams, respectively.

Proof. The argument is identical to the proof of Proposition B.1, applied to the finite collection of modes $\rho \in \Pi(H)$ with team spaces $X^{K(\rho)}$: tightness, weak compactness, continuity of pushforwards and continuity of the objective deliver existence of a planner optimum, and the affine perturbation argument shows that any optimum maximises the linear programme in which the planner reward is replaced by pf_ρ for $p = P(Y^*)$. The dual of that finite-mode measure linear programme is exactly the dual in Lemma B.1 applied to one mode for idle workers (team space X, reward $\bar{u}$, identity role map) and one mode per partition ρ (team space $X^{K(\rho)}$, reward pf_ρ, coordinate role maps). Hence strong duality, dual attainment, and complementary slackness follow from Lemma B.1, with equality on the support of idle workers and on the support of each active team measure. The active production modes are those for which (97) binds on a positive measure of teams. Rebundling and unbundling are therefore changes in the support of the measures $\{\mu_\rho\}_{\rho \in \Pi(H)}$ after restricting pre-AI modes to the surviving task set; cases with the same surviving-task partition before and after automation are no-redesign cases. □

References

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

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

Autor, D. and N. Thompson (2025): "Expertise," *Journal of the European Economic Association*, 23(4), 1203–1271.

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

Demirer, M., J. J. Horton, N. Immorlica, B. Lucier, and P. Shahidi (2026): "Chaining tasks, redefining work: A theory of AI automation," NBER Working Paper No. 34859.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024): "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 384(6702), 1306–1308.

Gans, J. S. and A. Goldfarb (2026): "O-Ring Automation," NBER Working Paper No. 34639.

Garicano, L., J. Li, and Y. Wu (2026): "Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries," CEPR Discussion Paper No. DP21453.

Kremer, M. (1993): "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics*, 108(3), 551–575.

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

Ocampo, S. (2022): "A Task-Based Theory of Occupations with Multidimensional Heterogeneity," CHCP Working Paper No. 2022-2, University of Western Ontario.