NBER WORKING PAPER SERIES

BUT I LIKE DOING THIS!
ENJOYABLE TASKS, CONTRACTING, AND AUTOMATION

Joshua S. Gans

Working Paper 35309
http://www.nber.org/papers/w35309

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

Thanks to Refine.ink, ChatGPT 5.5 and Claude Opus 4.8 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/w35309

But I Like Doing This! Enjoyable Tasks, Contracting, and Automation
Joshua S. Gans
NBER Working Paper No. 35309
June 2026
JEL No. J22, J24, J31, J33, O33

ABSTRACT

Workers sometimes enjoy productive tasks and voluntarily devote unpaid time to them. We study O-ring jobs in which firms can either price a complete task bundle or specify paid task floors while workers remain free to add time. For any allocation supported by both hourly contracts, the wage bill is identical: voluntary top-up is not a discount. The contracts differ because paid floors cannot cap attractive tasks below the worker's voluntary supply. This implementability constraint adds a containment motive for automation alongside replacement and scale effects. It also makes payroll measures incomplete: conditional on a common automated set and a common AI technology, jobs with the same payroll footprint can differ in worked time and task mix. Rich salaried bundle pricing removes the hourly-contract distortion.

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

1 Introduction

A software engineer spends their scheduled day writing documentation, answering support tickets, and fixing bugs. After hours, they return to a tool they have been refining. The additional time is unpaid in payroll records, but it need not be uncompensated in utility: working on the tool may itself be valuable to them at the margin.

In the costly-task benchmark for standard task-based models of automation, human tasks must be purchased from workers. Automation then asks which compensated tasks are replaced and how the productivity of retained labour changes. Our starting point is that some productive tasks need not have this form. A task that is pleasant at the margin reduces the compensation required for assigned work; a task sufficiently attractive relative to leisure can be supplied beyond the paid assignment. In the latter case, productive time need not coincide with purchased time.

In this paper, we ask how automation changes job design when some productive time may be worker-chosen rather than purchased by the firm, and what payroll data reveal in that setting. We study O-ring jobs (where tasks are complementary) in which some tasks may be intrinsically rewarding to the worker. The key contractual asymmetry is that a firm may be able to specify the composition of paid tasks without being able to cap workers' time on a sufficiently attractive task. To represent that asymmetry without treating required hours as coerced labour, we compare two hourly contracts. Under a *controlled-bundle hourly contract*, the firm implements a final paid task bundle through task shares and an hourly wage. Under a *paid-floor hourly contract*, the firm posts an hourly wage and a required task-share bundle for paid work; the worker chooses paid hours and may add unpaid time to particular tasks. Paid floors constrain the composition of paid work, not an exogenously imposed number of paid hours. Once that distinction is made, automation can affect both the allocation of compensated effort and the boundary between paid responsibilities and voluntary productive work.

The distinction matters for the empirical study of generative artificial intelligence and work. Three patterns recur. User-level time savings can be substantial while earnings and recorded hours barely move (Humlum and Vestergaard, 2025a; Bick et al., 2026). AI exposure can intensify work, extending the working day or expanding the scope of tasks (Ranganathan and Ye, 2026; Jiang et al., 2025). And the within-worker task mix can shift, with workers spending more time on core or higher-value tasks and less on auxiliary work (Hoffmann et al., 2025; Brynjolfsson et al., 2025). We provide a single contracting mechanism that can jointly generate these patterns in suitable parameter regions, while making clear that it does not predict each pattern universally.

We consider one firm, one worker, and many complementary tasks. Output is produced by an O-ring technology, so every task in the benchmark technology is essential, and each task's marginal product depends on the hours supplied to the others. The worker values leisure and may directly value or dislike any individual task. Our baseline contains the standard costly-task case: if every task is disliked at every relevant hour, every productive hour carries a positive compensation gap, paid hours coincide with worked hours, and paid floors do not restrict the firm's feasible allocations. That nested case is the costly-task O-ring automation model of Gans and Goldfarb (2025), in which every human task must be purchased. Relative to it, intrinsic task value changes automation in three ways: it re-weights which tasks are replaced, it can lead the firm to automate a task precisely because its voluntary expansion cannot be contained, and it breaks the equality of worked and paid time on which payroll measurement relies. The additional questions arise only when a task is sufficiently attractive relative to leisure to be expanded voluntarily. We do not require that every task be enjoyable for an initial interval, and we permit marginal enjoyment curves to cross rather than impose a universal ranking of tasks.

Our first result is a wage-bill equivalence. For any final allocation implementable under paid floors, the wage bill equals the wage bill required to implement the same allocation under a controlled-bundle hourly contract. Unpaid top-up is, therefore, not a cash saving on a given final allocation. The two hourly regimes differ in their feasible allocations: paid floors cannot implement a task bundle that caps a sufficiently attractive task below the worker's voluntary level. When that constraint binds, job design is described by a Kuhn-Tucker system rather than by the controlled-bundle first-order condition task by task.

Automation then operates through three margins, two inherited from the costly-task benchmark and one new. The *replacement margin* removes compensated work, and the *scale margin* raises the marginal product of retained human tasks; both are present in the standard model, but enjoyment re-weights replacement, since a task rewarding at the margin saves on the wage bill and is cheaper to retain. The *containment margin* has *no* counterpart without enjoyable tasks: under paid floors, automation can avoid retaining a task whose uncapped voluntary top-up would lower the firm's paid-floor value. Enjoyment thus places two opposing forces on the same task: a wage-bill credit that favours retention and an implementability loss that favours automation, whereas the benchmark gives the firm a single, cost-driven reason to automate. Their balance, not a ranking of tasks by enjoyment, governs selection: a clean cutoff in which less rewarding tasks are automated first follows only under additional dominance, participation, and implementability conditions, and otherwise selection need not be monotone in enjoyment. Long-run redesign can then widen the paid-worked-hours wedge, but it can also reduce or eliminate it.

The empirical implication is a payroll-only non-identification result. Conditional on a common observed automated set, a common AI technology, and the floor-accounting convention stated below, two post-automation configurations can have the same measured hourly wage, wage bill, and paid hours while differing in total worked time, unpaid top-up, and task mix, because task utilities are unobserved. The construction does not hold output, realised AI input, pre-period behaviour, or worker identity fixed; those data can break the equivalence. The result is not a claim that a worker with fixed preferences moves between those allocations without a payroll response; it states what payroll observables alone cannot identify.

We contribute to four literatures. We retain the task-based decomposition of technological change (Autor et al., 2003; Acemoglu and Autor, 2011; Autor, 2015; Acemoglu and Restrepo, 2018, 2019, 2020, 2022) but shift attention to how the worker values tasks at the margin (Acemoglu et al., 2025). We add a contracting mechanism to the empirical literature on generative AI at work (Brynjolfsson et al., 2025; Noy and Zhang, 2023; Peng et al., 2023; Dell'Acqua et al., 2026; Humlum and Vestergaard, 2025a,b; Bick et al., 2026; Ranganathan and Ye, 2026; Jiang et al., 2025; Hoffmann et al., 2025). We connect compensating differentials and meaningful work (Rosen, 1974, 1986; Mas and Pallais, 2017; Maestas et al., 2023; Nikolova and Cnossen, 2020; Cassar and Meier, 2018; Bailey et al., 2019) to the within-job intensive margin. Finally, we combine the multitasking-and-incomplete-contracts tradition of Holmstrom and Milgrom (1991) with O-ring production (Kremer, 1993) and allocation of time (Becker, 1965). The closest antecedent on the production side is Gans and Goldfarb (2025), who study automation in O-ring jobs while abstracting from task-level intrinsic utility. As noted above, adding task-level intrinsic utility and paid-floor implementation changes the automation problem from compensated task substitution alone to one involving implementability, voluntary expansion, and the informational content of payroll measures.

The rest of the paper is organised as follows. Sections 2–4 build the model and study automation. Section 5 states the payroll non-identification result. Section 6 studies the rich salaried benchmark. Section 7 discusses scope and empirical implications, and Section 8 concludes. Appendices provide supporting derivations, welfare comparisons, and the common-technology numerical construction.

2 Environment and hourly contracts

2.1 Technology, preferences, and notation

A representative firm produces a single output using tasks indexed by $i \in \mathcal{N} = \{1, \dots, n\}$. Human hours on task i are $h_i \geq 0$. We write total human worked time as $L \equiv \sum_{i \in \mathcal{N}} h_i$, so that $T - L$ is leisure when $T > 0$ is the worker's time endowment. This notation reserves $\mathcal{H}$ for a set of human-retained tasks after automation.

Before automation, output requires every benchmark task and is given by the O-ring technology

$$Y(h) = \underbrace{\prod_{i=1}^{n} h_i^{\beta}}_{\text{complementary human-task output}}, \qquad \beta \in (0, 1/n). \tag{1}$$

Every task is productive and essential in this benchmark technology; $n\beta < 1$ gives decreasing returns to scale.

When all performed hours are paid at wage w, worker utility separates labour income, leisure, and intrinsic task value:

$$U = \underbrace{wL}_{\text{labour income}} + \underbrace{v(T-L)}_{\text{leisure utility}} + \underbrace{\sum_{i \in \mathcal{N}} u_i(h_i)}_{\text{intrinsic task utility}}. \tag{2}$$

Under a paid-floor contract, only paid hours P are compensated, and utility becomes

$$U = \underbrace{wP}_{\text{paid labour income}} + \underbrace{v(T-L)}_{\text{leisure utility}} + \underbrace{\sum_{i \in \mathcal{N}} u_i(h_i)}_{\text{intrinsic task utility}}. \tag{3}$$

The functions u_i capture intrinsic task utility. At a specified hour level, a task is *rewarding at the margin* if $u_i'(h_i) > 0$, neutral at the margin if $u_i'(h_i) = 0$, and *disliked at the margin* if $u_i'(h_i) < 0$; no baseline ranking of tasks is imposed.

We use the marginal value of leisure repeatedly:

$$\lambda(L) \equiv \underbrace{v'(T-L)}_{\text{marginal value of leisure}}. \tag{4}$$

Table 1 records the core notation. Symbols used for a particular result or example are defined when they arise.

Symbol	Meaning
$i \in \mathcal{N} = \{1, \ldots, n\}$	Task index and benchmark task set.
$h_i, \quad L = \sum_i h_i, \quad T - L$	Human hours on task i; total worked time; leisure time.
$Y, \quad \beta$	Output and common O-ring task elasticity.
$u_i(h_i), \quad v(T-L), \quad \lambda(L)$	Intrinsic task utility; leisure utility; marginal value of leisure.
$C, \quad PF$	Controlled-bundle and paid-floor hourly regimes.
$w, \quad \rho, \quad P, \quad p_i = \rho_i P$	Hourly wage; paid-floor task shares; paid hours; realised paid floor on task i.
$E = L - P, \quad W, \quad \Pi^C$	Unpaid top-up; wage bill; controlled-bundle operating surplus.
$\mathcal{A}, \quad \mathcal{H} = \mathcal{N} \setminus \mathcal{A}$	Automated and retained human-task sets.
$a_j, \quad q_A, \quad c, \quad Y_{\mathcal{A}}(h)$	AI input; AI productivity; AI-input cost; AI-optimised output conditional on retained human hours.
$SR, \quad LR$	Short-run and long-run outcomes after automation.

Table 1: Core notation. Result-specific and numerical-example notation is introduced locally.

2.2 Maintained assumptions

Two assumptions discipline the model throughout. The first restricts preferences and the worker's feasible labour supply, while permitting tasks that are disliked from the first hour and marginal task utilities that cross.

Assumption 1 (Preferences and feasible work)**.** *The function $v : (0, T] \to \mathbb{R}$ is three times continuously differentiable, strictly increasing, and strictly concave, with $v'(s) \to +\infty$ as $s \downarrow 0$. Each $u_i : [0, T) \to \mathbb{R}$ is three times continuously differentiable and strictly concave. The sign of $u_i'(0)$ is unrestricted. Contracts considered below satisfy $0 < L < T$, non-negative implementing wages, and participation at the worker's outside option.*

Assumption 1 permits tasks disliked at the margin from the first hour, tasks initially rewarding at the margin and later tiring, and crossings in marginal task utility. An example used for calculations is $u_i(h) = \alpha_i + \theta_i h - \frac{\kappa}{2} h^2$, where $\kappa > 0$ and θ_i may have either sign.

The second assumption is a regularity condition that ensures the firm's optimisation problems are well posed. It separates the economic mechanism from global shape restrictions on primitives.

Assumption 2 (Existence, curvature, and regularity)**.** *Whenever a global optimum is asserted, the corresponding reduced firm objective has an interior maximiser and is strictly concave on the relevant interior domain. Whenever a paid-floor Kuhn–Tucker characterisation is asserted,* *the stated local solution satisfies a constraint qualification. The same requirements apply to the AI-optimised reduced objective when automation is present.*

This assumption separates the mechanism from global shape restrictions. For the non-automated controlled-bundle objective below, sufficient curvature restrictions include $2u_i''(h) + hu_i'''(h) \leq 0$ and $2v''(T-L) - Lv'''(T-L) \leq 0$, together with $n\beta < 1$. Paid-floor constraints need not define a convex feasible set, so the paid-floor first-order system is stated locally unless global optimality is separately established.

2.3 The two hourly contracts

Controlled-bundle hourly contract. For a target final allocation h with $L > 0$, let $\sigma_i = h_i/L$. The firm specifies the paid task bundle σ and posts an hourly wage. The worker chooses paid hours L along that bundle, maximising $wL + v(T-L) + \sum_i u_i(\sigma_i L)$. At an interior implemented target, the worker's first-order condition gives the compensating hourly wage

$$w^C(h) = \underbrace{\lambda(L)}_{\text{leisure opportunity cost}} - \underbrace{\sum_i \frac{h_i}{L} u_i'(h_i)}_{\text{average marginal task value}}, \tag{5}$$

and hence the wage bill is

$$W^C(h) = \underbrace{L\lambda(L)}_{\text{cost of forgone leisure}} - \underbrace{\sum_i h_i u_i'(h_i)}_{\text{intrinsic-utility credit}} = \sum_i h_i[\lambda(L) - u_i'(h_i)]. \tag{6}$$

Thus, a task that is more valuable to the worker at the margin requires less compensating pay, holding the final bundle fixed. The term "controlled bundle" refers to control of the final task composition of paid work under voluntary hourly labour supply; it does not mean that the worker is compelled to work without an implementing compensation condition.

Paid-floor hourly contract. The paid-floor contract is intended for settings in which compensated responsibilities or deliverables are verifiable, while exact caps on self-directed effort are costly, undesirable, or difficult to enforce. Research, software tools, creative work, and professional problem-solving are natural examples. The contract is less plausible for tightly scheduled hourly work or in environments where unpaid productive work is prevented or legally constrained. The model isolates this contractual asymmetry rather than treating it as universal.

The firm posts a wage w and a paid-floor bundle $\rho \in \Delta^{n-1}$. The worker chooses paid hours $P \geq 0$ and final task hours h according to

$$\begin{aligned} \max_{P,h} \quad & wP + v\left(T - \sum_i h_i\right) + \sum_i u_i(h_i) \\ \text{subject to} \quad & \underbrace{h_i \geq \rho_i P \ \ \forall i}_{\text{paid-floor requirements}} \quad , \quad \sum_i h_i \leq T. \end{aligned} \tag{7}$$

Because $v'(s) \to +\infty$ as $s \downarrow 0$, every solution considered below has strictly positive leisure, so the boundary $\sum_i h_i = T$ is never attained.

Realised paid floors are $p_i = \rho_i P$. Thus the firm prescribes the composition of paid responsibilities while the worker selects paid hours and may supply additional, unpaid task time. By concavity of the worker's objective and linearity of the constraints, the worker Kuhn–Tucker conditions are sufficient. At an interior solution with $P > 0$, let η_i be the multiplier on $h_i \geq \rho_i P$. The compensation gap, complementarity condition, and wage condition are

$$\eta_i = \underbrace{\lambda(L) - u_i'(h_i)}_{\text{compensation gap}} \geq 0,$$

$$h_i \geq \underbrace{p_i}_{\text{paid floor}} \quad , \qquad \underbrace{\eta_i(h_i - p_i) = 0}_{\text{top-up complementarity}} \quad , \tag{8}$$

$$w = \underbrace{\sum_i \rho_i \eta_i}_{\text{wage supporting paid hours}} \quad . \tag{9}$$

Multiplying (9) by P gives

$$W^{PF}(p,h) = wP = \sum_i \underbrace{p_i}_{\text{paid floor}} \underbrace{[\lambda(L) - u_i'(h_i)]}_{\text{compensation gap}}. \tag{10}$$

Additional hours on a zero-gap task do not enter the wage bill, but they do enter worked time and output.

3 Implementability and job design

3.1 Wage-bill equivalence

We first ask which final allocations the two contracts can implement, and at what wage bill. The next lemma shows that, on the set of allocations both regimes can sustain, the wage bill is identical. The two regimes differ on the set of feasible allocations, not on the price of any given allocation.

Lemma 1 (Implementation and wage-bill equivalence). *On the maintained domain on which the implied wage is non-negative and worker participation is satisfied, let h be an interior final task allocation with $L = \sum_i h_i < T$. It is implementable by a paid-floor contract with positive paid hours if and only if*

$$\underbrace{\lambda(L) - u_i'(h_i)}_{\text{compensation gap}} \geq 0 \qquad \textit{for all } i. \tag{11}$$

For every such allocation,

$$\begin{aligned} W^{PF}(h) = W^C(h) = \quad & \underbrace{L\lambda(L)}_{\text{cost of forgone leisure}} \quad - \quad \underbrace{\sum_i h_i u_i'(h_i)}_{\text{intrinsic-utility credit}} \\ = & \sum_i h_i[\lambda(L) - u_i'(h_i)]. \end{aligned} \tag{12}$$

If every compensation gap is zero, the same allocation is implementable with $P > 0$ and $w = 0$ by assigning a positive paid floor to any one task.

Proof. If h is implemented under a paid-floor contract, (8) yields $\eta_i = \lambda(L) - u_i'(h_i) \geq 0$, establishing necessity.

For sufficiency, write $g_i = \lambda(L) - u_i'(h_i) \geq 0$. If some $g_i > 0$, set $p_i = h_i$ whenever $g_i > 0$ and $p_i = 0$ whenever $g_i = 0$. Then $P = \sum_i p_i > 0$, set $\rho_i = p_i/P$, and set $w = \sum_i \rho_i g_i \geq 0$. Tasks with $g_i > 0$ bind their floors, and tasks with $g_i = 0$ may lie above zero floors; all worker Kuhn–Tucker conditions hold. If $g_i = 0$ for every task, select any k with $h_k > 0$, set $p_k = h_k$, $p_i = 0$ for $i \neq k$, $P = h_k$, $\rho_k = 1$, and $w = 0$. Again all Kuhn–Tucker conditions hold with positive paid hours.

Finally, a task with $g_i > 0$ must bind, so $p_i = h_i$; a task with $h_i > p_i$ has $g_i = 0$. Hence $p_i g_i = h_i g_i$ task by task and $W^{PF} = \sum_i p_i g_i = \sum_i h_i g_i = W^C$. □

Definition 1 (Minimal-floor accounting convention). *For an implementable paid-floor allocation, let*

$$g_i(h) = \lambda(L) - u_i'(h_i)$$

be task i's compensation gap. The payroll observables used below are generated by the following floor-selection rule: tasks with $g_i(h) > 0$ are recorded as paid floors at their final hours, $p_i = h_i$, and tasks with $g_i(h) = 0$ are recorded with zero paid floor, $p_i = 0$, unless an external institutional rule requires a positive floor. If all compensation gaps are zero and positive paid hours are required, one retained task is assigned an arbitrary positive floor. The recorded payroll objects are then

$$P = \sum_i p_i, \qquad W = \sum_i p_i g_i(h), \qquad w = W/P \ \textit{ when } P > 0, \qquad E = L - P.$$

Statements below involving paid hours, measured hourly wages, or unpaid top-up are conditional on this observed floor-selection rule. Alternative positive floors on zero-gap tasks can change P, w, and E without changing final hours, output, worker utility, or firm surplus.

The interpretation of Lemma 1 is that voluntary top-up is not a discount on a fixed final allocation. Two forces are at work. On any allocation both contracts can implement, the worker's marginal task utility offsets the compensating wage in exactly the same way under either regime, so the wage bill coincides. By contrast, paid floors restrict the set of allocations the firm can sustain, because the firm cannot cap a task that is rewarding at the margin below the worker's voluntary level. The wedge between regimes is therefore one of implementability rather than price. The paid/unpaid split is non-unique on zero-gap tasks, and Definition 1 pins down the payroll convention used for the measured objects in the automation and identification sections.

3.2 Controlled-bundle design

Substituting (6) into operating surplus separates productive value from the wage bill needed to implement the bundle:

$$\Pi^C(h) = \underbrace{Y(h)}_{\text{output}} - \underbrace{\left[L\lambda(L) - \sum_i h_i u_i'(h_i) \right]}_{\text{implementation wage bill}}. \tag{13}$$

At an interior optimum, each task balances its marginal product and task-value contribution against the marginal implementation cost:

$$\underbrace{\frac{\beta Y(h)}{h_i}}_{\text{marginal product}} + \underbrace{u_i'(h_i) + h_i u_i''(h_i)}_{\text{task-value contribution}} - \underbrace{\left[\lambda(L) - Lv''(T-L)\right]}_{\text{marginal implementation cost}} = 0 \qquad \text{for every } i. \tag{14}$$

The next proposition records the controlled-bundle optimum and notes that, in the absence of restrictions on the shape of task utilities, the optimum does not pin down a ranking of task hours.

Proposition 1 (Controlled-bundle allocation)**.** *Under Assumptions 1–2, the controlled-bundle contract has a unique interior optimum h^C satisfying* (14). *Without additional restrictions on the functions u_i, no general ranking of the components of h^C follows from the model.*

Proof. Differentiating (13) gives (14). Strict concavity and existence in Assumption 2 make the first-order conditions sufficient and identify a unique interior optimum. Because marginal task utilities may cross and may be negative at zero hours, an ordering of task hours requires additional restrictions not imposed in the baseline model. □

The proposition isolates two forces. A productive task raises output, while its intrinsic value reduces the compensation needed to implement hours on that task. Neither force orders hours across labels unless marginal task values can be ranked. A transparent special case supplies such a comparison without imposing it on the baseline primitives.

Corollary 1 (Parallel marginal-utility subclass)**.** *Suppose $u_i(h) = \alpha_i + \theta_i h - K(h)$, where K is three times continuously differentiable and strictly convex, and, for each $Y > 0$, the function $h \mapsto \frac{\beta Y}{h} - K'(h) - hK''(h)$ is strictly decreasing. If $\theta_i > \theta_j$, then $h_i^C > h_j^C$.*

Proof. In this subclass, (14) becomes $\frac{\beta Y^C}{h_i^C} + \theta_i - K'(h_i^C) - h_i^C K''(h_i^C) = \lambda(L^C) - L^C v''(T - L^C)$. If $\theta_i > \theta_j$ but $h_i^C \leq h_j^C$, strict monotonicity in h makes the left-hand side for i strictly exceed that for j, contradicting the common right-hand side. □

The corollary permits $\theta_i < 0$ for tasks disliked at the margin from the outset and does not require all tasks to be rewarding at the margin at small hours.

3.3 Paid-floor design

By Lemma 1, the firm facing paid floors solves the same reduced surplus objective as under controlled bundles, subject to

$$\underbrace{u_i'(h_i) - \lambda(L) \leq 0 \qquad \text{for every } i}_{\text{no profitable voluntary top-up below the target}} \quad . \tag{15}$$

We first ask when this implementability constraint is slack. If the controlled-bundle optimum already lies inside the implementable set, the two regimes coincide. The next proposition records this case and the implication for the paid/unpaid split.

Proposition 2 (When paid floors do not bind)**.** *If h^C satisfies* (15)*, then h^C is also the paid-floor optimum. Final task hours, output, wage bill, and firm surplus coincide across the two hourly contracts. If $u_i'(h_i^C) < \lambda(L^C)$ for all i, every performed task hour must be included in the paid floor. If equality holds for any task, the paid/unpaid split on that task is not identified by the final allocation or the wage bill.*

Proof. The paid-floor firm maximises the same objective (13) over allocations satisfying (15). If its unconstrained unique maximiser h^C satisfies those constraints, it remains optimal. The wage-bill statement follows from Lemma 1. Positive-gap tasks must bind their floors by complementarity; zero-gap tasks may be paid or supplied above their floors without changing the wage bill. $\square$

When this condition fails, the firm faces a genuine containment constraint: it cannot prescribe fewer hours on a task than the worker would choose to supply voluntarily. The next proposition characterises that constrained problem. The formal body records the Kuhn–Tucker conditions and the boundary direction; the economic content is the trade-off between marginal product and containment cost, which we draw out immediately after.

Proposition 3 (Paid-floor Kuhn–Tucker characterisation)**.** *Let h^{PF} be an interior regular local optimum of the paid-floor programme. There exist multipliers $\nu_i \geq 0$ satisfying*

$$\underbrace{u_i'(h_i^{PF}) - \lambda(L^{PF}) \leq 0, \qquad \nu_i\Big[u_i'(h_i^{PF}) - \lambda(L^{PF})\Big] = 0,}_{\textit{implementability and complementary slackness}} \tag{16}$$

and

$$\begin{aligned} &\underbrace{\frac{\beta Y^{PF}}{h_i^{PF}} + u_i'(h_i^{PF}) + h_i^{PF} u_i''(h_i^{PF}) - \Big[\lambda(L^{PF}) - L^{PF} v''(T - L^{PF})\Big]}_{\textit{unconstrained marginal surplus}} \\ &= \underbrace{\sum_j \nu_j \Big[\mathbf{1}\{i = j\} u_j''(h_j^{PF}) + v''(T - L^{PF})\Big]}_{\textit{containment adjustment}} \qquad \forall i. \end{aligned} \tag{17}$$

For a task on the voluntary boundary, $u_i'(h_i^{PF}) = \lambda(L^{PF})$, an increase in that task is feasible to first order and therefore satisfies

$$\underbrace{\frac{\beta Y^{PF}}{h_i^{PF}}}_{\textit{marginal product}} + \underbrace{h_i^{PF} u_i''(h_i^{PF}) + L^{PF} v''(T - L^{PF})}_{\textit{private containment term at the boundary}} \leq 0. \tag{18}$$

Under the minimal-floor convention, zero-gap tasks are assigned zero paid floor, and positive-gap *tasks are assigned their full final hours, except where an institutional paid-floor requirement is imposed externally. If every retained task has zero compensation gap and positive paid hours are required, an arbitrary positive floor is assigned to one retained task, as in Lemma 1.*

Proof. The gradient of the reduced objective is the left-hand side of (17), while the derivative of the jth constraint with respect to h_i is $\mathbf{1}\{i = j\}u_j''(h_j) + v''(T - L)$. The Kuhn–Tucker theorem gives (16)–(17) at a regular local optimum. If task i is on the boundary, increasing h_i decreases $u_i'(h_i)$ and increases $\lambda(L)$, so that direction relaxes its constraint. The directional derivative of the objective cannot be positive. Substituting $u_i'(h_i^{PF}) = \lambda(L^{PF})$ into that derivative yields (18). The minimal-floor convention follows from Lemma 1. □

The interpretation of Proposition 3 is organised around two private contracting forces. The marginal product of a retained human task pushes the firm to require more hours. The private containment term pushes in the opposite direction because expanding a boundary task changes the implementing wage bill and the slackness of the paid-floor implementability constraints. Equation (18) states this firm-side trade-off at the voluntary boundary: marginal product must be offset by curvature terms in the reduced paid-floor objective. This condition is not a claim that the marginal hour is socially wasteful. A binding voluntary boundary, therefore, has two distinct implications. First, final task hours need not solve the unconstrained controlled-bundle condition task by task: the containment adjustment in (17) shifts the marginal condition on every task in the bundle, not only on the constrained one. Second, under the minimal-floor accounting convention, paid responsibilities move away from zero-gap tasks and towards positive-gap tasks. A task rewarding at the margin need not sit on a voluntary boundary: if its marginal product is sufficiently high, the firm may require hours beyond the voluntary level. The boundary is relevant precisely for tasks whose implementability constraints bind. Proposition 3 is a local characterisation; the implementable set need not be convex, so the proposition does not by itself claim a unique global paid-floor allocation.

Figure 1 illustrates the mechanism in a two-task example. Task 1 is rewarding at the relevant allocation, and task 2 is not. The voluntary boundary is the locus on which the marginal utility from task 1 equals the marginal value of leisure. Allocations in the non-implementable region ask the worker to stop task 1 before they would voluntarily stop. The controlled-bundle allocation C lies in that region. Under paid floors, task 1 expands until the implementability boundary is reached at PF, and the paid floor can be shifted away from the zero-gap task.

The planner's condition in Appendix B implies $u_i'(h_i^{FB}) < \lambda(L^{FB})$ on every productive task. Thus, the efficient allocation is strictly self-enforcing: a complete bundle contract

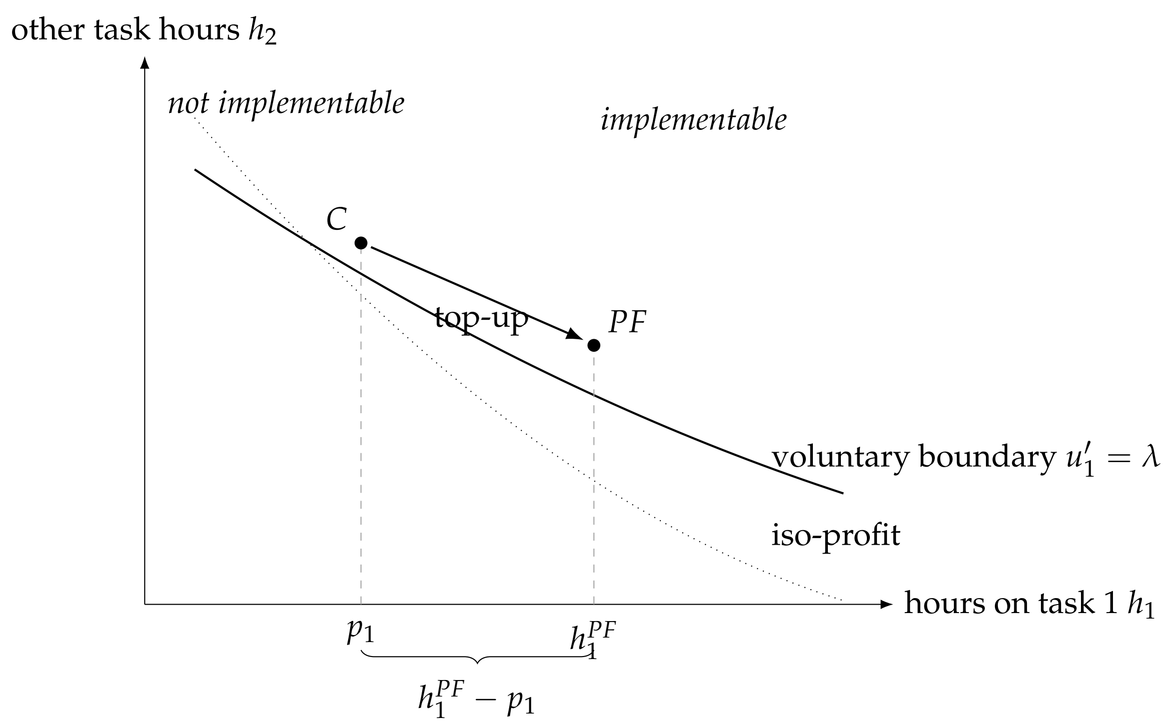

Figure 1: Mechanism in a two-task example. The controlled-bundle allocation C is not implementable under paid floors because it caps a task that is rewarding at the margin below the worker's voluntary level. The worker tops up task 1 until PF, where marginal task utility equals the marginal value of leisure. Paid minima shift away from the zero-gap task. Payroll-equivalent configurations of the kind constructed in Section 5 can differ in worked time L and unpaid top-up $E = L - P$.

can require all efficient task hours without inducing top-up. Persistent unpaid productive hours point to restricted contracts, incomplete measurement, norms, market power, or institutional rules governing paid time, rather than to the first-best allocation itself.

4 Automation: short run versus long run

We now introduce an AI technology. The timing distinction is between a short run, in which retained human paid requirements are inherited from the pre-automation job, and a long run, in which the firm may redesign the paid task bundle and select retained human tasks. The short-run paid-floor case treats the realised pre-automation paid obligations as fixed institutional or scheduling commitments: payroll systems, rosters, project assignments, or legal rules carry forward the realised floors p^0 until redesign occurs. It is therefore not a fresh optimisation of paid hours under the old scalable pair (w, ρ). Crossing timing with contractual regime gives the 2×2 grid in Figure 2. Under controlled bundles, paid and worked time coincide. Under paid floors, the worker may top up retained tasks; whether the wedge widens following redesign is an outcome, not an assumption.

4.1 The automation technology

Let $\mathcal{A} \subseteq \mathcal{N}$ be the set of automated tasks and let $\mathcal{H} = \mathcal{N} \setminus \mathcal{A}$ be the retained human-task set. For $j \in \mathcal{A}$, AI input is $a_j \geq 0$, effective task hours are $q_A a_j$, with $q_A > 0$, and input costs ca_j, with $c > 0$. Output is

$$Y(h, a; \mathcal{A}) = \underbrace{\prod_{i \in \mathcal{H}} h_i^{\beta}}_{\text{retained human-task inputs}} \cdot \underbrace{\prod_{j \in \mathcal{A}} (q_A a_j)^{\beta}}_{\text{automated-task inputs}} . \tag{19}$$

Let $m = |\mathcal{A}|$. Provided $m\beta < 1$, the AI first-order conditions give

$$a_j^* = \underbrace{\frac{\beta Y}{c}}_{\text{optimal automated input}} \qquad (j \in \mathcal{A}), \tag{20}$$

and, conditional on retained human hours,

$$Y_{\mathcal{A}}(h) = \Bigg[\underbrace{\prod_{i \in \mathcal{H}} h_i^{\beta}}_{\text{retained human contribution}} \cdot \underbrace{q_A^{m\beta} \left(\frac{\beta}{c}\right)^{m\beta}}_{\text{AI productivity–cost contribution}} \Bigg]^{1/(1-m\beta)} . \tag{21}$$

Optimised output net of AI input cost is $(1 - m\beta) Y_{\mathcal{A}}(h)$. By the envelope theorem,

$$\frac{\partial}{\partial h_i} \left[Y_{\mathcal{A}}(h) - c \sum_{j \in \mathcal{A}} a_j^*(h) \right] = \underbrace{\frac{\beta Y_{\mathcal{A}}(h)}{h_i}}_{\text{marginal value of a retained human task}} . \tag{22}$$

Accordingly, automation changes the scale of the marginal product but leaves the form of retained-human first-order expressions unchanged: Y is replaced by $Y_{\mathcal{A}}$.

	Controlled bundle	Paid floors
Short run	CELL 1 Retained human hours inherited: $h_i = h_i^0$ $L = P$ (no wedge)	CELL 3 Inherited realised floors p_i^0; worker tops up rewarding tasks $L \geq P$ (wedge opens)
Long run	CELL 2 Firm reoptimises required hours $L = P$ (no wedge)	CELL 4 Firm redesigns floors around top-up $L \geq P$ (wedge may change)

Figure 2: The four cells of Section 4, crossed by contractual regime (columns) and horizon (rows). Under controlled bundles, paid and worked hours coincide in both horizons. Under paid floors, the worker can top up retained tasks that are rewarding at the margin already in the short run; in the long run, the firm reoptimises which tasks carry the minima and which minima bind, and the unpaid-hours wedge may widen, shrink, or vanish.

4.2 The 2×2 grid

The four cells identify which features of the job are inherited on impact and which may be redesigned after adjustment.

Cell 1: controlled bundle, short run. Retained human hours are pinned at their pre-automation values h_i^0 for $i \in \mathcal{H}$. Paid and worked time coincide, $L^{C,SR} = P^{C,SR} = \sum_{i \in \mathcal{H}} h_i^0$, and the wage is $w^{C,SR} = \lambda(L^{C,SR}) - \sum_{i \in \mathcal{H}} (h_i^0 / L^{C,SR}) u_i'(h_i^0)$. Removing a task with a relatively low marginal task value raises average marginal task utility in the retained bundle and tends to reduce required compensation; lower retained human work also raises leisure and tends to lower λ.

Cell 2: controlled bundle, long run. The firm reoptimises retained human hours. Each retained task satisfies

$$\underbrace{\frac{\beta Y_{\mathcal{A}}(h)}{h_i}}_{\text{automation-adjusted marginal product}} + \underbrace{u_i'(h_i) + h_i u_i''(h_i)}_{\text{task-value contribution}} = \underbrace{\lambda(L) - L v''(T - L)}_{\text{marginal implementation cost}}, \qquad i \in \mathcal{H}, \tag{23}$$

together with (20). Automation changes both the scale of output and, when selection is endogenous, membership in $\mathcal{H}$.

Cell 3: paid floors, short run. The short-run exercise freezes the realised pre-automation paid-floor vector p^0, rather than offering a newly optimised scalable paid-floor schedule. This is a lock-in assumption about observed paid obligations: retained floors are inherited as absolute minima until the firm can redesign the paid-floor bundle. Given inherited realised floors p_i^0, the worker chooses voluntary top-up and final task hours satisfy

$$\begin{gathered}\underbrace{h_i^{PF,SR} \geq p_i^0}_{\text{inherited paid floor}}, \qquad \underbrace{\lambda(L^{PF,SR}) - u_i'(h_i^{PF,SR})}_{\text{compensation gap}} \geq 0, \\ \underbrace{(h_i^{PF,SR} - p_i^0)\,[\lambda(L^{PF,SR}) - u_i'(h_i^{PF,SR})] = 0}_{\text{top-up complementarity}}, \qquad i \in \mathcal{H}.\end{gathered} \tag{24}$$

Paid hours are $P^{PF,SR} = \sum_{i\in\mathcal{H}} p_i^0$, whereas worked time is $L^{PF,SR} = \sum_{i\in\mathcal{H}} h_i^{PF,SR}$. Removing compensated tasks can lower paid hours while top-up on retained tasks leaves worked time less responsive, or even raises it in some parameter regions.

Cell 4: paid floors, long run. The firm redesigns around implementability. Conditional on $\mathcal{A}$, it chooses a final retained-human allocation to maximise the AI-optimised reduced objective subject to $u_i'(h_i) \leq \lambda(L)$ for $i \in \mathcal{H}$. Retained human tasks are characterised locally by Proposition 3, using $Y_{\mathcal{A}}$. Under the minimal-floor accounting convention in Definition 1, the paid and unpaid components are

$$\begin{aligned} P^{PF,LR} &= \underbrace{\sum_{i\in\mathcal{H}:\, u_i'(h_i^{PF,LR})<\lambda(L^{PF,LR})} h_i^{PF,LR}}_{\text{positive-gap hours included in paid floors}}, \\ E^{PF,LR} &= \underbrace{\sum_{i\in\mathcal{H}:\, u_i'(h_i^{PF,LR})=\lambda(L^{PF,LR})} h_i^{PF,LR}}_{\text{zero-gap voluntary hours}}. \end{aligned} \tag{25}$$

If every retained task has zero compensation gap and positive paid hours are required, an arbitrary positive paid floor is assigned to one retained task, as in Definition 1. Redesign can shift paid responsibilities away from zero-gap tasks. It does not generally imply $E^{PF,LR} \geq E^{PF,SR}$, because redesign can also reduce retained hours or automate a task that would otherwise be topped up.

Figure 3 depicts one representative path in which the paid–worked-hours wedge opens and then grows under paid-floor redesign. It is an illustration of the mechanism, not a general comparative-static sign result.

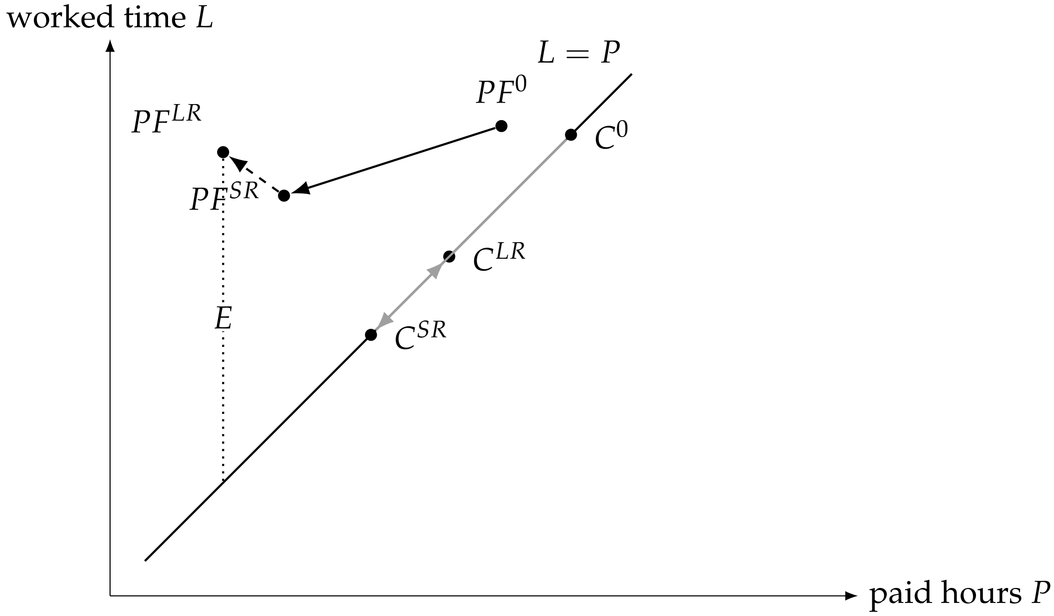

Figure 3: Illustrative paths in (P, L) space. Controlled-bundle outcomes C remain on the 45-degree line because paid and worked hours coincide. Paid-floor outcomes PF may move above that line when workers add unpaid top-up to retained tasks. Superscripts 0, SR, and LR denote pre-automation, short-run, and long-run outcomes. The dotted segment at PF^{LR} is the unpaid-hours wedge $E = L - P$.

4.3 Automation selection: replacement, scale, and containment

Without restrictions on task utility, there is no general statement that automation selects a tail of tasks ranked by enjoyment. For a restricted comparison, say that task i ranks above task j if $u_i'(x) \geq u_j'(x)$ at every retained-hour level relevant to a proposed swap. A total marginal-retention ranking is a relabelling for which this comparison holds whenever $i < j$.

Proposition 4 (Automation swaps, replacement, and containment)**.** *Fix the number m of automated tasks and suppose technology and AI costs are symmetric across task labels. Call a label swap* admissible *if the relabelled allocation remains in the maintained domain: the implementing wage is non-negative, worker participation is satisfied, and, under paid floors, the final allocation also satisfies the paid-floor implementability inequalities.*

1. *Under a controlled-bundle hourly contract, let an optimum automate task i and retain task j at human hours x. An admissible swap of their labels, keeping all quantities fixed, raises the firm's objective by $x[u_i'(x) - u_j'(x)]$. Under a total marginal-retention ranking, a tail optimum exists if every inversion-removing swap is admissible; with strict inequalities, no optimum contains a strict admissible inversion.*

2. *Under a paid-floor contract, the same comparison holds for admissible swaps. Therefore a tail optimum exists under a total marginal-retention ranking whenever every inversion-removing swap preserves the maintained domain and paid-floor implementability. If a swap is inadmissible because participation would fail or because retaining the task with higher marginal value at those hours would violate $u_i'(x) \leq \lambda(L)$, no tail conclusion follows. In the latter case, automation of that task can serve a private containment role.*

Proof. For a fixed number of automated tasks, swapping task labels preserves effective inputs, output, and AI cost. On an admissible controlled-bundle swap, the only changed term in reduced surplus is the retained task's wage-bill-saving contribution $xu_k'(x)$, yielding the stated difference. Iteratively removing admissible weak inversions produces a tail optimum when every inversion-removing swap is admissible, and strict improvements rule out strict admissible inversions. The comparison is not asserted for inadmissible swaps: participation depends on wage income and the level of task utility, not only on the marginal value $u_i'(x)$.

Under paid floors, Lemma 1 gives the same reduced objective on the implementable set, again restricted to the maintained participation and non-negative-wage domain. If the swapped allocation remains admissible, the same payoff comparison applies. If it is inadmissible because the retained task with higher marginal value at the displaced task's hours violates $u_i'(x) \leq \lambda(L)$, the swap argument is unavailable. This implementability loss is the containment channel. □

To make the selection role of containment explicit, define the reduced value of an automated set $\mathcal{A}$ before the paid-floor implementability restriction as

$$R_{\mathcal{A}}(h) \equiv Y_{\mathcal{A}}(h) - c \sum_{j \in \mathcal{A}} a_j^*(h) - \left[L\lambda(L) - \sum_{i \in \mathcal{H}} h_i u_i'(h_i) \right], \tag{26}$$

where $L = \sum_{i \in \mathcal{H}} h_i$. Let

$$V^C(\mathcal{A}) = \max_h R_{\mathcal{A}}(h), \qquad V^{PF}(\mathcal{A}) = \max_{h:\, u_i'(h_i) \leq \lambda(L)\ \forall i \in \mathcal{H}} R_{\mathcal{A}}(h),$$

with both maxima taken over the maintained non-negative-wage and participation domain, and let $D(\mathcal{A}) = V^C(\mathcal{A}) - V^{PF}(\mathcal{A}) \geq 0$ be the paid-floor containment loss.

Proposition 5 (Strict containment selection)**.** *Fix a common set of already automated tasks B and two remaining tasks i and j. Let $\mathcal{A}_i = B \cup \{i\}$ denote the set that automates i, and let*

$\mathcal{A}_j = B \cup \{j\}$ denote the set that automates j. Then

$$V^{PF}(\mathcal{A}_i) - V^{PF}(\mathcal{A}_j) = \left[V^C(\mathcal{A}_i) - V^C(\mathcal{A}_j)\right] + \left[D(\mathcal{A}_j) - D(\mathcal{A}_i)\right]. \tag{27}$$

If the controlled-bundle value favours retaining i, $V^C(\mathcal{A}_j) > V^C(\mathcal{A}_i)$, but

$$D(\mathcal{A}_j) - D(\mathcal{A}_i) > V^C(\mathcal{A}_j) - V^C(\mathcal{A}_i),$$

then the paid-floor firm strictly prefers automating i. The selection is driven by the larger containment loss from retaining i.

Proof. The identity follows by substituting $V^{PF}(\mathcal{A}) = V^C(\mathcal{A}) - D(\mathcal{A})$ for $\mathcal{A}_i$ and $\mathcal{A}_j$. The strict-preference condition is the inequality $V^{PF}(\mathcal{A}_i) > V^{PF}(\mathcal{A}_j)$ written using (27). □

The interpretation of Propositions 4 and 5 runs through three named selection margins. The replacement and scale margins are visible under controlled bundles: holding the number of automated tasks fixed, the firm tends to retain tasks whose human performance carries a higher marginal task value, because each retained unit of such a task saves on the wage bill needed to implement it. Under a total marginal-retention ranking and admissible inversion-removing swaps, this delivers a tail optimum in which lower-ranked tasks are automated first. The containment margin is specific to paid floors. When retaining a task creates a large implementability loss because its voluntary boundary binds, automation can be optimal even if the controlled-bundle value would favour retaining that task.

The replacement and scale effects are visible in the optimised operating surplus. If a total marginal-retention ranking justifies indexing candidate automated sets by the number m of automated lower-ranked tasks, the gain from adding an automated task decomposes as

$$\begin{aligned}\Delta\Pi^C(m) = \; & \underbrace{Y_m - Y_{m-1}}_{\text{output and scale effect}} \; - \; \underbrace{[W_m^C - W_{m-1}^C]}_{\text{replacement effect on compensation}} \\ & - \underbrace{c\,[m a_m - (m-1)a_{m-1}]}_{\text{incremental AI cost}},\end{aligned} \tag{28}$$

where the last term includes reoptimisation of input on previously automated tasks. Under paid floors, containment is an additional private reason that a technically symmetric task may be automated. It concerns the firm's inability to cap a retained task at the reduced-surplus optimum, not an assumption that additional task output is intrinsically harmful or socially wasteful. Its cost arises through the implementation of the wage bill and the active paid-floor constraints. Task-specific technical suitability, verification cost, risk, or regulation would add further selection margins outside this benchmark.

5 Payroll non-identification

The central empirical point is that payroll records commonly reveal the hourly wage, wage bill, and paid hours, but do not usually reveal total worked time, voluntary top-up, or within-worker task composition. Replacement of compensated tasks and voluntary expansion of retained tasks leave footprints on different margins. Replacement can appear in paid hours and the wage bill; voluntary expansion appears in worked time, unpaid top-up, and task mix. Since zero-gap voluntary hours enter neither P nor W directly, observationally similar payroll records can conceal economically different job allocations. Since $w = W/P$, the payroll triplet contains only two independent observables. We retain all three entries because wage rates, paid hours, and wage bills are commonly reported separately.

Consider two post-automation configurations with the same automated task and the same AI technology. In Configuration A, both retained human tasks are included in the paid floor. In Configuration B, one retained task that is rewarding at the margin lies on the voluntary boundary and is supplied beyond a zero paid floor, while the other retained task carries the paid floor. A labour economist observing only (w, W, P) sees the same payroll footprint in the two configurations. Time-use or task-level data reveal different worked time, different unpaid time, and a different task mix.

Proposition 6 (Payroll-only non-identification under the minimal-floor convention)**.** *Conditional on an observed automated set $\mathcal{A}$, a common AI technology (q_A, c), and the minimal-floor accounting convention in Definition 1, there exist two regular local paid-floor configurations with the same time endowment T, leisure utility v, production exponent β, and quadratic curvature parameter κ, such that the observed payroll vector (w, W, P) is identical while total human worked time L, unpaid time $E = L - P$, and retained task shares differ. The two configurations differ in latent marginal task-value parameters θ_i. The construction matches payroll observables only: output, realised AI-input use, pre-period behaviour, and worker identity are not held fixed.*

Proof. Appendix C supplies an explicit construction with two retained human tasks and one common automated task. Both configurations use $T = 10$, $v(s) = 8 \log s$, $\beta = 0.2$, $\kappa = 0.001$, $c = 1$, and $q_A = 111.389724$. Configuration A has final human hours $(1.2, 0.8)$ and paid floors equal to final hours. Configuration B has final human hours $(4, 2)$ and paid floors $(0, 2)$ under Definition 1. The appendix selects task-utility parameters so that implementability and the firm's regular local optimality conditions hold under the common AI technology. The observed payroll values coincide: $P = 2$, $W = 0.358114$, and $w = 0.179057$. The non-payroll outcomes differ: $(L^A, E^A) = (2, 0)$ and $(L^B, E^B) = (6, 4)$, the retained task shares are $(0.6, 0.4)$ and $(2/3, 1/3)$, output is $Y^A = 2.150486$ and $Y^B =$

	Configuration A	Configuration B
Observed in payroll under Definition 1		
Hourly wage w	0.179057	0.179057
Wage bill W	0.358114	0.358114
Paid hours P	2	2
Not identified by payroll alone		
Worked time L	2	6
Unpaid top-up $E = L - P$	0	4
Task shares	$(0.6, 0.4)$	$(2/3, 1/3)$
Output Y	2.150486	3.653771
AI input $a^* = \beta Y/c$	0.430097	0.730754

Table 2: Two post-automation configurations with a common AI technology and the same payroll footprint under the minimal-floor convention. Output and realised AI-input use differ, so those observables would separate the configurations.

3.653771, and realised AI input is $a^{A,*} = 0.430097$ and $a^{B,*} = 0.730754$ because $a^* = \beta Y/c$. This establishes the payroll-only claim. □

The proposition is pointwise and conditional on the observed automated set and on the minimal-floor accounting convention. It is a cross-sectional or latent-structure non-identification statement: it does not imply that a fixed worker with fixed preferences can transition between the two allocations while leaving payroll unchanged. It allows latent marginal task-value parameters, which payroll does not record, to differ across configurations. It also does not establish endogenous selection of the observed automated set and does not match output or realised AI-input use. Observing output, AI-input use or expenditure, worked time, task shares, worker identity, pre-period behaviour, or worker well-being can break the equivalence.

5.1 Three margins that do separate the configurations

Three types of data can separate the configurations.

The sign and magnitude of changes in $L - P$. Replacement can reduce paid and worked time together; voluntary expansion can raise $E = L - P$ even when paid hours are flat. Testing this margin requires time-use or self-report data on actual working time, not payroll alone. The qualitative evidence in Ranganathan and Ye (2026) and time-use evidence in Jiang et al. (2025) speak directly to this margin.

Within-worker task mix. A retained task mix may shift towards tasks with higher marginal intrinsic value for different reasons. Under replacement, the shift is driven by subtracting compensated tasks from the bundle. Under voluntary expansion, it is driven by the worker pushing a retained task rewarding at the margin past its paid floor. The within-worker evidence in Hoffmann et al. (2025), in which developers with Copilot access spend more time on core coding and less on auxiliary project-management tasks, is evidence on this margin. Distinguishing subtraction from expansion requires task-level hours or activity measures.

Worker-reported intensity and well-being. Configurations with the same payroll vector can differ in total hours and in the utility value of retained tasks. Worker-reported intensity and well-being may therefore differ across them. Meaning-of-work measures in Nikolova and Cnossen (2020), willingness-to-pay estimates for working conditions in Maestas et al. (2023), and retention or sentiment measures in Brynjolfsson et al. (2025) all sit on this margin.

6 Salaried compensation

A salaried contract can price a complete job bundle rather than using an hourly wage to implement labour supply at the margin. Let a required human task vector h be offered with fixed salary S. If the worker's outside utility is $\bar{U}$, the minimum salary required for acceptance separates the participation requirement from the non-wage value of the bundle:

$$S^s(h) = \underbrace{\bar{U}}_{\text{outside utility}} - \underbrace{\left[v(T-L) + \sum_i u_i(h_i) \right]}_{\text{non-wage utility from the bundle}}, \tag{29}$$

so

$$\frac{\partial S^s(h)}{\partial h_i} = \underbrace{\lambda(L)}_{\text{leisure opportunity cost}} - \underbrace{u_i'(h_i)}_{\text{marginal task value}}. \tag{30}$$

We now ask what a salaried contract that prices the complete bundle can achieve. The next proposition shows that exact bundle pricing implements the first-best human allocation, removes the voluntary-hours wedge at the efficient bundle, and, under an incremental-utility dominance order, restores a tail conclusion for automation.

Proposition 7 (Bundle pricing and the first best)**.** *Under Assumptions 1–2, exact bundle pricing* *by salary implements the first-best human allocation. At any productive first-best allocation,*

$$\underbrace{\frac{\beta Y^{FB}}{h_i^{FB}}}_{\text{marginal product}} + \underbrace{u_i'(h_i^{FB})}_{\text{marginal task value}} = \underbrace{\lambda(L^{FB})}_{\text{leisure opportunity cost}} \qquad \forall i, \tag{31}$$

and therefore $u_i'(h_i^{FB}) < \lambda(L^{FB})$ for every productive task. The same allocation is implementable by paid floors with every required hour paid, so voluntary top-up vanishes at the efficient allocation.

Under automation, suppose that for two tasks i and j, the intrinsic-utility increment from human performance satisfies $u_i(x) - u_i(0) \geq u_j(x) - u_j(0)$ at every retained human hour x relevant to a swap. If task i is automated while task j is retained, swapping their labels weakly lowers the required salary without changing output or AI cost. An incremental-utility dominance order therefore implies existence of a tail automated set under rich salaried bundle pricing.

Proof. The salaried firm maximises $Y(h) - S^s(h)$. Differentiating using (30) yields (31), which is also the planner's first-order condition for maximising $Y(h) + v(T - L) + \sum_i u_i(h_i)$. Strict concavity gives uniqueness where asserted. Since $\beta Y^{FB}/h_i^{FB} > 0$ on productive tasks, (31) implies $u_i'(h_i^{FB}) < \lambda(L^{FB})$. Under a salaried floor contract that offers the same fixed transfer conditional on supplying floors $p_i = h_i^{FB}$, those strict inequalities ensure that the accepted worker has no incentive to add voluntary time; the first-best bundle is therefore self-enforcing.

With automation, a label swap keeps production and AI cost unchanged. Relative to zero human hours on the automated task, retaining task i rather than j increases intrinsic utility by $[u_i(x) - u_i(0)] - [u_j(x) - u_j(0)] \geq 0$ and hence weakly reduces the salary required for acceptance. Iterating swaps gives the tail conclusion under an incremental-utility dominance order.[1] □

The interpretation is straightforward. Under a salary, the firm prices a complete bundle rather than financing the labour-supply margin hour by hour. The hourly wage-bill distortion disappears, and so does voluntary top-up at the efficient bundle. Salary alone is not sufficient: the benchmark requires rich bundle pricing and effective control of the task bundle. Statutory rules, collective agreements, monitoring limitations, or professional norms may prevent such implementation even for salaried workers.

Occupations differ in their proximity to this benchmark. Where required hours and tasks are easily specified and monitored, the paid-worked-hours wedge generated by this mechanism should be small. Where task effort is delegated, and tasks rewarding at the

[1] For example, in the parallel family $u_i(h) = \alpha_i + \theta_i h - K(h)$ with ordered θ_i, marginal and incremental-utility dominance are both satisfied; the constants α_i cancel from utility increments.

margin can be pursued beyond paid responsibilities, AI adoption can be associated with a larger wedge. The prediction is conditional: adoption changes the wedge only when retained task values and contractual institutions make voluntary top-up operative.

7 Discussion

Marginal task values and worker heterogeneity. Our intensive-margin allocations depend on $u_i'(h_i)$, not on utility levels alone. Empirically, both objects may vary across workers. Sorting and selection, therefore, matter: payroll invariance may reflect a change in worker composition as well as top-up by a given worker. We isolate the within-job mechanism rather than providing a full labour-market incidence model.

Which jobs fit the paid-floor contract? The scope of the paid-floor benchmark is deliberately limited. As discussed in Section 2, the mechanism applies when compensated responsibilities can be specified but self-directed effort on a retained activity is difficult to cap. Empirical applications should therefore distinguish delegated knowledge work from settings in which unpaid productive work is prevented or legally constrained.

Intrinsic utility is one mechanism for additional work. Workers may spend unrecorded time on work because they enjoy it, but also because of promotion incentives, deadlines, professional norms, fear of displacement, or managerial pressure. The present model isolates intrinsic task utility. Dynamic incentives and coercive or quasi-coercive unpaid labour require distinct modelling and empirical tests.

Technology and task structure. O-ring production makes complementarity transparent, but it is a special case: AI can substitute for, augment, create, or eliminate tasks. Similarly, the symmetric AI technology in Section 4 isolates enjoyment-based selection and abstracts from task-specific technical automability, verification, risk, and compliance cost. The automation-cutoff results should therefore be read as conditional benchmark results.

Identifying evidence. Payroll alone cannot distinguish required effort from voluntary top-up in Proposition 6. The result is deliberately limited to payroll observables under the stated floor-accounting convention. Time-use measures, task logs, output, AI-input use or expenditure, worker identity, pre-period behaviour, worker-reported intensity and well-being, retention, and explicit restrictions on task preferences can separate the mechanisms. The empirical prediction is not that unpaid hours must rise with AI; it is that a paid-worked-hours wedge can arise and may move differently across occupations depending on contracting institutions and retained task values.

Boundary behaviour of automation. As $c \to \infty$, automated input becomes prohibitively costly. A loss of AI's productivity advantage alone, such as $q_A \to 1$, does not generally eliminate adoption: automation may remain profitable if its unit cost is sufficiently low relative to the human implementation cost. Full automation can become attractive as AI cost becomes sufficiently small, subject to any modelled role for human participation or human-only task value.

8 Conclusion

Intrinsic task utility can alter the interpretation of observed hours and earnings after automation. Under a scalable paid-floor hourly contract, workers can add unpaid time to tasks that are attractive at the margin. The resulting wedge changes implementable task allocations, not the wage-bill cost of a fixed implementable allocation. Automation can remove compensated tasks, increase the value of retained human inputs, or avoid the private paid-floor loss from retaining a task whose voluntary boundary binds.

We do not impose a universal ranking of task enjoyment and do not predict that automation must widen unpaid hours. Instead, we identify the contracting and preference conditions under which that outcome can occur. Conditional on a common observed automated set, a common AI technology, and the stated floor-accounting convention, payroll variables alone need not identify total worked time or task mix when task utilities are latent; output, AI-input use, and worker-level panel information can break that equivalence. Rich salaried bundle pricing provides the limiting benchmark in which the hourly implementation distortion and the voluntary-hours wedge disappear at the first best.

Appendices

A Supporting derivations

A.1 Second-order terms in the controlled-bundle problem

For (13), the Hessian has diagonal entries $\partial^2\Pi^C/\partial h_i^2 = \beta(\beta-1)Y/h_i^2 + 2v''(T-L) - Lv'''(T-L) + 2u_i''(h_i) + h_i u_i'''(h_i)$ and off-diagonal entries $\partial^2\Pi^C/\partial h_i \partial h_j = \beta^2 Y/(h_i h_j) + 2v''(T-L) - Lv'''(T-L)$ for $i \neq j$. Under the sufficient curvature restrictions following Assumption 2, the task-utility and hourly wage-bill terms do not overturn the concavity of Cobb–Douglas production.

A.2 Quadratic illustration of paid-floor design

Under the quadratic specification $u_i(h) = \alpha_i + \theta_i h - \frac{\kappa}{2}h^2$, with $\kappa > 0$, task i has marginal task utility $u_i'(h) = \theta_i - \kappa h$, and its contribution to the controlled-bundle marginal condition is $\theta_i - 2\kappa h$. If $\theta_i > \lambda(L)$, its voluntary boundary conditional on total worked time is $h_i^V(\lambda) = (\theta_i - \lambda)/\kappa$. If $\theta_i \leq \lambda(L)$, there is no positive voluntary lower bound at that value of leisure. The parameter θ_i may be negative, so the subclass permits tasks disliked from the first hour. Under the parallel-cost restriction with ordered θ_i, higher-θ_i tasks have higher controlled-bundle hours and, where positive, higher voluntary lower bounds. Without that restriction, task-hour rankings need not be monotone and marginal utility curves may cross.

A.3 AI-optimised human Hessian

With m automated tasks, optimised output net of AI cost is $(1-m\beta)Y_{\mathcal{A}}(h)$. Its human marginal derivative is $\beta Y_{\mathcal{A}}/h_i$, its diagonal second derivative is $\frac{\beta Y_{\mathcal{A}}}{h_i^2}\left(\frac{\beta}{1-m\beta} - 1\right)$, and its off-diagonal second derivative is $\frac{\beta^2 Y_{\mathcal{A}}}{(1-m\beta)h_i h_j}$ for $i \neq j$. These expressions are used in Appendix C.

A.4 Impacts on wages, hours, income, and utility

Under controlled bundles, the hourly wage is $w^C = \lambda(L^C) - \sum_{i\in\mathcal{H}}(h_i/L^C)u_i'(h_i)$. Automation affects this wage through a leisure force and a composition force. If retained human worked time falls, leisure rises, and λ falls. If automation removes tasks with low marginal task value relative to retained tasks, average marginal task utility rises, and the compensating wage tends to fall. In the long run, a productivity-scale effect can push retained human hours upward, raising λ and potentially reversing the wage movement.

Under paid floors, the measured hourly wage is $w^{PF} = W^{PF}/P^{PF}$. It can move differently from total income or total worked time. If redesign reduces paid floors on zero-gap tasks while leaving the wage bill unchanged, P^{PF} falls and measured w^{PF} rises mechanically. If automation removes tasks with positive compensation gaps, W^{PF} falls. Measured hourly wages are, therefore, especially incomplete statistics for the paid-floor regime.

For any final allocation implementable under both hourly contracts, Lemma 1 implies the same wage bill. Differences in worker income across regimes arise because the contracts can lead to different final allocations, not because one prices the same allocation more cheaply. At a regime-specific allocation, worker utility is $W + v(T - L) + \sum_{i \in \mathcal{H}} u_i(h_i)$, and operating surplus is $Y - c \sum_{j \in \mathcal{A}} a_j - W$.

Outcome	Controlled bundle, SR	Controlled bundle, LR	Paid floors, SR	Paid floors, LR	Empirical link
Hourly wage	tends downward if lower-value compensated tasks are removed.	Ambiguous once scale effects feed back into hours.	Ambiguous because W and P can move differently.	Ambiguous; floor redesign can raise $w = W/P$ even as paid hours fall.	Earnings
Paid hours P	Down mechanically on automated inherited tasks.	Ambiguous.	Down if automated tasks carried positive paid floors.	May fall further as floors move away from zero-gap tasks; no universal sign.	Paid time
Worked time L	Down mechanically in the inherited-hours short run.	Ambiguous.	May fall by less than P, remain flat, or rise.	May remain high, fall, or rise depending on retention and top-up.	Intensity
Task mix	Tasks removed by subtraction.	Reoptimised retained mix.	May shift towards retained tasks with higher marginal value.	May shift through replacement, redesign, or containment.	Task mix

Table 3: Conditional comparative statics of automation under the two hourly contracts. "Ambiguous" records that replacement, scale, redesign, and containment can offset one another.

B Welfare benchmarks

B.1 First-best allocation

The total surplus from a human task allocation before automation is $\mathcal{S}(h) = Y(h) + v(T - L) + \sum_i u_i(h_i)$. Its interior first-order condition balances productive value and intrinsic task value against the value of leisure:

$$\underbrace{\frac{\beta Y^{FB}}{h_i^{FB}}}_{\text{marginal product}} + \underbrace{u_i'(h_i^{FB})}_{\text{marginal task value}} = \underbrace{\lambda(L^{FB})}_{\text{leisure opportunity cost}} \quad \text{for all } i. \tag{32}$$

At the efficient allocation, the marginal product of a task plus the worker's own marginal task utility equals the marginal value of leisure. Because $\beta Y^{FB}/h_i^{FB} > 0$ for every productive task, one has $u_i'(h_i^{FB}) < \lambda(L^{FB})$.

B.2 Local social inefficiency under controlled-bundle hourly implementation

Evaluating the social marginal value at the controlled-bundle hourly allocation and using (14),

$$\begin{aligned} \left.\frac{\partial \mathcal{S}}{\partial h_i}\right|_{h=h^C} &= \frac{\beta Y^C}{h_i^C} + u_i'(h_i^C) - \lambda(L^C) \\ &= \Big[\lambda(L^C) - L^C v''(T - L^C)\Big] - \Big[u_i'(h_i^C) + h_i^C u_i''(h_i^C)\Big] + u_i'(h_i^C) - \lambda(L^C) \\ &= \underbrace{-L^C v''(T - L^C)}_{\text{leisure-curvature wedge}} + \underbrace{-h_i^C u_i''(h_i^C)}_{\text{task-utility-curvature wedge}} > 0. \end{aligned} \tag{33}$$

Thus, the controlled-bundle hourly allocation is locally inefficient in every positive task direction: holding other task hours fixed, a marginal increase in any task raises total surplus. Without additional monotone comparative-static conditions, this local result does not assert $h_i^{FB} > h_i^C$ coordinate by coordinate.

B.3 The meaning of unpaid time under paid floors

At a paid-floor voluntary boundary, $u_i'(h_i) = \lambda(L)$. Holding other task hours fixed, the social marginal value of increasing that task is therefore $\beta Y/h_i$ before automation, or $\beta Y_{\mathcal{A}}/h_i$ conditional on an automated set. It is positive for a productive task. The boundary condition in Proposition 3 is consequently not a total-surplus test. It is a firm-side condition in the reduced hourly-contract problem, where expanding a boundary task also changes implementation costs and active paid-floor constraints.

In a paid-floor equilibrium, unpaid productive time can be real in payroll data: $L^{PF} > P^{PF}$. But it is not a free input in equilibrium. It arises on tasks where the compensation gap is zero, so the worker's marginal intrinsic utility exactly offsets forgone leisure. If firms can use complete bundle contracts with fixed transfers, the value of those hours is capitalised into the compensation package and the first-best bundle is implemented. Persistent unpaid time therefore points to restricted contracts, incomplete measurement, professional norms, labour-market power, or institutional rules governing paid work, rather than to a failure of worker optimisation.

C Common-technology construction for Proposition 6

Take two retained human tasks and one automated task, so $m = 1$, and set $T = 10$, $v(s) = 8 \log s$, $\beta = 0.2$, and $\kappa = 0.001$, with quadratic task utilities $u_i(h) = \theta_i h - \kappa h^2/2$. Both configurations automate the same third task and share the AI technology $c = 1$ and $q_A = 111.389724$, so $\gamma \equiv q_A/c = 111.389724$. With one automated task, (21) implies $Y(h_1, h_2) = [\beta\gamma h_1 h_2]^{\beta/(1-\beta)}$.

Configuration A: two paid retained tasks. Let $(h_1^A, h_2^A) = (1.2, 0.8)$ and $L^A = P^A = 2$, and set paid floors equal to final hours. The output expression gives $Y^A = 2.150486$. Selecting $\theta_1^A = 0.893986$ and $\theta_2^A = 0.713979$ produces compensation gaps $g_1^A = \lambda(L^A) - u_1'(h_1^A) = 0.107214$ and $g_2^A = 0.286821$. Both gaps are positive, and the retained-human first-order conditions hold up to displayed rounding. Payroll variables are $W^A = 1.2g_1^A + 0.8g_2^A = 0.358114$ and $w^A = W^A/P^A = 0.179057$. Using the Hessian expressions above, the Hessian of the AI-optimised reduced objective is approximately $\begin{pmatrix} -0.538509 & -0.200496 \\ -0.200496 & -0.818520 \end{pmatrix}$, with eigenvalues -0.433974 and -0.923055. Hence, Configuration A is a strict local optimum conditional on the observed automated set.

Configuration B: one voluntary retained task. Let $(h_1^B, h_2^B) = (4, 2)$, $L^B = 6$, $(p_1^B, p_2^B) = (0, 2)$, and $P^B = 2$. The same value of γ gives $Y^B = 3.653771$. Now $\lambda(L^B) = 2$ and $-v''(T - L^B) = 0.5$. Choose $\theta_1^B = 2.004000$ and $\theta_2^B = 1.822943$. Then task 1 is on the voluntary boundary, $u_1'(4) = 2 = \lambda(L^B)$, while task 2 has positive compensation gap $g_2^B = \lambda(L^B) - u_2'(2) = 0.179057$. Therefore $W^B = 2g_2^B = 0.358114$ and $w^B = W^B/P^B = 0.179057$, matching Configuration A.

For the active voluntary constraint on task 1, its derivatives with respect to h_1 and h_2 are -0.501 and -0.5. The corresponding reduced-objective derivatives are -2.821311 and -2.815680. With multiplier $\nu_1 = 5.631360$, stationarity holds up to rounding because these equal $\nu_1(-0.501)$ and $\nu_1(-0.5)$. The Hessian of the AI-optimised Lagrangian is approximately $\left(\begin{smallmatrix} -1.128414 & -1.069324 \\ -1.069324 & -1.231176 \end{smallmatrix}\right)$, with eigenvalues -0.109238 and -2.250353. Thus, Configuration B is a strict regular local optimum conditional on the same observed automated set.

Finally, both configurations have the same (q_A, c) and hence the same γ. Equivalently, applying (21) to either allocation gives $q_A/c = \beta^{-1}[Y^{1-\beta}/(h_1 h_2)^{\beta}]^{1/\beta} = 111.389724$ up to displayed rounding. Yet $(L^A, E^A) = (2, 0)$ and $(L^B, E^B) = (6, 4)$, their task shares differ, and their realised output and AI input differ. This verifies Proposition 6 as a payroll-only construction under a common AI technology, conditional on the common automated task set.

References

ACEMOGLU, D. AND D. H. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by D. Card and O. C. Ashenfelter, Amsterdam: Elsevier, vol. 4B, 1043–1171.

ACEMOGLU, D., F. KONG, AND P. RESTREPO (2025): "Tasks at Work: Comparative Advantage, Technology and Labor Demand," in *Handbook of Labor Economics*, Amsterdam: Elsevier, vol. 6, 1–114.

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, 1488–1542.

——— (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33, 3–30.

——— (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128, 2188–2244.

——— (2022): "Tasks, Automation, and the Rise in US Wage Inequality," *Econometrica*, 90, 1973–2016.

AUTOR, D. H. (2015): "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29, 3–30.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118, 1279–1333.

BAILEY, C., R. YEOMAN, A. MADDEN, M. THOMPSON, AND G. KERRIDGE (2019): "A Review of the Empirical Literature on Meaningful Work: Progress and Research Agenda," *Human Resource Development Review*, 18, 83–113.

BECKER, G. S. (1965): "A Theory of the Allocation of Time," *The Economic Journal*, 75, 493–517.

BICK, A., A. BLANDIN, AND D. J. DEMING (2026): "The Rapid Adoption of Generative AI," *Management Science*, published online 20 January 2026.

BRYNJOLFSSON, E., D. LI, AND L. R. RAYMOND (2025): "Generative AI at Work," *Quarterly Journal of Economics*, 140, 889–942.

CASSAR, L. AND S. MEIER (2018): "Nonmonetary Incentives and the Implications of Work as a Source of Meaning," *Journal of Economic Perspectives*, 32, 215–238.

DELL'ACQUA, F., E. MCFOWLAND III, E. R. MOLLICK, H. LIFSHITZ, K. C. KELLOGG, S. RAJENDRAN, L. KRAYER, F. CANDELON, AND K. R. LAKHANI (2026): "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality," *Organization Science*, 37, 403–423.

GANS, J. S. AND A. GOLDFARB (2025): "O-Ring Automation," Working paper.

HOFFMANN, M., S. BOYSEL, F. NAGLE, S. PENG, AND K. XU (2025): "Generative AI and the Nature of Work," Working Paper 25-021, Harvard Business School.

HOLMSTROM, B. AND P. MILGROM (1991): "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization*, 7, 24–52.

HUMLUM, A. AND E. VESTERGAARD (2025a): "Large Language Models, Small Labor Market Effects," Working Paper 33777, National Bureau of Economic Research, revised October 2025.

——— (2025b): "The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities among Workers," *Proceedings of the National Academy of Sciences*, 122, e2414972121.

JIANG, W., J. PARK, R. J. XIAO, AND S. ZHANG (2025): "AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents," Working Paper 33536, National Bureau of Economic Research.

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

MAESTAS, N., K. J. MULLEN, D. POWELL, T. VON WACHTER, AND J. B. WENGER (2023): "The Value of Working Conditions in the United States and the Implications for the Structure of Wages," *American Economic Review*, 113, 2007–2047.

MAS, A. AND A. PALLAIS (2017): "Valuing Alternative Work Arrangements," *American Economic Review*, 107, 3722–3759.

NIKOLOVA, M. AND F. CNOSSEN (2020): "What Makes Work Meaningful and Why Economists Should Care About It," *Labour Economics*, 65, 101847.

Noy, S. and W. Zhang (2023): “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science*, 381, 187–192.

Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer (2023): “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,” Preprint 2302.06590, arXiv.

Ranganathan, A. and X. M. Ye (2026): “AI Doesn’t Reduce Work—It Intensifies It,” *Harvard Business Review*.

Rosen, S. (1974): “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy*, 82, 34–55.

——— (1986): “The Theory of Equalizing Differences,” in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter and R. Layard, Amsterdam: North-Holland, vol. 1, 641–692.