

NBER WORKING PAPER SERIES

O-RING AUTOMATION

Joshua S. Gans
Avi Goldfarb

Working Paper 34639
<http://www.nber.org/papers/w34639>

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

We thank Refine.ink, ChatGPT 5.2 Pro, and Claude Opus 4.5 for research assistance. Funding from the SSHRC is gratefully acknowledged. Responsibility for all errors remains our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34639>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Joshua S. Gans and Avi Goldfarb. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

O-Ring Automation
Joshua S. Gans and Avi Goldfarb
NBER Working Paper No. 34639
January 2026
JEL No. J23, J24, O33

ABSTRACT

We study automation when tasks are quality complements rather than separable. Production requires numerous tasks whose qualities multiply as in an O-ring technology. A worker allocates a fixed endowment of time across the tasks performed; machines can replace tasks with given quality, and time is allocated across the remaining manual tasks. This “focus” mechanism generates three results. First, task-by-task substitution logic is incomplete because automating one task changes the return to automating others. Second, automation decisions are discrete and can require bundled adoption even when automation quality improves smoothly. Third, labour income can rise under partial automation because automation scales the value of remaining bottleneck tasks. These results imply that widely-used exposure indices, which aggregate task-level automation risk using linear formulas, will overstate displacement when tasks are complements. The relevant object is not average task exposure but the structure of bottlenecks and how automation reshapes worker time around them.

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

Avi Goldfarb
University of Toronto
Rotman School of Management
and NBER
agoldfarb@rotman.utoronto.ca

1 Introduction

Automation is often evaluated through a task lens: identify what tasks are performed in a job, determine which tasks can be performed by machines, and infer how labour demand and wages change as machines take over (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2019b). A common formalisation, going back to Zeira (1998), abstracts from cross-task complementarities: tasks (or activities) are effectively separable, so automating one task does not change the productivity of labour on the others.

Many production processes are not like that. In a wide class of processes, tasks are *quality complements*. A single weak link can spoil the product, delay the workflow, or undermine the value of effort elsewhere. Kremer (1993) formalised this idea with the O-ring production function, where output is multiplicative in task qualities. Under such a technology, a change in the quality of one task scales the marginal value of quality in every other task.

This paper studies automation in a stripped-down environment designed to isolate that interaction. A production process consists of n essential tasks. In the benchmark, a single worker performs all tasks and allocates a fixed time endowment h across them. More time on a task increases its quality. Automation becomes available as a capital alternative for any task: for a rental cost r per task, a machine completes that task at fixed quality. The firm chooses which tasks to automate. If at least one task remains manual, the firm and worker then bargain bilaterally over wages (we use Nash bargaining). When a task is automated, the worker reallocates their fixed time to the remaining manual tasks, increasing their quality.¹ Because output is the product of task qualities, this reallocation changes the return to automating any given task.

The model delivers three messages. First, *task-by-task substitution logic is incomplete*. Automating a task not only replaces the quality of that task; it also changes the worker’s time allocation and thus the quality of all remaining manual tasks. This generates a “focus” mechanism: automating some tasks can raise the worker’s contribution on the remaining tasks.

Second, *automation decisions are discrete and bundle-like*. Even when automation

¹This mechanism has a within-worker division-of-labour flavour. Becker and Murphy (1992) emphasise gains from splitting tasks across workers; here the relevant margin is reallocating time within a single worker’s task bundle.

quality improves smoothly, the firm’s optimal number of automated tasks changes via thresholds. Moreover, because the base level of output rises with automation, it can be unprofitable to automate the first task at the margin while still being optimal to automate many tasks in total. This undermines simple marginal adoption rules.

Third, *labour income need not fall under partial automation*. When not all tasks are automated, increases in automation quality can raise labour income because automation scales the value of the remaining labour bottlenecks. This is the opposite of what a purely substitutive, separable task accounting would suggest.

The paper is organised as follows. Section 2 sets up the model and defines the equilibrium concept for labour income. Section 3 solves the worker’s time allocation for any set of automated tasks. Section 4 characterises the firm’s automation choice, first for homogeneous automation quality and then for heterogeneous task-specific quality. Section 5 derives implications for labour income. Section 6 discusses what this implies for prospective and retrospective task-based assessments of automation. Section 7 concludes.

2 Model Setup

A firm produces a single good (price normalised to one) using a process composed of $n \geq 1$ essential tasks, indexed by $s \in \{1, \dots, n\}$. Each task produces a *quality* $q_s > 0$. Output is the product of task qualities, as in an O-ring technology:

$$Y = \prod_{s=1}^n q_s. \quad (1)$$

A single worker supplies a fixed time endowment $h > 0$ (hours) to the firm. If task s is performed manually, the worker assigns $h_s \geq 0$ hours to it and produces quality

$$q_s^L = a h_s, \quad (2)$$

where $a > 0$ is labour productivity (common across tasks). Time is scarce:

$$\sum_{s \in \mathcal{M}} h_s = h, \quad (3)$$

where \mathcal{M} is the set of manual tasks.

For each task s , there exists a capital technology that can perform the task at fixed quality $\theta_s > 0$ at a rental cost $r > 0$ per task. If task s is automated, its quality is

$$q_s^K = \theta_s, \tag{4}$$

and it requires no worker time. Let \mathcal{A} denote the set of automated tasks and let $k = |\mathcal{A}|$ be the number of automated tasks. Then $|\mathcal{M}| = n - k$.

The firm chooses which tasks to automate (i.e. chooses \mathcal{A}), hires the worker for h hours when at least one task remains manual, and time is allocated efficiently across manual tasks.²

A competitive “wage equals marginal product” benchmark is not very informative in this environment. When more than one task is manual, output is multiplicative in the task-level hour allocations, so changes in one task’s hours scale the productivity of all other hours. Moreover, when $k < n$ the firm cannot produce without the worker, so a naive “pay the worker their marginal contribution” rule would assign the worker essentially all net output $Y(\mathcal{A}) - r|\mathcal{A}|$, making labour the residual claimant on capital, which is not an appealing outcome.

Instead, we assume the firm and worker engage in bilateral Nash bargaining.³ If at least one task remains manual ($k < n$), the worker is required for production and the firm and worker engage in bilateral wage bargaining. We assume *efficient bargaining*: conditional on reaching agreement, the time allocation across manual tasks maximises output given \mathcal{A} . This allows us to solve for the efficient time allocation and output first, and then pin down the division of surplus. Let $\beta \in (0, 1)$ denote the worker’s bargaining weight and let $u_0 \geq 0$ denote the worker’s disagreement payoff (the value of spending the h hours elsewhere). If bargaining fails, production does not occur when $k < n$ and the firm earns zero (and incurs no rental costs).

We model wage determination using the Nash bargaining solution (Nash, 1950). Let $W(\mathcal{A})$ denote the *total wage bill* paid to the worker when the automation set is

²This can be interpreted as the firm directing the worker’s time allocation, or as the worker optimising given a contract that makes them the residual claimant for output net of automation rental costs. The results below only require that the time allocation is efficient conditional on \mathcal{A} .

³Automation rental costs $r|\mathcal{A}|$ are incurred only if production occurs (e.g. pay-per-use). Hence the firm’s disagreement payoff when $k < n$ is zero, and automation does not create a hold-up problem in this baseline setup. If automation costs were instead sunk upon adoption, the firm’s disagreement payoff would be $-r|\mathcal{A}|$ and the bargaining stage would need to be modified accordingly.

\mathcal{A} . When $k < n$, Nash bargaining chooses $W(\mathcal{A})$ to maximise

$$(W(\mathcal{A}) - u_0)^\beta (Y(\mathcal{A}) - r|\mathcal{A}| - W(\mathcal{A}))^{1-\beta}. \quad (5)$$

Whenever the match surplus is positive (i.e. $Y(\mathcal{A}) - r|\mathcal{A}| > u_0$), the solution is interior and yields

$$W(\mathcal{A}) = u_0 + \beta(Y(\mathcal{A}) - r|\mathcal{A}| - u_0), \quad (6)$$

and firm profit

$$\Pi(\mathcal{A}) = Y(\mathcal{A}) - r|\mathcal{A}| - W(\mathcal{A}) = (1 - \beta)(Y(\mathcal{A}) - r|\mathcal{A}| - u_0). \quad (7)$$

If all tasks are automated ($k = n$), the worker is not required. In that case, there is no wage bargaining, the firm produces $Y(\mathcal{A}) = \prod_{s=1}^n \theta_s$, and earns $\Pi(\mathcal{A}) = Y(\mathcal{A}) - rn$. The worker is not employed by the firm and receives u_0 elsewhere. When $k < n$, an associated per-hour wage is $w(\mathcal{A}) = W(\mathcal{A})/h$.

To focus on the interesting case, we make the following assumption throughout:

Assumption 1 (Gains from trade). *For any automation choice \mathcal{A} with $k < n$, the match surplus is strictly positive: $Y(\mathcal{A}) - r|\mathcal{A}| > u_0$.*

Hence, whenever the worker is required for production, the Nash bargaining solution is interior and agreement is reached.

3 Time allocation across manual tasks

Given this model set-up, the first step in analysis is to characterise the worker's allocation of their time across tasks. Fix an automation set \mathcal{A} and let \mathcal{M} be the remaining manual tasks, with $m = |\mathcal{M}| = n - k$. Conditional on \mathcal{A} , output is

$$Y(\mathcal{A}, \{h_s\}_{s \in \mathcal{M}}) = \left(\prod_{s \in \mathcal{A}} \theta_s \right) \left(\prod_{s \in \mathcal{M}} a h_s \right). \quad (8)$$

The firm (or worker) chooses $\{h_s\}_{s \in \mathcal{M}}$ to maximise (8) subject to (3). The solution equalises time across manual tasks.

Lemma 1 (Efficient time allocation). *Fix a non-empty manual set \mathcal{M} with $m = |\mathcal{M}| \geq 1$. The allocation that maximises output subject to (3) is*

$$h_s^* = \frac{h}{m} \text{ for all } s \in \mathcal{M}. \quad (9)$$

Proof. Because the automated-task term $\prod_{s \in \mathcal{A}} \theta_s$ and the constant a^m do not depend on the time allocation, the problem reduces to maximising $\prod_{s \in \mathcal{M}} h_s$ subject to $\sum_{s \in \mathcal{M}} h_s = h$ and $h_s \geq 0$. The arithmetic–geometric mean inequality implies

$$\prod_{s \in \mathcal{M}} h_s \leq \left(\frac{1}{m} \sum_{s \in \mathcal{M}} h_s \right)^m = \left(\frac{h}{m} \right)^m,$$

with equality if and only if all h_s are equal, which yields (9). □

Substituting (9) into (8) gives a simple expression for output as a function of the automation set.

Corollary 1 (Output conditional on automation). *If $k < n$ (so $m = n - k \geq 1$), output equals*

$$Y(\mathcal{A}) = \left(\prod_{s \in \mathcal{A}} \theta_s \right) \left(\frac{ah}{n-k} \right)^{n-k}. \quad (10)$$

If $k = n$, output is $Y(\mathcal{A}) = \prod_{s=1}^n \theta_s$.

Equation (10) makes the key interaction transparent: automating tasks (raising k) increases output directly through $\prod_{s \in \mathcal{A}} \theta_s$ and indirectly by reducing the number of manual tasks $n - k$, which increases the quality of every remaining manual task via the factor $ah/(n - k)$.

4 Automation choice

The firm chooses $\mathcal{A} \subseteq \{1, \dots, n\}$ to maximise its expected payoff from wage bargaining. When $k < n$, profit is given by the Nash bargaining outcome in (7). When $k = n$, the worker is not required and the firm appropriates the full net output.

$$\max_{\mathcal{A} \subseteq \{1, \dots, n\}} \Pi(\mathcal{A}). \quad (11)$$

We first solve the homogeneous case $\theta_s = \theta$ for all s to highlight the core logic. We then allow θ_s to vary across tasks and derive a task-by-task ranking result.

4.1 Homogeneous automation quality

Assume $\theta_s = \theta$ for all tasks. Then output depends only on the number of automated tasks k .

Definition 1 (Output, net output, and bargained payoffs as a function of k). *For $k \in \{0, 1, \dots, n\}$ define*

$$Y(k) = \begin{cases} \theta^k \left(\frac{ah}{n-k} \right)^{n-k}, & k \leq n-1, \\ \theta^n, & k = n, \end{cases} \quad (12)$$

and net output (revenue net of automation rental costs) as

$$S(k) = Y(k) - rk. \quad (13)$$

When $k \leq n-1$ the worker is required. Nash bargaining implies a wage bill

$$W(k) = u_0 + \beta(S(k) - u_0) = (1 - \beta)u_0 + \beta S(k), \quad (14)$$

and firm profit

$$\Pi(k) = S(k) - W(k) = (1 - \beta)(S(k) - u_0). \quad (15)$$

When $k = n$ the worker is not hired, $W(n) = 0$, and the firm earns $\Pi(n) = S(n) = \theta^n - rn$.

The firm's automation problem, therefore, reduces to choosing $k \in \{0, \dots, n\}$ to maximise $\Pi(k)$. Conditional on employing the worker ($k \leq n-1$), $\Pi(k)$ is an affine transformation of $S(k)$, so the profit-maximising *partial* automation choice coincides with the k that maximises $S(k)$ over $\{0, \dots, n-1\}$. The new wedge created by bargaining is at the extensive margin: full automation ($k = n$) avoids wage bargaining and lets the firm keep all net output.

4.1.1 Benchmark: frictionless automation ($r = 0$)

When automation is frictionless (i.e., $r = 0$), the output function $Y(k)$ has a simple shape. Conditional on employing the worker ($k \leq n - 1$), maximising profit is equivalent to maximising $Y(k)$ because $\Pi(k) = (1 - \beta)(Y(k) - u_0)$ when $r = 0$.

Proposition 1 (Log-concavity of output). *For $k \in \{0, 1, \dots, n - 1\}$, $\ln Y(k)$ is strictly concave in k under a continuous relaxation. Consequently, the discrete sequence $\{Y(k)\}_{k=0}^n$ is log-concave and unimodal.*

Proof. For $k < n$, taking logs of (12) yields

$$\ln Y(k) = k \ln \theta + (n - k) [\ln(ah) - \ln(n - k)].$$

Differentiating twice with respect to k under a continuous relaxation gives

$$\frac{d^2}{dk^2} \ln Y(k) = -\frac{1}{n - k} < 0,$$

so $\ln Y(k)$ is strictly concave on $[0, n)$. This implies that $\{Y(k)\}_{k=0}^{n-1}$ is log-concave and unimodal. When $n \geq 2$, log-concavity at the last interior index also holds because

$$Y(n - 1)^2 = \theta^{2n-2}(ah)^2 = 4Y(n - 2)Y(n),$$

by direct substitution from Definition 1 (since $(ah/2)^2 = (ah)^2/4$) and when $n = 1$ the claim for $\{Y(0), Y(1)\}$ is immediate. \square

The continuous first-order condition characterises the peak.

Proposition 2 (Frictionless optimum (conditional on employing the worker)). *Suppose $r = 0$ and restrict attention to outcomes with $k \leq n - 1$ (so the worker is employed and bargaining occurs). Under a continuous relaxation, the profit-maximising number of manual tasks is*

$$m^* = n - k^* = \frac{ah}{\theta e}. \quad (16)$$

The optimal integer k^ is the element of $\{0, 1, \dots, n - 1\}$ closest to $n - ah/(\theta e)$, with the boundary cases $k^* = 0$ if $ah/(\theta e) \geq n$ and $k^* = n - 1$ if $ah/(\theta e) \leq 1$.*

Proof. From the expression for $\ln Y(k)$, the first derivative is

$$\frac{d}{dk} \ln Y(k) = \ln \theta + \ln(n - k) - \ln(ah) + 1.$$

Setting the derivative to zero yields $\ln(\theta(n - k)e) = \ln(ah)$ and hence $n - k = ah/(\theta e)$. If $ah/(\theta e) \in (1, n)$, unimodality implies the discrete optimum is one of the integers closest to $n - ah/(\theta e)$; if $ah/(\theta e) \geq n$ then the constrained optimum is $k^* = 0$, and if $ah/(\theta e) \leq 1$ then the constrained optimum is $k^* = n - 1$. \square

Equation (16) is the cleanest summary of the “focus” mechanism: higher automation quality θ reduces the optimal number of manual tasks because machines raise output both directly and by allowing labour to concentrate.

4.1.2 A task-by-task replacement threshold

Even with frictionless automation ($r = 0$), it is not automatic that the firm automates *every* task. When only a few manual tasks remain, each manual task receives a large share of time and can be performed at high quality. This creates a rising “barrier” to automating the last tasks.

Let $m = n - k$ be the number of manual tasks. Compare output before and after automating one additional task (so $m \mapsto m - 1$).

Proposition 3 (Rising quality threshold for automating the next task). *Fix $m \in \{2, \dots, n\}$. With $r = 0$, automating one additional task weakly increases output if and only if*

$$\theta \geq \underline{\theta}(m) \equiv \frac{ah}{m \left(\frac{m}{m-1}\right)^{m-1}} = \frac{ah}{m} \left(\frac{m-1}{m}\right)^{m-1}. \quad (17)$$

Moreover, $\underline{\theta}(m)$ is strictly decreasing in m and satisfies

$$\lim_{m \rightarrow \infty} m \underline{\theta}(m) = \frac{ah}{e} \quad \text{and} \quad \underline{\theta}(2) = \frac{ah}{4}.$$

Proof. With m manual tasks (so $k = n - m$ automated tasks), output is

$$Y_m = \theta^{n-m} \left(\frac{ah}{m}\right)^m.$$

After automating one additional task, there are $m - 1$ manual tasks and output is

$$Y_{m-1} = \theta^{n-(m-1)} \left(\frac{ah}{m-1} \right)^{m-1}.$$

Automation weakly raises output if and only if $Y_{m-1} \geq Y_m$, which is equivalent to

$$\frac{Y_{m-1}}{Y_m} = \frac{\theta}{ah} m \left(\frac{m}{m-1} \right)^{m-1} \geq 1.$$

Rearranging yields $\theta \geq \underline{\theta}(m)$ as defined in (17).

To see that the threshold falls in m , define the continuous extension

$$g(x) = x \left(\frac{x}{x-1} \right)^{x-1} \quad \text{for } x > 1.$$

Then $\underline{\theta}(m) = ah/g(m)$. Moreover,

$$\ln g(x) = x \ln x - (x-1) \ln(x-1), \quad \frac{d}{dx} \ln g(x) = \ln \left(\frac{x}{x-1} \right) > 0,$$

so $g(x)$ is strictly increasing on $(1, \infty)$, hence $g(m)$ is strictly increasing for integers $m \geq 2$ and therefore $\underline{\theta}(m) = ah/g(m)$ is strictly decreasing in m .

Finally,

$$m \underline{\theta}(m) = \frac{ah}{\left(\frac{m}{m-1} \right)^{m-1}} = \frac{ah}{\left(1 + \frac{1}{m-1} \right)^{m-1}} \xrightarrow{m \rightarrow \infty} \frac{ah}{e},$$

using $\lim_{t \rightarrow \infty} \left(1 + \frac{1}{t} \right)^t = e$. Also, $\underline{\theta}(2) = ah/(2 \cdot 2) = ah/4$. □

Proposition 3 formalises a point that task-by-task exercises tend to miss: the relevant comparison for automating a task is not between machine quality and some fixed “human quality”. Human quality is endogenous because time is reallocated when tasks are automated. The last remaining tasks are performed with disproportionately high care and are, therefore, disproportionately hard to replace.

4.1.3 Costly automation ($r > 0$)

With $r > 0$, automation is a discrete choice: the firm picks an integer $k \in \{0, \dots, n\}$, paying rk in rental costs. When at least one task remains manual ($k \leq n - 1$), the worker is required and Nash bargaining implies

$$\Pi(k) = (1 - \beta)(S(k) - u_0), \quad S(k) \equiv Y(k) - rk,$$

so *conditional on employing the worker* the firm chooses k to maximise $S(k)$ over $\{0, \dots, n - 1\}$.

A useful way to see why a task-by-task adoption test can mislead is to look at the *local* (one-step) gain from automating one more task. Because output is multiplicative and time is reallocated across the remaining manual tasks, the marginal return to automation is inherently *state dependent*: it depends on how many tasks are already automated.

Proposition 4 (Local return to automating one additional task). *Fix $k \in \{0, 1, \dots, n - 2\}$ and let $m = n - k$ be the number of manual tasks before automating one more (so $m \geq 2$). The one-step change in surplus from increasing automation from k to $k + 1$ is*

$$\Delta S(k + 1) \equiv S(k + 1) - S(k) = Y(k) \left[\frac{\theta}{ah} m \left(\frac{m}{m - 1} \right)^{m-1} - 1 \right] - r. \quad (18)$$

Equivalently, automating one additional task is locally profitable at k if and only if

$$\theta \geq \underline{\theta}(m) \left(1 + \frac{r}{Y(k)} \right), \quad (19)$$

where $\underline{\theta}(m)$ is the frictionless (output) threshold from (17).

Proof. By definition $S(k) = Y(k) - rk$, so

$$\Delta S(k + 1) = S(k + 1) - S(k) = (Y(k + 1) - Y(k)) - r.$$

For $k \leq n - 2$, both k and $k + 1$ satisfy $k \leq n - 1$, so output is given by (12). Writing $m = n - k$,

$$Y(k) = \theta^k \left(\frac{ah}{m} \right)^m, \quad Y(k + 1) = \theta^{k+1} \left(\frac{ah}{m - 1} \right)^{m-1}.$$

Hence

$$\frac{Y(k+1)}{Y(k)} = \frac{\theta}{ah} m \left(\frac{m}{m-1} \right)^{m-1},$$

and therefore

$$Y(k+1) - Y(k) = Y(k) \left[\frac{\theta}{ah} m \left(\frac{m}{m-1} \right)^{m-1} - 1 \right].$$

Substituting into $\Delta S(k+1) = (Y(k+1) - Y(k)) - r$ yields (18).

Finally, $\Delta S(k+1) \geq 0$ is equivalent to

$$\frac{\theta}{ah} m \left(\frac{m}{m-1} \right)^{m-1} \geq 1 + \frac{r}{Y(k)}.$$

Multiplying both sides by $ah / \left[m \left(\frac{m}{m-1} \right)^{m-1} \right]$ gives $\theta \geq \underline{\theta}(m) \left(1 + \frac{r}{Y(k)} \right)$, i.e. (19). \square

Proposition 4 is easiest to read by separating two forces in (19). The term $\underline{\theta}(m)$ is the *focus barrier*: when fewer manual tasks remain (smaller m), each receives more time and is performed at higher quality, raising the quality requirement for machines to be worth using on the next task. The multiplicative wedge $(1 + r/Y(k))$ is the *cost wedge*: holding r fixed, a higher current output level $Y(k)$ dilutes the effective burden of paying r for one more automated task. Along any range where $Y(k)$ rises with k , the cost wedge shrinks as k rises, while the focus barrier increases mechanically because $m = n - k$ falls.

The implication of such state dependence is that automation is likely to be optimal as a bundled transformation of a number of tasks.

Proposition 5 (Bundled adoption can be privately optimal). *Suppose $n \geq 3$ and consider the partial-automation problem $\max_{k \in \{0,1,\dots,n-1\}} S(k)$, where $S(k) = Y(k) - rk$. If*

$$\max\{0, Y(1) - Y(0)\} < \frac{Y(2) - Y(0)}{2}, \quad (20)$$

then there exists a rental cost $r > 0$ such that $S(1) < S(0)$ yet $\max_{k \geq 2} S(k) > S(0)$.

In particular, any r satisfying

$$\max\{0, Y(1) - Y(0)\} < r < \frac{Y(2) - Y(0)}{2}$$

delivers $S(1) < S(0) < S(2)$.

Proof. We have

$$S(1) - S(0) = Y(1) - Y(0) - r, \quad S(2) - S(0) = Y(2) - Y(0) - 2r.$$

If $r > \max\{0, Y(1) - Y(0)\}$ then $r > 0$ and $S(1) - S(0) < 0$, i.e. $S(1) < S(0)$. If also $r < \frac{Y(2) - Y(0)}{2}$ then $S(2) - S(0) > 0$, i.e. $S(2) > S(0)$. Condition (20) ensures this interval for r is non-empty. \square

Hence, a marginal adoption test based only on whether it is profitable to automate a single task (starting from $k = 0$) can fail: automation may arrive as a discrete bundle rather than one task at a time.⁴ Because $\Delta S(k + 1)$ depends on k through both $m = n - k$ and $Y(k)$, the sign of the first step $\Delta S(1)$ need not predict the sign of later steps. For instance, it can easily be the case that the first unit of automation is not worth paying for, but once some tasks are automated, the same rental cost buys a larger proportional scale-up of output.

Turning now to comparative statics, while the global argmax must generally be computed over $k \in \{0, \dots, n\}$, some directional results are sharp for the partial-automation problem. Let $\tilde{k}^*(\theta, h, r)$ denote any maximiser of $S(k)$ over $k \in \{0, \dots, n - 1\}$ (ties broken toward the largest maximiser), i.e. the optimal *partial* automation choice conditional on employing the worker. Then \tilde{k}^* is (i) weakly increasing in θ and (ii) weakly decreasing in r .

The effect of worker time h is generally ambiguous when $r > 0$: higher h raises the quality of each remaining manual task (reducing the replacement incentive), but it also raises the current output level $Y(k)$ and thus dilutes the proportional burden of paying the fixed rental cost r (which can increase the incentive to automate).⁵ Formally, the monotonicity in r is immediate because $S(k) = Y(k) - rk$ is affine in r with slope $-k$, so a higher r penalises higher k more. Monotonicity in θ does not rely on supermodularity: instead, for any $0 \leq j < k \leq n - 1$, the difference $S(k; \theta, h, r) - S(j; \theta, h, r)$ crosses zero at most once in θ (Lemma 2). Hence, as automation quality improves, the optimal *partial* automation choice $\tilde{k}^*(\theta, h, r)$ is weakly increasing in θ .

Lemma 2 (Single crossing in automation quality for partial automation). *Fix (a, h, n, r) and consider the partial-automation objective $S(k; \theta, h, r) = Y(k; \theta, h) - rk$ over*

⁴This is distinct from the notion that, for instance, artificial intelligence adoption may require system-side change; see Agrawal et al. (2024).

⁵For example, with $n = 5$, $a = 1$, $\theta = 0.5$, and $r = 0.1$, the maximiser of $S(k)$ over $k \in \{0, \dots, n - 1\}$ is $k = 0$ at $h = 3$ but $k = 1$ at $h = 4$.

$k \in \{0, 1, \dots, n-1\}$. For any $0 \leq j < k \leq n-1$, there exists a unique threshold $\bar{\theta}_{j,k} > 0$ such that

$$S(k; \theta, h, r) \geq S(j; \theta, h, r) \iff \theta \geq \bar{\theta}_{j,k}.$$

Consequently, the largest maximiser $\tilde{k}^*(\theta, h, r) \in \arg \max_{k \in \{0, \dots, n-1\}} S(k; \theta, h, r)$ is weakly increasing in θ .

Proof. For $\ell \leq n-1$, write $C_\ell \equiv \left(\frac{ah}{n-\ell}\right)^{n-\ell} > 0$, so $S(\ell; \theta, h, r) = C_\ell \theta^\ell - r\ell$. Fix $0 \leq j < k \leq n-1$ and define

$$D_{k,j}(\theta) \equiv S(k; \theta, h, r) - S(j; \theta, h, r) = C_k \theta^k - C_j \theta^j - r(k-j).$$

As $\theta \downarrow 0$, we have $D_{k,j}(\theta) \rightarrow -r(k-j) < 0$. As $\theta \rightarrow \infty$, $D_{k,j}(\theta) \rightarrow +\infty$ because $k > j$.

Moreover,

$$D'_{k,j}(\theta) = kC_k \theta^{k-1} - jC_j \theta^{j-1} = \theta^{j-1} (kC_k \theta^{k-j} - jC_j),$$

which has at most one zero on $(0, \infty)$. Hence $D_{k,j}$ is decreasing on $(0, \theta^*)$ and increasing on (θ^*, ∞) for $\theta^* = (jC_j / (kC_k))^{1/(k-j)}$ (and $\theta^* = 0$ when $j = 0$). Therefore $D_{k,j}$ crosses zero exactly once at some $\bar{\theta}_{j,k} > 0$, implying the stated iff condition.

For monotonicity of the largest maximiser, let $\theta_2 > \theta_1$ and let $k_1 = \tilde{k}^*(\theta_1, h, r)$. For any $j < k_1$, we have $S(k_1; \theta_1, h, r) \geq S(j; \theta_1, h, r)$, so by the single-crossing property $S(k_1; \theta_2, h, r) \geq S(j; \theta_2, h, r)$. Thus no $j < k_1$ can be the *largest* maximiser at θ_2 , implying $\tilde{k}^*(\theta_2, h, r) \geq k_1$. \square

Bargaining affects the *extensive margin* but not the *partial-automation* ranking. For any $k \leq n-1$, $\Pi(k) = (1-\beta)(S(k) - u_0)$ is an affine transformation of $S(k)$, so the identity of \tilde{k}^* does not depend on (β, u_0) . However, a higher worker bargaining weight β reduces the firm's profit from *every* outcome that still employs labour (all $k \leq n-1$) while leaving $\Pi(n) = S(n)$ unchanged. Consequently, as β rises, full automation (when feasible) becomes weakly more attractive relative to any partial-automation outcome: the set of (θ, h, r) for which the global optimum is $k = n$ expands with β .

4.1.4 When does the firm fully automate?

Full automation corresponds to $k = n$, in which case the worker is not required and wage bargaining is avoided. It is optimal if and only if the firm's profit from full automation weakly exceeds its profit from the best alternative that still requires labour.

Proposition 6 (Condition for full automation). *Full automation ($k = n$) is optimal if and only if*

$$S(n) = \theta^n - rn \geq \max_{k \in \{0, 1, \dots, n-1\}} \Pi(k) = (1 - \beta) \max_{k \in \{0, 1, \dots, n-1\}} \{S(k) - u_0\}. \quad (21)$$

Equivalently,

$$\max_{k \in \{0, 1, \dots, n-1\}} S(k) \leq u_0 + \frac{S(n)}{1 - \beta}. \quad (22)$$

Proof. Condition (21) is the definition of optimality under the profit objective. Substituting $\Pi(k) = (1 - \beta)(S(k) - u_0)$ for $k \leq n - 1$ yields the right-hand equality in (21). Rearranging (21) gives (22). \square

A useful necessary condition, obtained by comparing only $k = n$ and $k = n - 1$, is

$$\theta^n - rn \geq (1 - \beta)[\theta^{n-1}(ah) - r(n - 1) - u_0]. \quad (23)$$

In the special case $r = 0$ and $u_0 = 0$, condition (23) reduces to $\theta \geq (1 - \beta)ah$. Unlike the surplus-maximising benchmark, full automation can, therefore, be privately optimal even when $\theta < ah$: the firm may prefer a lower-output all-capital technology because it avoids sharing surplus with the worker.

4.2 Heterogeneous automation quality and a task ranking

Now allow θ_s to differ across tasks. The firm chooses a subset \mathcal{A} of tasks to automate. Because manual tasks are symmetric in the time allocation problem, only the multiset of automated qualities matters.

Proposition 7 (Sorting: automate the highest-quality tasks). *Suppose automation costs are identical across tasks (cost r per automated task). In any optimal automation*

set \mathcal{A}^* , if tasks are indexed so that $\theta_1 \geq \theta_2 \geq \dots \geq \theta_n$, then there exists an integer k^* such that

$$\mathcal{A}^* = \{1, 2, \dots, k^*\}.$$

That is, the firm automates the k^* tasks with the highest machine quality.

Proof. Consider any feasible automation set \mathcal{A} of size $k < n$. Output under (10) is proportional to $\prod_{s \in \mathcal{A}} \theta_s$, while the manual component depends only on k through $(ah/(n-k))^{n-k}$. If there exist $i \notin \mathcal{A}$ and $j \in \mathcal{A}$ with $\theta_i > \theta_j$, swapping i into \mathcal{A} and j out strictly increases $\prod_{s \in \mathcal{A}} \theta_s$ and leaves costs unchanged, so it strictly increases net surplus. Repeating yields the sorted form. \square

Proposition 7 delivers the task-by-task criterion in the only sense that is robust in this environment: tasks are not decided independently, but conditional on choosing *how many* tasks to automate, the firm automates the tasks with the highest machine quality. The non-separability lies in the choice of k^* , not in which tasks are chosen given k^* .

5 Labour income under partial automation

This section asks a narrow question: *conditional on not fully automating*, what happens to labour income as automation quality improves?

For clarity we return to the homogeneous case $\theta_s = \theta$. When $k \leq n-1$ (so labour is required), the Nash bargaining outcome in (14) implies that the total wage bill paid by the firm is an affine function of net output:

$$W(k) = (1 - \beta)u_0 + \beta S(k),$$

and the associated per-hour wage is $w(k) = W(k)/h$.

Proposition 8 (Bargained labour income rises with automation quality under partial automation). *Fix (a, h, r, n, β, u_0) . Consider the best partial-automation wage bill*

$$\bar{W}(\theta) = \max_{k \in \{0, 1, \dots, n-1\}} W(k; \theta) = (1 - \beta)u_0 + \beta \max_{k \in \{0, 1, \dots, n-1\}} \{Y(k; \theta) - rk\}.$$

Then $\bar{W}(\theta)$ is weakly increasing in θ . Moreover, $\bar{W}(\theta)$ is constant on any interval of θ over which $k = 0$ is optimal, and strictly increasing on any interval over which a

fixed optimal $k^* \in \{1, \dots, n-1\}$ prevails.

Proof. For any fixed $k \in \{1, \dots, n-1\}$, $Y(k; \theta)$ is strictly increasing in θ (it is proportional to θ^k); for $k = 0$, $Y(0; \theta)$ is constant in θ . Hence each function $Y(k; \theta) - rk$ is weakly increasing in θ , and the pointwise maximum over $k \in \{0, 1, \dots, n-1\}$ is weakly increasing. Because $\bar{W}(\theta)$ is an affine function of that maximum, it inherits the same weak monotonicity. On any interval where the maximiser is constant at some $k^* \geq 1$, the maximum is strictly increasing, and so is $\bar{W}(\theta)$. \square

Proposition 8 is the core “remainder” result: holding fixed that the worker remains necessary, better automation raises bargained labour income because it scales the value of the manual bottlenecks, and the worker captures a share β of that incremental surplus.

This does *not* imply labour income rises without bound as automation improves, because sufficiently high automation quality can make full automation optimal (Proposition 6). When the firm fully automates, the worker is not hired by the firm and the wage bill from this job is zero (the worker falls back on u_0 elsewhere). In that sense, the model predicts a potentially non-monotone relationship between automation quality and labour income: rising under partial automation, and collapsing when labour is no longer required. Moreover, because full automation avoids bargaining, this collapse can occur at a lower automation quality when the worker’s bargaining weight β is high.

6 Implications for task-based empirical work

The model is deliberately spare, so the implications are correspondingly sharp. In an O-ring job with fixed worker time, task-based inferences that treat tasks independently can be wrong because tasks may be quality complements and because worker time can be reallocated.

This contrasts with much empirical work on the labour market impact of automation, which often assumes task separability. Specifically, the standard task-based approach typically assumes that a job can be decomposed into a vector of discrete, independent tasks, and that the automation of one task represents a pure substitution of capital for labour, leaving the production function of the remaining human tasks largely unchanged (e.g., Autor et al., 2003; Acemoglu and Restrepo, 2018, 2019a).

This “additive” view of the labour process simplifies aggregation but ignores the interdependence of work activities.

Here we consider the empirical implications of the model and, in particular, the interpretation of recent studies. We consider, first, studies that anticipate the impact of artificial intelligence (or AI) on labour markets, followed by those that study past automation and find evidence consistent with some of the issues raised by the model here.

6.1 Studies that anticipate AI’s impact

By assuming task separability, researchers have developed tools for anticipating the impact of task automation on labour markets. However, these tools will be misleading under an O-ring production function. Starting with the probabilistic classifications of occupation risk developed by Frey and Osborne (2017), and including the patent-text analysis of Webb (2020) and the rubrics of Felten et al. (2021), researchers have sought to quantify the “exposure” of human labour to machine substitution.

These papers do not account for two core features of the mechanism in this paper: (i) tasks are *quality complements* (O-ring production), and (ii) the worker has a *fixed endowment of time* that must be allocated across remaining manual tasks. Under this framework, the automation of a task is not merely a subtraction of labour demand; it is a reallocation event. When a machine takes over a subset of tasks, the worker’s fixed time endowment is concentrated onto the remaining manual tasks, endogenously raising their quality and, potentially, the marginal product of the worker. This “focus effect” implies that standard empirical measures of exposure and wage impacts are biased: prospective indices that treat tasks as additive bundles will systematically overstate displacement risks by failing to account for the rising value of human comparative advantage in remaining “bottleneck” tasks.⁶

Endogenous human quality. A common benchmark compares machine quality to “human quality” on the task. With fixed worker time, human quality on a task is $a(h/m)$, where m is the number of manual tasks *after* automation. That is an endogenous equilibrium object. Automating some tasks increases human quality on the

⁶This is a distinct critique from the one offered in Agrawal et al. (2025), which emphasised that such tools measure short-term effects as implementation tasks get automated.

remaining tasks, potentially making those remaining tasks *less* susceptible to automation than they were initially. An exposure calculation that holds human quality fixed will tend to overstate the degree of job automation. Put differently, prospective studies typically compare machine capability (θ) to a fixed human capability (a_{human}) and classify a task as “exposed” if $\theta > a_{\text{human}}$. In the O-ring fixed-time model, human capability is dynamic: as tasks are automated, m decreases and effective human quality $a(h/m)$ rises. A task that looks automatable today (when the worker is split across many tasks) may become non-automatable tomorrow (when the worker focuses on only a few tasks).

Complementarity, bottlenecks, and aggregation. Because output is multiplicative, the profitability of automating a task depends on the qualities of all other tasks. Proposition 5 shows that the first unit of automation can be privately unprofitable even when automating many tasks is optimal, and Proposition 3 shows a rising threshold for automating the next task. A method that classifies tasks one-by-one can, therefore, miss situations where automation arrives as a bundle (for example, when a platform or integrated system automates multiple steps at once). More basically, when tasks are quality complements, the value of a job is often governed by the least-automatable task (the bottleneck), not the average automatability of its components.

Prospective studies. The most cited study in this domain, Frey and Osborne (2017), estimates that a large share of employment is at high risk of computerization. Their methodology relies on identifying “engineering bottlenecks”—attributes of tasks that make them difficult for current AI to perform—and then using a Gaussian process classifier to estimate automation probabilities for detailed occupations.

The O-ring fixed-time model suggests a different interpretation of bottlenecks. Frey and Osborne effectively treats bottlenecks as static binary shields: if a job involves a hard-to-automate component (e.g., negotiation or social intelligence), the probability of automation is low. This ignores the intensive margin central to this paper. If non-bottleneck tasks in a job are automated, the worker does not necessarily lose their job; instead, they reallocate their time to the bottleneck task, raising its quality from $a(h/n)$ to $a(h/m)$ and potentially increasing the job’s value. Consider a purchasing manager: as administrative components (data retrieval, scheduling, doc-

umentation) are automated, the manager can become a “super-negotiator,” spending a much larger share of time on high-value interactions. Conversely, if the remaining bottleneck is weak, automation may progress further. The point is that the relevant object is not “does the job contain a bottleneck?” but rather how automation changes the *quality and value* of the bottleneck task through time reallocation.

A second generation of prospective studies constructs continuous exposure scores at the task level and aggregates them to occupations. Webb (2020) measures overlap between task descriptions and AI patents; Felten et al. (2021) map AI applications to O*NET abilities; Eloundou et al. (2024) use GPT-4 and human experts to score tasks by potential time savings. Despite differences in measurement, all three aggregate using weighted linear averages:

$$\text{Exposure}_{\text{occ}} = \sum_{t \in \text{Tasks}} w_t \cdot \text{Exposure}_t. \quad (24)$$

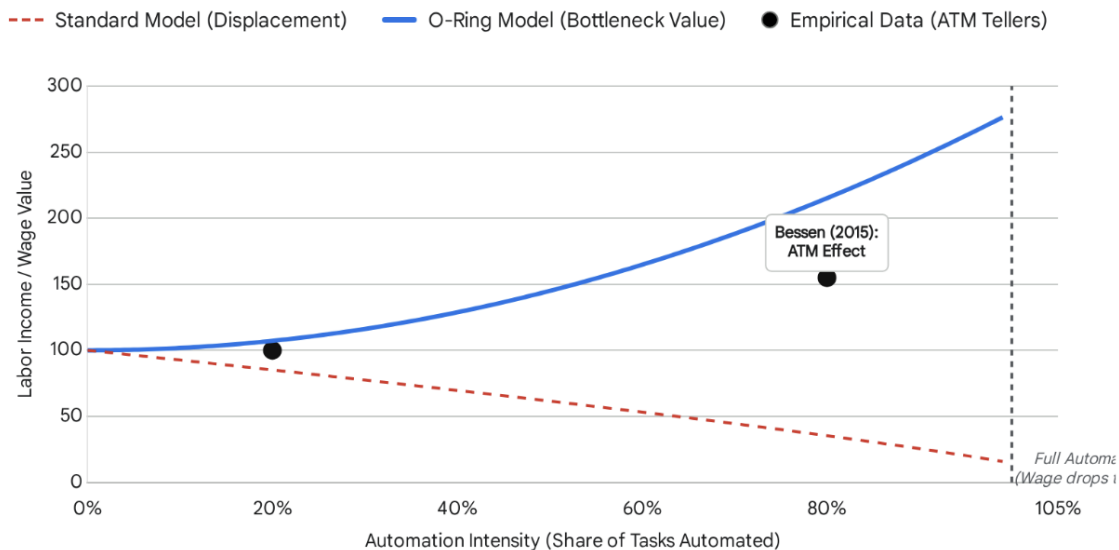
This summation is mathematically inconsistent with O-ring production $Y = \prod_s q_s$. In an O-ring context, if an occupation consists of ten tasks and nine are highly exposed, a linear index reads “90% exposed.” But if the tenth task is a binding bottleneck, automation of the other nine tasks reallocates time into the bottleneck and does not need to eliminate the job. Averaging scores dilutes the protective power of bottlenecks and tends to overstate displacement risk for jobs with heterogeneous task difficulty. Eloundou et al. (2024) acknowledge that task exposure could mean either substitution or complementarity, but lack a mechanism linking the two; the fixed-time O-ring model provides one.

6.2 Evidence for the focus effect and quality complementarity

Several case studies and occupation-level narratives suggest that the “focus effect” and quality complementarity operate in practice.

The case of bank tellers and ATMs, documented by Bessen (2015) and discussed by Autor (2015), provides a canonical illustration. The widespread deployment of ATMs automated routine cash-handling tasks. Under a separable task model, this should have produced sharp displacement. Yet teller employment did not collapse; rather, the occupation shifted toward “relationship banking” and higher-value cus-

The Automation Wage Paradox: Displacement vs. Bottleneck Value



Theoretical predictions of labor income as automation intensity increases. The Standard Task Model (dashed red) predicts declining wages due to displacement. The O-Ring Model (solid blue) predicts rising wages under partial automation as the remaining 'bottleneck' tasks become more valuable (Proposition 8). The empirical case of Bank Tellers (1980–2010) aligns with the O-Ring trajectory.

Figure 1: The automation wage paradox: the standard separable displacement view predicts declining labour income with automation intensity, while the O-ring fixed-time model predicts rising labour income under partial automation as remaining bottleneck tasks become more valuable (Proposition 8). The ATM/teller case is often interpreted as consistent with this pattern (Bessen, 2015; Autor, 2015).

tomers interaction. In the language of this paper, automation removed some tasks from the time budget, reallocating effort toward remaining tasks and raising the quality of the human bottleneck. Importantly, this is not merely a story about “new tasks” appearing exogenously; it is also a story about the endogenous reallocation and intensification of effort in tasks that were already present but not previously central. See Figure 1.

More generally, as Agrawal et al. (2023) discuss, augmentation and automation are not necessarily substitutes. The distinction between automation (replacing tasks) and augmentation (increasing the value of tasks) emphasised by Autor et al. (2022) can

be interpreted through an O-ring lens: what looks like “augmentation” at the occupational level can arise endogenously when automation of non-bottleneck components increases the marginal product of the remaining bottleneck tasks by concentrating worker time. In high-skill settings, the same logic is visible in domains such as radiology: when AI automates components like detection or triage, human effort can shift toward integrative diagnosis and communication. At the same time, evidence suggests this reallocation is not frictionless. Agarwal et al. (2023) show that human-AI collaboration can fail to achieve potential gains when workers misupdate beliefs about AI or treat AI signals inappropriately, highlighting a friction absent from the model: realising the focus effect requires the worker (and organisation) to effectively reallocate effort toward the remaining bottleneck.

Finally, evidence from professional services often points in the same direction. In legal services, technologies that automate discovery and document review are frequently argued to reweight lawyer time toward strategy and advisory components rather than eliminating the profession; Markovic (2019) surveys evidence consistent with this kind of task reshaping. In the O-ring interpretation, automating routine complements can increase the value of the remaining abstract bottleneck tasks by reallocating time and raising quality in the bottleneck.

The analysis in this paper suggests that much of the task-based empirical literature has operated under a restrictive assumption of separability. By treating tasks as independent units that can be peeled away one by one, researchers have built indices that likely overstate the risk of technological unemployment and can misjudge the sign and timing of wage changes.

7 Conclusion

This paper shows that the economics of automation differ when tasks are quality complements and worker time is fixed. Automating a task does not merely substitute capital for labour on that task. It also changes the allocation of labour time across the remaining tasks and, therefore, the quality of those tasks, which feeds back into the return to automation itself.

Three results are central. First, task-by-task substitution logic is incomplete: automating a task changes the worker’s time allocation and, in turn, the quality of all remaining manual tasks, creating a “focus” mechanism that standard expos-

ure measures miss. Second, optimal automation is discrete and can require bundled adoption rather than marginal one-by-one adoption, even when automation quality improves smoothly. Third, conditional on partial automation, improvements in automation quality can raise bargained labour income because automation scales the value of remaining labour bottlenecks and the worker captures a share of the resulting match surplus.

A practical implication is that prospective measures of automation exposure may be misleading when applied to jobs with quality-complementary tasks. The linear aggregation formulas used in widely-cited exposure indices are mathematically inconsistent with O-ring production: they treat a job as the sum of its task-level risks rather than as a system governed by bottlenecks. When tasks are complements and worker time is fixed, automating some tasks raises the quality of the remaining tasks, potentially making them harder to automate than they initially appeared. Exposure indices that hold human task quality fixed will therefore tend to overstate displacement risk. The relevant object is not the average automatability of a job's tasks but the structure of its bottlenecks and how automation reshapes the allocation of worker time around them.

References

- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4B, 1043–1171.
- 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.
- (2019a): “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33, 3–30.
- (2019b): “The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand,” *Cambridge Journal of Regions, Economy and Society*, 13, 25–35.
- AGARWAL, N., A. MOEHRING, P. RAJPURKAR, AND T. SALZ (2023): “Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology,” Working Paper 31422, National Bureau of Economic Research, July 2023 (revised version exists).
- AGRAWAL, A., J. S. GANS, AND A. GOLDFARB (2023): “Do we want less automation?” *Science*, 381, 155–158.
- (2024): “Artificial intelligence adoption and system-wide change,” *Journal of Economics & Management Strategy*, 33, 327–337.
- (2025): “The Economics of Bicycles for the Mind,” Tech. rep., National Bureau of Economic Research.
- 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., C. CHIN, A. SALOMONS, AND B. SEEGMILLER (2022): “New Frontiers: The Origins and Content of New Work, 1940–2018,” Working Paper 30389, National Bureau of Economic Research.
- 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.

- BECKER, G. S. AND K. M. MURPHY (1992): “The Division of Labor, Coordination Costs, and Knowledge,” *The Quarterly Journal of Economics*, 107, 1137–1160.
- BESSEN, J. (2015): “Toil and Technology: Innovative Technology is Displacing Workers to New Jobs Rather than Replacing Them Entirely,” *Finance & Development*, 52, march 2015.
- ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2024): “GPTs are GPTs: Labor market impact potential of LLMs,” *Science*, 384, 1306–1308.
- FELTEN, E. W., M. RAJ, AND R. SEAMANS (2021): “Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses,” *Strategic Management Journal*, 42, 2195–2217.
- FREY, C. B. AND M. A. OSBORNE (2017): “The Future of Employment: How Susceptible Are Jobs to Computerisation?” *Technological Forecasting and Social Change*, 114, 254–280.
- KREMER, M. (1993): “The O-Ring Theory of Economic Development,” *The Quarterly Journal of Economics*, 108, 551–575.
- MARKOVIC, M. (2019): “Rise of the Robot Lawyers?” *Arizona Law Review*, 61, 325–350.
- NASH, J. F. (1950): “The Bargaining Problem,” *Econometrica*, 18, 155–162.
- WEBB, M. (2020): “The Impact of Artificial Intelligence on the Labor Market,” Tech. rep., Working paper, accessed via author website; also circulated as SSRN working paper (Abstract ID: 3482150).
- ZEIRA, J. (1998): “Workers, Machines, and Economic Growth,” *The Quarterly Journal of Economics*, 113, 1091–1117.