The Horse Condition: Labor Demand in the Age of AI

Brian C. Albrecht*

May 6, 2026

Will AI do to human workers what the internal combustion engine did to horses? A large literature measures what AI can technically do. That is a question about displacement inside activities. But aggregate labor demand depends on what consumers buy with the savings. Any cost reduction frees up spending. Labor demand persists as long as that spending lands on anything with human labor somewhere in its value chain. I formalize this through an accounting identity: employment equals total expenditure times the embodied human-labor share, divided by the wage. I derive a support bound on the embodied labor share. The share can vanish only if consumer spending on every activity with embodied labor bounded away from zero itself vanishes. I identify seven mechanisms that determine whether it does. All seven resolved against horses. All remain open for humans.

JEL-Classification: J23, J24, O33, E25

Keywords: AI, automation, labor demand, derived demand, embodied labor, horse condition, expenditure migration

1 Introduction

Every generation of automation technology revives the worry that machines will do to human workers what tractors did to horses (Leontief 1983; Brynjolfsson and McAfee 2015). A growing empirical literature measures which tasks AI can perform, which occupations are exposed, and what happens to workers in them.[1] These are questions about what happens inside production. They do not, by themselves, pin down *aggregate* labor

*International Center for Law & Economics. Email: mail@briancalbrecht.com.

1. On task-level exposure, see Brynjolfsson, Mitchell, and Rock (2018); Eloundou et al. (2024). On occupation-level exposure, see Frey and Osborne (2017); Felten, Raj, and Seamans (2021); Pizzinelli et al. (2023). On labor market outcomes for exposed workers, see Webb (2020); Kogan et al. (2021); Hampole et al. (2025).

demand. What happens to aggregate labor demand depends also on what people buy when AI lowers prices. The answer depends on the share of human labor *embodied* in what consumers buy, meaning the total human labor used at every stage of the value chain, from raw inputs to final delivery.

Within labor demand, cheaper AI does two things at once. Firms substitute AI for workers, which reduces labor demand per unit of output. Cheaper AI also lowers output prices, output expands, and the expansion pulls labor demand back up. Whether labor demand rises or falls depends on which effect is larger. This is the Hicks-Marshall decomposition of derived demand into substitution and scale effects (Hicks 1932; Marshall 1890). The scale effect runs through final-product demand, because lower costs expand output and raise firms' demand for inputs, including labor. Hamermesh's summary of the labor-demand evidence shows why this term cannot be treated as incidental, since much of the movement in employment demand comes from changes in demand for final products (Hamermesh 1993, 349). Substitution nonetheless dominates the theoretical literature because it is analytically richer and more policy-relevant. How strongly the scale channel operates depends on demand elasticity, how much more output people buy when prices fall. Bessen (2019) shows that in early textiles, steel, and autos, demand was elastic enough that automation raised employment. Demand later became inelastic, and automation reduced it. Whether demand is elastic enough varies by activity, but the standard task and growth models (Acemoglu and Restrepo 2019; Jones and Tonetti 2026; Aghion, Jones, and Jones 2019) cannot capture this variation because they aggregate into a single final good, so consumer reallocation across activities with different labor content plays no role. I instead decompose final expenditure. Aggregate labor demand depends both on how much human labor remains in each activity after AI and on where spending moves when relative prices change.

Consider two sectors. AI cuts radiology costs and software development costs by the same percentage. Nobody gets recreational MRIs, so radiology output barely expands and the substitution effect dominates. Radiologist demand falls. Software is different. Many firms and households commission software they could not afford before, output expands sharply, and the scale effect dominates. Developer demand rises. Same cost saving, opposite sign. And the aggregate is not just the sum across sectors. It depends on where consumers send the dollars they save. If AI saves a household \$200 a month on automated services, aggregate labor demand depends on whether that \$200 goes to more automated services or to a piano teacher. The radiology and software cases are about the human labor share within an activity. The piano-teacher case is about how spending shifts across activities with different shares.

To be precise, aggregate labor demand satisfies an accounting identity. Let X denote total expenditure. Let w denote the wage. Let S denote the human-labor share embodied in final expenditure, averaged across activities. Then labor demand is $L = XS/w$. For labor demand to collapse, it is not enough for AI to displace workers inside activities. Every dollar of spending, wherever it lands, must lose all its embodied human labor. The identity has three moving parts. S itself is an expenditure-weighted average, $S = \sum_j m_j s_j$, so it depends on both the labor content inside each activity and the spending shares across them. Displacement inside activities pushes s_j down. But spending can migrate toward activities where s_j remains high, and new tasks can raise s_j from below. The rest of the paper works out these (and additional) margins of adjustment.

Several mechanisms determine s_j, the embodied labor share inside each activity. The substitution elasticity decomposes into within-task substitution and the density of tasks near the human-AI cost margin, which is a frontier hazard rate. But substitution need not be total. Humans and AI often produce in teams, retaining positive human labor content (team production). Workers can retrain into new comparative advantages (skill replenishment). New tasks with human content can appear (reinstatement). And in activities where mistakes are costly, human-supervised production wins on full cost even at zero AI direct cost (reliability).

Where spending migrates when AI lowers prices determines m_j, the expenditure shares across activities. Whether the scale effect or substitution effect dominates determines X/w, aggregate expenditure relative to the wage. Each channel can resolve either way. The horse case is the configuration in which all seven went against horse labor at once. For humans, the question is whether any margin resolves differently in an expenditure-weighted and persistent way. Table 1 states the seven conditions and how each resolved for horses.

I call the configuration in which all seven resolve against human labor the *horse condition*. The list is not exhaustive. Consumer preferences for human-provided services, political protection of labor-intensive industries, and workers' ownership of capital all map to the support bound's primitives but lie outside the seven. The organizing result is a support bound on S. For S to vanish, expenditure mass on every set of activities with embodied labor share bounded away from zero must itself vanish. Whether that happens is the empirical question the seven mechanisms address.

Table 1: The horse condition: seven conditions for labor collapse

Condition	Horses	Humans
Affects embodied labor share s_j		
Frontier hazard Is AI actually cheaper for the tasks that matter?	Near-perfect substitutes	Largely unknown
Reinstatement Do new human tasks appear?	No new horse tasks	Historically positive
Team complementarity Do humans and AI work together?	No	Large and growing
Skill replenishment Can workers retrain into new advantages?	No	Yes, but fragile
Reliability stakes Do mistakes matter enough to keep humans?	Low-stakes field work	Medicine, law, engineering
Affects expenditure shares m_j		
Expenditure migration Where does the freed-up money go?	Away from horse-intensive	Open empirical question
Affects aggregate expenditure X/w		
Scale vs. substitution Do people buy more when it gets cheaper?	Inelastic demand	Varies by activity

A labor demand floor follows from the identity $L = XS/w$, provided expenditure relative to wages does not itself vanish. The bound is mechanical and says nothing about AI capability. It says where to look. Automation can itself shrink the set of human-intensive activities by pushing their embodied labor below any threshold, and it can redirect spending away from whatever remains. The bound does not predict whether either happens. Does any mechanism sustain expenditure on human-intensive activities? Does AI erode the human labor content of those activities faster than spending and new tasks replenish it?

If we specialize to a homothetic CES benchmark and look at local automation price declines, the within-activity and cross-activity margins collapses even further into a single threshold. Fixed-wage labor demand rises with automation whenever $\eta > \Sigma$, where η is the elasticity of final demand and Σ is an effective substitution threshold combining within-activity substitution with between-activity reallocation. This generalizes the textbook Hicks-Marshall condition $\eta > \sigma$ to a multi-activity economy. The horse condition can be stated two ways. The local CES result gives a single-threshold summary. The support bound gives the global organizing condition. Acemoglu (2025) reaches a similar local threshold from a calibrated production-side model. The expenditure identity here shows further that collapse requires no mechanism to sustain expenditure on human-intensive activities.

Section 2 introduces the expenditure identity, the support bound, and the horse condition. Section 3 establishes what production-side analysis can and cannot determine. Section 4 develops four mechanisms that sustain s_j. Sections 5–6 decompose labor-demand changes and derive the CES threshold. Section 7 maps each mechanism to the support condition.

2 The Expenditure Identity and the Horse Condition

The economy contains many activities that differ in human labor intensity. When AI makes some goods cheap, people buy different things. Whether those changes raise or lower aggregate labor demand depends on the embodied labor in what people buy.

A dollar spent on a hand-built ceramic mug embodies mostly human labor. A dollar spent on cloud storage embodies almost none. The aggregate labor share depends on which dollars consumers spend and how much human labor is embodied in what they buy.

Index activities by $j = 1, \dots, J$. Each activity has an expenditure share and an embodied labor share:

$$m_j = \frac{p_j y_j}{X}, \qquad s_j = \frac{w\ell_j}{p_j}. \tag{1}$$

Here $X = \sum_j p_j y_j$ is nominal final expenditure, and ℓ_j is direct-plus-indirect embodied labor per final unit of activity j. Under competitive pricing with nonnegative factor payments, $s_j \in [0, 1]$. With other primary factors or pure profits the aggregation identity $wL = XS$ still holds, but the residual $1 - s_j$ is not necessarily AI capital or any single automatable input. Note that the embodied qualification matters. Direct labor shares identify direct final-stage labor demand. Aggregate final-expenditure labor demand requires labor embodied through the input-output network.

Activity-level labor demand is

$$L_j = y_j \ell_j = \frac{X}{w} m_j s_j. \tag{2}$$

Therefore aggregate labor demand is

$$L = \frac{X}{w} S, \qquad S = \sum_j m_j s_j. \tag{3}$$

Log-differentiating (3) gives

$$d \ln L = d \ln X + d \ln S - d \ln w. \tag{4}$$

At a fixed wage,

$$d \ln L = d \ln X + d \ln S. \tag{5}$$

I hold the wage fixed throughout. Allowing wages to adjust scales the quantity response but does not change its sign (Appendix B.7).

Let q index automation progress along a path. Let $e(q) = d \ln X / dq$ denote expenditure growth and $\kappa(q) = -d \ln S / dq$ the rate at which embodied labor declines. Then

$$\frac{d \ln L}{dq} = e(q) - \kappa(q). \tag{6}$$

Labor demand grows when expenditure grows faster than embodied labor declines. What determines κ?

The identity shifts attention from tasks to expenditure. A sector can have high technical exposure but small aggregate importance if it absorbs little expenditure. Conversely, a low-exposure sector that absorbs large expenditure and remains human-intensive can matter more for aggregate labor demand than a highly exposed sector that nobody spends money on. The automation question is not only what AI can do. It is what automation does to the path of S.

So far labor is a scalar aggregate. Read the identity either as a representative-wage labor-demand schedule or as a wage-bill identity. Distributional claims require applying the same accounting by labor type. When there are multiple labor types, each with its own wage, the identity becomes

$$L^a = \frac{X}{w_a} \sum_j m_j s_{ja}. \tag{7}$$

Some labor types may collapse while the aggregate rises. The scalar expression can also be read as a wage-bill identity, $wL = XS$.

2.1 The support bound

When can S vanish? The answer is accounting, but it is worth being precise. Since $S = \sum_j m_j s_j$ and each $s_j \in [0,1]$, any subset of activities with labor share at least b and expenditure share at least $\underline{m}$ puts a floor on S.

For any $b > 0$,

$$S = \sum_j m_j s_j \geq b \sum_{j:s_j \geq b} m_j. \tag{8}$$

If expenditure mass $\underline{m}$ remains on activities with $s_j \geq b$, then $S \geq \underline{m}\, b$. So $S \to 0$ requires that for every $b > 0$, expenditure mass on the set $\{j : s_j \geq b\}$ vanishes. Conversely, if $s_j \in [0,1]$ and expenditure mass above every positive labor-share threshold vanishes, then $S \to 0$.

A surviving island is concrete. Suppose 15% of consumer expenditure goes to activities with at least 30% embodied human labor: healthcare with human practitioners, in-person education, skilled trades, live performance. Then $S \geq 0.30 \times 0.15 = 0.045$. The aggregate human labor share cannot fall below 4.5%, no matter what happens in the other 85% of the economy. The bound is mechanical and requires no assumption about AI capability.

For labor demand itself to collapse, either S must vanish or X/w must. The support bound pins down what $S \to 0$ requires: for every $b > 0$, expenditure mass on activities with $s_j \geq b$ must vanish. I call the configuration in which both channels resolve against labor the *horse condition*. This is a floor on aggregate labor demand, not a welfare result. It does not imply high wages, smooth transitions, or equal protection across labor types. The type-specific identity in (7) can be applied separately, and some labor types may have no support even when aggregate human labor demand remains positive.

The rest of the paper asks what keeps s_j bounded away from zero inside activities (Sections 3–4) and what keeps m_j nonvanishing on high-s_j activities (Sections 5–6).

3 Within-Activity Labor Demand

Inside each activity j, cheaper AI changes s_j at rate $\gamma_j = a_j - x_j$, where a_j is the price decline and x_j is the within-activity displacement rate. This section derives the local mechanics, first from the Hicks-Marshall formula for a single activity, then from the task-assignment model that microfounds the substitution elasticity. I hold the wage fixed throughout. Appendix B.7 closes the equilibrium.

3.1 The Hicks-Marshall decomposition

Consider one activity with two inputs: human labor L and AI capital K_A. Let s denote labor's cost share, η the elasticity of final demand, and σ the Hicks elasticity of substitution between labor and AI. Under constant returns to scale, competitive pricing, and a downward-sloping demand curve, the standard Hicks-Marshall decomposition (Appendix A) gives the own-wage elasticity of labor demand:

$$\lambda_w = s\eta + (1-s)\sigma. \tag{9}$$

The first term is the scale effect, the second is the substitution effect. When wages rise, both reduce labor demand and they add.

The AI question reverses the experiment. Let the rental price of AI fall by $q_A > 0$ at a fixed wage. The same system yields

$$d\ln L = (1-s)(\eta - \sigma)\, q_A. \tag{10}$$

Now the two channels oppose each other. The scale channel contributes $+(1-s)\eta q_A$: cheaper AI lowers cost, lowers price, and expands output, pulling labor demand up. The substitution channel contributes $-(1-s)\sigma q_A$: cheaper AI induces firms to replace labor with AI capital. The sign depends on which force dominates.

$$\eta > \sigma \implies d\ln L > 0, \tag{11}$$

$$\sigma > \eta \implies d\ln L < 0. \tag{12}$$

When wages rise, both channels reduce labor demand and η affects only the magnitude. When a substitute input gets cheaper, the two channels oppose each other. For wages, the sign is known without η. For cheaper AI, it is not.

Every AI-and-jobs prediction is implicitly a claim about two numbers. The first is σ, how easily AI replaces the input. The second is η, how much more output people want when it gets cheaper. Predicting displacement from σ alone gets the sign wrong whenever $\eta > \sigma$.

When AI displaces tasks, labor demand falls, pushing wages down. Lower wages make humans cheaper relative to AI, pushing some tasks back. Wage feedback dampens the employment response in both directions. Let $\epsilon_S \geq 0$ denote the elasticity of labor supply. The equilibrium employment change is $\epsilon_S/(\lambda_w + \epsilon_S)$ times the fixed-wage shift, which is less than one. Wage adjustment scales the quantity response but does not change its sign. Appendix B.7 derives the formal mapping for the CES case.

3.2 Task assignment and the frontier hazard

The Hicks-Marshall formula gives two parameters (η and σ) but says nothing about which tasks are at risk, or why an AI that aces a demo might not actually displace much labor. The task-assignment framework, developed for offshoring by Grossman and Rossi-Hansberg (2008) and applied to automation by Acemoglu and Restrepo (2019), models which tasks humans perform, which machines perform, and how the allocation shifts when AI gets cheaper. Displacement depends not on how many tasks AI *can* do, but on how many economically important tasks sit near the margin where human and AI costs are close.

Tasks are indexed by $i \in [0, 1]$ with importance weights $\alpha(i) \geq 0$, $\int_0^1 \alpha(i)\, di = 1$. Human productivity at task i is $a_L(i) > 0$. AI productivity is $a_A(i) \geq 0$. Human unit task cost is $w/a_L(i)$; when $a_A(i) > 0$, AI unit task cost is $r_A/a_A(i)$.

Humans perform task i when their cost is lower: $w/a_L(i) \leq r_A/a_A(i)$. For tasks with $a_A(i) > 0$, define human comparative advantage as $\chi(i) = \ln[a_L(i)/a_A(i)]$; for tasks AI cannot perform, set $\chi(i) = +\infty$. Let the relative cost cutoff be $\omega = \ln(w/r_A)$. Then humans perform task i if and only if $\chi(i) \geq \omega$. High $\chi(i)$ means humans are relatively productive on task i. A fall in r_A raises ω, pushing tasks toward AI.

Let G denote the α-weighted CDF of $\chi(i)$, $G(x) = \int_0^1 \alpha(i)\mathbf{1}\{\chi(i) \leq x\}\, di$, and assume G is absolutely continuous at the cutoffs of interest, with density $g = G'$. The human task share is

$$S_L(\omega) = 1 - G(\omega). \tag{13}$$

Differentiating: $dS_L/d\omega = -g(\omega)$. A marginal change in the cutoff displaces human tasks in proportion to $g(\omega)$, the density of tasks near the economic margin, not the total mass of tasks AI could technically perform.

Definition 1 (Frontier hazard rate)**.** Define

$$\phi(\omega) = \frac{g(\omega)}{S_L(\omega)}. \tag{14}$$

Then $d \ln S_L = -\phi(\omega)\, d\omega$. A large ϕ means many important tasks cluster near the switching margin. Most AI exposure studies count how many tasks AI could technically perform. But only what is near the margin matters for displacement.

Aggregate task services with CES elasticity $\sigma_T \geq 0$. By Shephard's lemma, the cost-share derivative has two components. The first is ordinary CES substitution holding assignment fixed. The second is frontier reassignment. Let $\pi(\omega)$ denote the density, in total cost-share units, of tasks exactly at the assignment frontier. Define the cost-share frontier hazard:

$$h(\omega) = \frac{\pi(\omega)}{s_L(\omega)}. \tag{15}$$

The raw task hazard $\phi = g/S_L$ counts tasks near the margin. The cost-share hazard $h = \pi/s_L$ weights them by economic importance. The two coincide under Cobb-Douglas task aggregation ($\sigma_T = 1$) but differ in general.

Proposition 1 (Task decomposition of the substitution elasticity)**.** *Under CES task aggregation with elasticity σ_T, positive factor shares $s_L, s_A > 0$, differentiability, and no atom at the cutoff, the local Hicks elasticity of substitution between labor and AI is*

$$\sigma = \sigma_T + \frac{h(\omega)}{s_A(\omega)}. \tag{16}$$

The first component is the within-assignment task-service elasticity. The second is the frontier reassignment term: the cost-share hazard h scaled by the AI cost share s_A.

Proof. See Appendix B.1. ■

Acemoglu, Kong, and Restrepo (2025) prove a more general version. The substitution elasticity between any two factor groups equals within-task substitution plus reassignment at the margin. Empirical estimates of capital-labor substitution are typically below unity (Oberfield and Raval 2021), so $\sigma_T < 1$ is the empirically relevant case.

Return to the local fixed-wage Hicks-Marshall condition from Section 3.1. In the one-product model, cheaper AI raises labor demand iff $\eta > \sigma$. The proposition above says what σ requires. Technical exposure E does not pin down either component. σ_T is a property of the task aggregator, not of AI capability. And h/s_A depends on how many tasks cluster near the economic cutoff, which exposure studies do not measure.

Most empirical work on AI and labor measures *exposure*: how many tasks could AI technically perform? Task-level studies map machine-learning, patent, or language-model capabilities to work activities (Brynjolfsson, Mitchell, and Rock 2018; Webb 2020; Eloundou et al. 2024; Hampole et al. 2025). Occupation-level studies aggregate task and ability measures into job, occupation, industry, or worker-group exposure (Frey and Osborne 2017; Arntz, Gregory, and Zierahn 2017; Felten, Raj, and Seamans 2021; Pizzinelli et al. 2023). But exposure is not displacement.

Let $\mathcal{A} \subseteq [0, 1]$ be the set of tasks AI could technically perform at some acceptable quality level. Technical exposure is

$$E = \int_{i \in \mathcal{A}} \alpha(i)\, di. \tag{17}$$

E counts capability. It is silent on whether firms actually adopt, whether workers are cheaper anyway, and whether demand expands enough to offset whatever displacement occurs.

If the cutoff moves from ω_0 to $\omega_1 > \omega_0$, the tasks that actually switch are those in the band $[\omega_0, \omega_1)$:

$$I(\omega_0, \omega_1) = \int_{\mathcal{A}} \mathbf{1}\{\omega_0 \le \chi(i) < \omega_1\}\, \alpha(i)\, di. \tag{18}$$

Proposition 2 (Exposure $\neq$ incidence)**.** *$I \le E$, but E does not determine I. Two economies with the same E can have different I.*

Proof. Fix E. Place all exposed task mass at $\chi(i) = 10$ with $\omega_1 = 1$. Then $I = 0$. Alternatively, place it uniformly on $[0, 1]$ with $\omega_0 = 0$, $\omega_1 = 1$. Then $I = E$. ■

Between exposure and displacement sit several distinct steps. Firms must actually adopt AI for the tasks in question. The tasks that switch must be near the economic cutoff, not just technically feasible. And even when tasks do switch, scale effects from cheaper output can offset the displacement. Each step attenuates the link from E_j to the equilibrium employment change.

Horses had narrow comparative advantage concentrated near the cutoff. The frontier hazard was high. For humans, exposure is broad but the frontier hazard is largely unknown. Two economies with the same E can have very different $g(\omega)$ and therefore very different displacement. The gap between what is measured and what matters is not about being careful. It is a missing variable.

The task model microfounds x_j for the support bound: the displacement rate inside each activity decomposes as $x_j = \sigma_{Tj} s_{Aj} + h_j$. A small frontier hazard h_j keeps x_j low and makes it easier for $s_j(q)$ to remain bounded away from zero in expenditure-weighted activities.

4 Sustaining s_j: Four Production-Side Mechanisms

The frontier hazard and within-task substitution determine how fast displacement erodes s_j. Four additional mechanisms can counteract that erosion, keeping s_j bounded away from zero inside specific activities. Each matters for the horse condition only if it operates on a set of activities with nonvanishing expenditure mass. Formally, it must help sustain $s_j(q) \geq b$ on a set with $\sum_j m_j(q) \geq \underline{m}$. What matters is how each, and any additional margin not spelled out explicitly, enters the support condition. I do not model these mechanisms formally; each would require its own model.

4.1 New tasks and reinstatement

So far, the task set is fixed. In practice, technological change also creates new tasks. Automobiles displaced carriage drivers but created mechanics, traffic engineers, and driving instructors. “AI creates lots of new activity” does not imply “AI creates lots of new human work.” Reinstatement is measured by the human content of new tasks, not by the total number of new products.

Horses had $R \approx 0$ in expenditure-weighted terms. No expenditure-significant new horse-complementary tasks appeared. For humans, reinstatement has been empirically positive across the 19th and 20th centuries (Autor 2015; Acemoglu and Restrepo 2019). Autor et al. (2024) document that new job titles account for a large share of employment growth. But they also find that the demand-eroding effects of automation have intensified over the past four decades while the demand-increasing effects of augmentation have not. Whether reinstatement continues under AI is a forward-looking question that the framework identifies but cannot answer.

4.2 Team production

The binary human-versus-AI assignment is too stark for generative AI. In practice, AI often creates a team technology where humans and AI work together on the same task (Brynjolfsson, Li, and Raymond 2023; Dell'Acqua et al. 2023). If team production raises productivity by factor m_i while retaining human labor content $\beta_T(i) \in (0,1)$ per baseline task unit, team AI raises task-level labor demand iff $\ln Y_i' - \ln Y_i > \ln m_i - \ln \beta_T(i)$, where primes denote post-adoption values. Team production complements labor only if output expansion outweighs the net labor saving per unit. Most AI productivity studies estimate the task-level productivity gain but not the output expansion or the change in human labor content. Horses had no team tasks. Humans have many and the set is growing. For the support bound, team production keeps s_j bounded away from zero when retained team labor content is bounded away from zero in activities with nonvanishing spending.

4.3 Apprenticeship and skill formation

Human comparative advantage $\chi(i)$ depends on skill, and workers build skill through experience. If AI automates the entry-level tasks that junior workers learn from, it can raise current output while undermining the supply of skilled workers in the future. As workers lose entry-level experience, their comparative advantage erodes, effectively raising the frontier hazard in future periods. In the short run, senior workers gain. In the long run, the supply of senior workers falls because fewer juniors complete the training path.

Horses did not learn. Their comparative advantage was physical and fixed. Humans maintain comparative advantage through experience, education, and retraining. But that maintenance is not automatic. If AI removes the bottom rungs of the skill ladder, the current static picture overstates human resilience. Whether the training path adjusts determines whether human-intensive islands persist.

4.4 Reliability

When tasks are complements and failure on one component destroys the value of all others (Kremer 1993), the relevant cost is not direct production cost but cost per successful unit. Let mode m have direct cost C_m, success probability Q_m, and failure damage Δ. The full expected cost per successful unit is

$$\tilde{c}_m = \frac{C_m + (1 - Q_m)\Delta}{Q_m}. \tag{19}$$

If AI is less reliable than human-supervised production ($Q_A < Q_H$), then even at zero direct AI cost ($C_A = 0$), all-AI production is cheaper only if

$$\Delta < \frac{Q_A C_H}{Q_H - Q_A}. \tag{20}$$

When damage stakes exceed this threshold, human-supervised production survives regardless of how cheap AI becomes (Albrecht and Thompson 2026). Horses were displaced from low-stakes, mechanically independent tasks. Human cognitive tasks are often high-stakes and complementary: surgery, aviation, structural engineering. Reliability says which islands of human-intensive production exist. Expenditure says whether anyone visits them.

5 Expenditure Migration and Labor Saving

Sections 3–4 developed the production side: what determines s_j inside an activity. This section turns to the demand side. What determines m_j, and hence whether expenditure flows toward or away from human-intensive activities?

What happens to S when AI gets cheaper? Three forces can change the aggregate labor share. Total spending can grow. Spending can migrate across activities with different embodied labor. The embodied labor inside each activity can change. These are distinct margins, and they can point in different directions.

The local notation separates the three channels:

$$e = \frac{d \ln X}{dq}, \qquad a_j = -\frac{d \ln p_j}{dq}, \qquad x_j = -\frac{d \ln \ell_j}{dq},$$
$$g_j = \frac{d \ln m_j}{dq}, \qquad \gamma_j = \frac{d \ln s_j}{dq}, \qquad \mu_j^L = \frac{m_j s_j}{S}. \tag{21}$$

At a fixed wage, $\gamma_j = a_j - x_j$: the embodied labor share in an activity rises when price falls faster than labor content per unit.

Proposition 3 (Local labor-content decomposition)**.** *At a fixed wage,*

$$\frac{d \ln L}{dq} = e + \frac{\mathrm{Cov}_m(s_j, g_j)}{S} + \mathbb{E}_{\mu^L}[\gamma_j], \tag{22}$$

where

$$\mathrm{Cov}_m(s_j, g_j) = \sum_j m_j (s_j - S) g_j. \tag{23}$$

Proof. See Appendix B.2. ∎

The decomposition separates three forces.

$$\begin{aligned} \text{labor-demand shift} = {} & \text{expenditure growth} \\ & + \text{expenditure migration} \\ & + \text{within-activity labor saving.} \end{aligned}$$

The migration term is positive when expenditure shares grow for activities whose embodied labor share is above the expenditure-weighted average. It is negative when expenditure shares grow for activities with low embodied labor.

When AI makes digital content nearly free, the direction of migration depends on where consumers redirect spending. If they spend their savings on restaurant meals, live events, and personal training, expenditure migrates toward human-intensive activities and the covariance is positive. If they spend on more AI-generated content and automated services, it migrates away. The sign is an empirical question about preferences, not a consequence of technology.

This is the demand-side analogue of the frontier hazard in the task model. In the task model, what matters is density near a cost cutoff. In the expenditure model, it is the covariance between human labor intensity and expenditure-share growth.

To see this, note that in the single-final-good case, $J = 1$, so $m_1 = 1$, $S = s_1$, and the migration covariance $\mathrm{Cov}_m(s_j, g_j)$ is identically zero. The standard single-good setup mechanically shuts down the expenditure-migration margin. This is why the standard task and growth models cannot distinguish sectors with identical cost savings but opposite labor-demand outcomes.

The migration term is so far an accounting object. It depends on the covariance between embodied labor and expenditure-share growth, but I have not said where g_j comes from. Without a demand system, migration is a counterfactual. We cannot recover it from baseline data alone. An estimated demand system closes the gap.

Proposition 4 (Demand-system recovery of migration)**.** *Suppose expenditure shares satisfy* $m_j = m_j(p_1, \dots, p_J, X)$*. Let*

$$B_{jk} = \frac{\partial \ln m_j}{\partial \ln p_k}, \qquad \zeta_j = \frac{\partial \ln m_j}{\partial \ln X}, \qquad a_k = -\frac{d \ln p_k}{dq}.$$

Then the expenditure-share response to an automation shock is

$$g_j = -\sum_k B_{jk} a_k + \zeta_j e, \tag{24}$$

and the expenditure-migration component is

$$M = \frac{\mathrm{Cov}_m \left(s_j, -\sum_k B_{jk} a_k + \zeta_j e \right)}{S}. \tag{25}$$

Proof. See Appendix B.3. ■

For horses, the expenditure side was simple. People wanted cheaper plowing, not more plowing-by-horses. No one responded to cheaper tractor plowing by buying more horse-drawn carriage rides. Expenditure migrated away from horse-intensive activities, not toward them. For humans, the expenditure response is an open empirical question. If spending moves toward human-intensive services, care, and experiences after AI makes digital goods cheap, migration supports labor demand. If spending moves toward automated goods and services, migration works the other way.

Where does the task model from Section 3.2 enter the demand side? Through the within-activity share term. Recall that $a_j(q) = -d \ln p_j(q)/dq$ is the price decline in activity j, and $x_j(q) = -d \ln \ell_j(q)/dq$ is the net task-content displacement rate: how fast embodied labor per unit falls due to within-task substitution, frontier reassignment, and any offsetting reinstatement. Since $s_j = w\ell_j / p_j$, holding the wage fixed gives

$$\gamma_j(q) = \frac{d \ln s_j(q)}{dq} = a_j(q) - x_j(q). \tag{26}$$

The embodied labor share of an activity rises when its price falls faster than its labor content ($a_j > x_j$) and falls when displacement outpaces cheapening ($x_j > a_j$). In the no-reinstatement task benchmark from Section 3.2, displacement has the form

$$x_j(q) = \sigma_{Tj} s_{Aj} + h_j, \tag{27}$$

where $\sigma_{Tj} s_{Aj}$ is ordinary within-assignment substitution and h_j is the cost-share hazard at the human-AI assignment frontier. Each ingredient on the right-hand side comes from an existing literature. Task assignment and displacement come from the task framework (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011; Acemoglu and Restrepo 2018). Reinstatement of new human tasks is formalized in Acemoglu and Restrepo (2019) and Acemoglu, Kong, and Restrepo (2025). Bottlenecks and weak links come from AI-growth models (Aghion, Jones, and Jones 2019; Jones and Tonetti 2026). I add the aggregation step, placing those production-side objects inside the expenditure identity so the limit applies to the aggregate embodied labor share, not to a single task frontier. Which activities have labor shares bounded away from zero is a production question. Whether those activities receive nonvanishing spending is an expenditure question. The identity connects them.

Proposition 5 (Aggregate labor-share limit)**.** *For the fixed-wage labor-demand schedule, the decline rate of the aggregate human labor share is*

$$\kappa(q) = -\frac{\mathrm{Cov}_m(s_j, g_j)}{S} + \mathbb{E}_{\mu^L}[x_j - a_j]. \tag{28}$$

Along an infinite automation path, $S(q) \to 0$ *if and only if*

$$\int_0^Q \kappa(q)dq \to +\infty. \tag{29}$$

For the same schedule, human labor demand goes to zero only if

$$X(q)S(q) \to 0. \tag{30}$$

Proof. See Appendix B.4. ∎

For the labor share to vanish, displacement net of migration must accumulate without bound. For labor demand itself to vanish, total expenditure cannot grow fast enough to compensate.

6 The CES Threshold

How elastic is elastic enough? The local decomposition says labor demand depends on expenditure growth, migration, and within-activity labor saving. The local decomposition does not deliver a scalar threshold on its own. The CES specialization closes the migration term and produces a memorable local benchmark. All results in this section hold the wage fixed. Appendix B.7 maps the fixed-wage demand shift into equilibrium wages and employment.

For any automation shock vector (a_j) and any demand system, the general local decomposition (proved in Appendix B.5) is

$$\frac{d\ln L}{dq} = e + a + \frac{\mathrm{Cov}_m(s_j, a_j)}{S} + \frac{\mathrm{Cov}_m(s_j, g_j)}{S} - \mathbb{E}_{\mu^L}[x_j], \qquad a = \sum_j m_j a_j. \tag{31}$$

Three forces support labor demand. Expenditure grows and prices fall (scale), automation cheapens human-intensive activities disproportionately (price incidence), and spending shifts across activities with different embodied labor (migration). One force works against it: AI displaces labor inside activities.

Now specialize. Activities enter a homothetic CES final-demand aggregator with elasticity of substitution ε. Within each activity j, human labor and automation substitute locally with elasticity σ_j. Two moments summarize the economy:

$$V = \mathrm{Var}_m(s_j) = \sum_j m_j (s_j - S)^2, \qquad W = \sum_j m_j s_j (1 - s_j)\sigma_j. \tag{32}$$

Corollary 1 (CES automation threshold)**.** *In the homothetic CES environment, with a local fixed-wage shock and within-activity displacement $x_j = \sigma_j a_j$:*

(i) For any shock vector (a_j), the fixed-wage labor-demand effect is

$$\frac{d\ln L}{dq} = \eta a + \frac{\varepsilon\, \mathrm{Cov}_m(s_j, a_j)}{S} - \mathbb{E}_{\mu^L}[\sigma_j a_j]. \tag{33}$$

(ii) Under the normalization $a_j = 1 - s_j$, this reduces to

$$\frac{d\ln L}{dq} = (1 - S)(\eta - \Sigma), \tag{34}$$

where

$$\Sigma = (1-\rho)\bar{\sigma}_A + \rho\varepsilon, \qquad \rho = \frac{V}{S(1-S)}, \tag{35}$$

and

$$\bar{\sigma}_A = \frac{\sum_j m_j s_j (1-s_j)\sigma_j}{\sum_j m_j s_j (1-s_j)}. \tag{36}$$

When $\sigma_j = \sigma_0$ *for all* j, $\Sigma = (1-\rho)\sigma_0 + \rho\varepsilon$.

Proof. See Appendix B.6. ■

The aggregate Σ is a convex combination of two objects. The within-activity component $\bar{\sigma}_A$ is a weighted average of activity-level substitution elasticities. The between-activity component ε enters through reallocation (Oberfield and Raval 2021). The services sector is not safe just because haircuts are hard to automate. If consumers can substitute away from personal services toward automated alternatives (ε large), the aggregate substitution threshold rises even though no individual service is highly automatable.

The threshold $\eta > \Sigma$ generalizes the two-factor $\eta > \sigma$ to heterogeneous activities. Acemoglu (2025) reaches a similar threshold from a calibrated macro perspective. Appendix B.7 shows how this fixed-wage demand shift maps into equilibrium wages and employment under isoelastic labor supply. A declining labor share does not mean declining labor demand. Under this CES normalization, the labor share declines when $\Sigma > 1$, but labor demand falls only when $\Sigma > \eta$. When $\eta > 1$, the second condition is stronger, and between the two thresholds the labor share is falling but aggregate labor demand is rising. When $\eta < 1$, the ordering reverses.

The clean $\eta > \Sigma$ threshold relies on homotheticity. With non-homothetic demand, expenditure growth changes shares through Engel effects. Let $\zeta_j = \partial \ln m_j / \partial \ln X$ be the expenditure elasticity of activity j's share. The migration term picks up an additional piece:

$$M^X = \frac{e\,\mathrm{Cov}_m(s_j, \zeta_j)}{S}. \tag{37}$$

If human-intensive activities are luxuries (ζ_j high where s_j is high), the covariance is positive and automation-driven income growth shifts spending toward human-intensive activities, making the threshold easier to clear. If human-intensive activities are inferior, the Engel term compounds displacement. For labor type a, the relevant object is $\mathrm{Cov}_m(s_{ja}, \zeta_j)$ from (7). The aggregate Engel term can be positive while the type-specific term is negative for low-skill labor.

7 The Mechanisms Behind the Support Condition

The support bound (Section 2.1) identifies the condition under which collapse is possible. The mechanisms developed in Sections 3–6 are each a reason the lower bound in (8) might persist. They are not logically independent. None is automatically sufficient without magnitude, persistence, and expenditure weight.

The horse collapse had exactly the opposite structure. Tractors were close substitutes for horse power, demand for horse-intensive services was inelastic, and consumers could substitute from horse-powered to engine-powered versions of the same service. No expenditure-significant new horse-complementary tasks appeared. Failure stakes were too low to protect horse labor on full-cost grounds. Recreational horse uses survived, but they did not retain enough expenditure mass to matter for the aggregate. In the notation above, for every $b > 0$, expenditure mass on activities with horse labor share at least b vanished.

The framework identifies the expenditure-weighted conditions that would have to hold for a horse-like collapse of human labor demand.

A Single-Activity Derivation

Two inputs produce output, human labor L and AI capital K_A. Let w denote the wage and r_A the rental price of AI capital. With constant returns to scale the unit cost function $c(w, r_A)$ summarizes the technology. Competitive pricing gives $p = c(w, r_A)$. Final demand is $Y = D(p)$.

Four relationships pin down the local fixed-wage labor-demand schedule. All changes are evaluated locally and written as log-differentials.

1. *Price-cost link.* By Shephard's lemma, unit cost is a share-weighted average of input prices:
$$d \ln p = s \, d \ln w + (1 - s) \, d \ln r_A, \tag{38}$$
where $s = wc_w / c$ is labor's cost share and $1 - s$ is AI capital's share.
2. *Output demand.*
$$d \ln Y = -\eta \, d \ln p, \tag{39}$$
where $\eta > 0$ is the elasticity of final demand.
3. *Substitution.* The Hicks elasticity of substitution $\sigma = d \ln(K_A / L) / d \ln(w / r_A)$ governs how the input mix responds to relative prices:
$$d \ln L - d \ln K_A = -\sigma(d \ln w - d \ln r_A). \tag{40}$$
4. *Output-input link.* With CRS, output change is a share-weighted average of input changes:
$$d \ln Y = s \, d \ln L + (1 - s) \, d \ln K_A. \tag{41}$$

Set $d \ln r_A = 0$ and solve for $d \ln L$ as a function of $d \ln w$. From (40), $d \ln L - d \ln K_A = -\sigma \, d \ln w$. From (41), $d \ln K_A = (d \ln Y - s \, d \ln L) / (1 - s)$. Substituting and using (38)–(39) gives the own-wage elasticity (9). Setting $d \ln w = 0$ and $d \ln r_A = -q_A$ gives the cheaper-AI result (10).

B Proofs

B.1 Task decomposition (Proposition 1)

Proof. By Shephard's lemma, $K_A/L = (s_A/s_L)(w/r_A)$, so

$$\sigma = 1 + \frac{d\ln(s_A/s_L)}{d\ln(w/r_A)}.$$

Under CES task aggregation with a continuous assignment frontier:

$$\frac{d\ln(s_A/s_L)}{d\ln(w/r_A)} = (\sigma_T - 1) + \frac{\pi(\omega)}{s_L(\omega)\, s_A(\omega)}.$$

The first term is within-assignment substitution; the second is frontier reassignment. Adding 1 gives $\sigma = \sigma_T + \pi/(s_L s_A) = \sigma_T + h/s_A$. ■

B.2 Local labor-content decomposition (Proposition 3)

Proof. From (3), $d\ln L/dq = e + (1/S)dS/dq$. Since

$$\frac{dS}{dq} = \sum_j s_j \frac{dm_j}{dq} + \sum_j m_j \frac{ds_j}{dq},$$

and $dm_j/dq = m_j g_j$, $ds_j/dq = s_j \gamma_j$, the first term is $\sum_j m_j s_j g_j/S$. Because expenditure shares sum to one, $\sum_j m_j g_j = 0$, so this equals $\text{Cov}_m(s_j, g_j)/S$. The second is $\mathbb{E}_{\mu^L}[\gamma_j]$. ■

B.3 Demand-system recovery of migration (Proposition 4)

Proof. Taking the total derivative of $\ln m_j(p_1, \ldots, p_J, X)$ with respect to q gives

$$\frac{d\ln m_j}{dq} = \sum_k \frac{\partial \ln m_j}{\partial \ln p_k} \frac{d\ln p_k}{dq} + \frac{\partial \ln m_j}{\partial \ln X} \frac{d\ln X}{dq}.$$

Substitute B_{jk}, $a_k = -d\ln p_k/dq$, and $e = d\ln X/dq$ to obtain (24). Substituting (24) into the migration term in (22) gives (25). ■

B.4 Aggregate labor-share limit (Proposition 5)

Proof. By Proposition 3, $d\ln S/dq = \mathrm{Cov}_m(s_j, g_j)/S + \mathbb{E}_{\mu^L}[\gamma_j]$. Substitute $\gamma_j = a_j - x_j$ and use $\kappa = -d\ln S/dq$. Integrating gives the share condition. The labor-demand condition follows from $L = XS/w$. ∎

B.5 General local decomposition (equation (31))

Proof. From Proposition 3, $d\ln L/dq = e + \mathrm{Cov}_m(s_j, g_j)/S + \mathbb{E}_{\mu^L}[\gamma_j]$. Substituting $\gamma_j = a_j - x_j$ and expanding:

$$\mathbb{E}_{\mu^L}[a_j] = a + \frac{\mathrm{Cov}_m(s_j, a_j)}{S},$$

where the equality uses $\sum_j m_j s_j a_j = Sa + \mathrm{Cov}_m(s_j, a_j)$. Collecting terms gives (31). ∎

B.6 CES automation threshold (Corollary 1)

Proof. Part (i). Homothetic CES demand gives $g_j = (\varepsilon - 1)(a_j - a)$ and aggregate expenditure response $e = (\eta - 1)a$, so $e + a = \eta a$. Since adding a constant does not affect covariance, $\mathrm{Cov}_m(s_j, g_j) = (\varepsilon - 1)\,\mathrm{Cov}_m(s_j, a_j)$. Substituting into (31) collapses the two covariance terms:

$$\frac{\mathrm{Cov}_m(s_j, a_j)}{S} + \frac{(\varepsilon - 1)\,\mathrm{Cov}_m(s_j, a_j)}{S} = \frac{\varepsilon\,\mathrm{Cov}_m(s_j, a_j)}{S}.$$

Under CES within-activity production at fixed wage, $x_j = \sigma_j a_j$ by Shephard's lemma, giving (33).

Part (ii). Under $a_j = 1 - s_j$, $a = 1 - S$ and $\mathrm{Cov}_m(s_j, a_j) = -V$. The displacement term is $\mathbb{E}_{\mu^L}[\sigma_j(1 - s_j)] = W/S$. Substituting:

$$\frac{d\ln L}{dq} = \eta(1 - S) - \frac{\varepsilon V}{S} - \frac{W}{S} = (1 - S)\left[\eta - \frac{\varepsilon V + W}{S(1 - S)}\right].$$

Let $A = S(1 - S) - V = \sum_j m_j s_j (1 - s_j)$. Since $W = \bar{\sigma}_A A$,

$$\frac{W + \varepsilon V}{S(1 - S)} = (1 - \rho)\bar{\sigma}_A + \rho\varepsilon.$$

∎

B.7 Equilibrium wage and employment response (Proposition 6)

The body of the paper works with fixed-wage labor-demand shifts. Allowing wages to adjust scales the quantity response but does not change its sign. Write the fixed-wage demand shift as

$$A = (1 - S)(\eta - \Sigma)dq. \tag{42}$$

The own-wage elasticity of labor demand in the multi-activity model is

$$\Lambda_w = S\eta + (1 - S)\Sigma, \tag{43}$$

which generalizes the single-activity formula $\lambda_w = s\eta + (1 - s)\sigma$ from (9), and labor supply has elasticity $\epsilon_S \geq 0$:

$$d \ln L^s = \epsilon_S d \ln w.$$

Proposition 6 (Equilibrium wage and employment response)**.** *In the CES heterogeneous-activity model with isoelastic labor supply,*

$$d \ln w = \frac{(1 - S)(\eta - \Sigma)}{\epsilon_S + S\eta + (1 - S)\Sigma} dq, \tag{44}$$

and

$$d \ln L = \frac{\epsilon_S}{\epsilon_S + S\eta + (1 - S)\Sigma}(1 - S)(\eta - \Sigma)dq. \tag{45}$$

Proof. Write labor demand as $d \ln L^d = A - \Lambda_w d \ln w$. Market clearing requires $d \ln L^d = d \ln L^s$, so $A - \Lambda_w d \ln w = \epsilon_S d \ln w$. Solving for $d \ln w$ and substituting into labor supply gives (44) and (45). If the denominator is positive, the sign of both responses is governed by $\eta - \Sigma$. ■

References

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40 (121): 13–58. https://doi.org/10.1093/epolic/eiae042.

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics,* vol. 4B, 1043–1171. Elsevier.

Acemoglu, Daron, Fredric Kong, and Pascual Restrepo. 2025. "Tasks at Work: Comparative Advantage, Technology and Labor Demand." In *Handbook of Labor Economics,* 6:1–114. Elsevier.

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488–1542.

———. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3–30.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. 2019. "Artificial Intelligence and Economic Growth." In *The Economics of Artificial Intelligence: An Agenda,* 237–282. University of Chicago Press.

Albrecht, Brian C., and Henry Thompson. 2026. "The Cost of Successful Output: Reliability, Liability, and AI Adoption." Working paper.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2017. "Revisiting the Risk of Automation." *Economics Letters* 159:157–160. https://doi.org/10.1016/j.econlet.2017.07.001.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. "New Frontiers: The Origins and Content of New Work, 1940–2018." *Quarterly Journal of Economics* 139 (3): 1399–1465.

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

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4): 1279–1333.

Bessen, James. 2019. "Automation and Jobs: When Technology Boosts Employment." *Economic Policy* 34 (100): 589–626. https://doi.org/10.1093/epolic/eiaa001.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2023. "Generative AI at Work." NBER Working Paper 31161.

Brynjolfsson, Erik, and Andrew McAfee. 2015. "Will Humans Go the Way of Horses? Labor in the Second Machine Age." *Foreign Affairs* 94 (4): 8–14.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108:43–47. https://doi.org/10.1257/pandp.20181019.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Working Paper 24-013.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of Large Language Models." *Science* 384 (6702): 1306–1308. https://doi.org/10.1126/science.adj0998.

Felten, Edward W., Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (12): 2195–2217. https://doi.org/10.1002/smj.3286.

Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114:254–280. https://doi.org/10.1016/j.techfore.2016.08.019.

Grossman, Gene M., and Esteban Rossi-Hansberg. 2008. "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review* 98 (5): 1978–1997.

Hamermesh, Daniel S. 1993. *Labor Demand.* Princeton University Press.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market." NBER Working Paper 33509. https://doi.org/10.3386/w33509.

Hicks, John R. 1932. *The Theory of Wages.* London: Macmillan.

Jones, Charles I., and Christopher Tonetti. 2026. "Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion." Stanford GSB and NBER Working Paper, Version 0.4.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. 2021. "Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations." NBER Working Paper 29552.

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

Leontief, Wassily. 1983. "National Perspective: The Definition of Problem and Opportunity." In *The Long-Term Impact of Technology on Employment and Unemployment,* 3–7. Washington, DC: National Academy Press.

Marshall, Alfred. 1890. *Principles of Economics.* London: Macmillan.

Oberfield, Ezra, and Devesh Raval. 2021. "Micro Data and Macro Technology." *Econometrica* 89 (2): 703–732.

Pizzinelli, Carlo, Augustus Panton, Marina Mendes Tavares, Mauro Cazzaniga, and Longji Li. 2023. "Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications." IMF Working Paper 2023/216. https://doi.org/10.5089/9798400254802.001.

Webb, Michael. 2020. "The Impact of Artificial Intelligence on the Labor Market." Stanford University Working Paper. https://doi.org/10.2139/ssrn.3482150.