Automation and the Training Divide: Firm Heterogeneity and Internal Labour Reallocation

Perinaz Alp*

July 8, 2026

Abstract

Whether automation destroys jobs or upgrades them depends on a decision that most aggregate models leave in the background: the firm's choice of whether to retrain, reassign, or replace the workers whose tasks a machine can now perform. This paper builds a dynamic, heterogeneous-firms model in which automation, firm-sponsored training, and an *endogenous, firm-specific* promotion threshold jointly determine which workers acquire skills, which move into non-routine roles, and which are displaced. The central analytical contribution is a single-crossing result (Lemma 1) showing that the skill threshold separating routine from non-routine assignment exists and is unique on the entire skill line *without any restriction on the support of the skill distribution*—a sharpening of the standard assignment argument made possible by defining marginal revenue products per unit of worker mass rather than as skill derivatives. Building on this foundation, the paper delivers four results with direct empirical content: the threshold falls with firm productivity and rises with the skill premium (Proposition 1); a binding "routine core" of tasks that survives automation strictly raises the threshold, and does so most severely in small firms (Proposition 2); optimal displacement truncates the skill distribution from below (Proposition 3); and the sign of the automation–threshold relationship is governed by a transparent elasticity condition, $\eta \geq \sigma/(1 - s_R)$ (Proposition 4). A numerical illustration calibrated to a plausible routine revenue share ($s_R \approx 0.59$) shows these forces to be quantitatively large: a 20% productivity advantage raises the promoted share of the routine workforce from 40% to 52%, while doubling automation capital moves the threshold by more than 20% in *either* direction depending on within-bundle substitutability. The framework reconciles classical human-capital theory with the empirical regularity that automation transforms far more jobs than it eliminates.

Keywords: automation, human capital, firm-sponsored training, task assignment, firm heterogeneity, job polarisation.
JEL codes: J24, O33, D21, E24, L23.

1 Introduction

Technological change has reshaped labour demand over recent decades, with automation technologies progressively displacing routine-intensive work, compressing wage distributions, and driving job polarisation across economies. Policy responses have consistently leaned on firm-sponsored training and upskilling as the primary instrument for managing this adjustment, and with good reason: a substantial share of job-related training is financed and delivered by

*Department of Economics, University College London. Correspondence: perinaz.alp@ucl.ac.uk. I am grateful to Onur Demiray for ongoing collaboration on the broader research project of which this chapter forms a part.

employers (Acemoglu and Pischke, 1999). Yet this reliance sits uneasily with the prediction of classical human-capital theory, which holds that a firm should *withdraw* training from tasks facing imminent automation, because the expected return on the investment collapses once a machine can do the job.

The data tell a more nuanced story. Automation frequently *transforms* rather than eliminates jobs, leaving workers employed in roles that demand a higher and qualitatively different set of skills. Reconciling this pattern with theory requires a model in which the training decision, the reassignment decision, and the displacement decision are made jointly and endogenously, inside the firm, in response to the same automation shock. Existing frameworks analyse the *aggregate* employment effects of automation but typically abstract from the endogenous training choices of firms, or treat them as a residual. A unified framework that links automation shocks, firm training choices, and internal worker mobility within a single optimising problem has not previously been developed.

Contribution. This paper closes that gap and, in doing so, makes four contributions that we believe are novel.

1. **An assignment threshold with no boundedness condition.** By recasting the marginal revenue product of a worker as a *per-mass* (Gâteaux) derivative of revenue rather than a derivative with respect to the skill level, we obtain a net-assignment-gap function that crosses zero *exactly once* on the positive half-line. The promotion threshold therefore exists and is unique with no restriction whatsoever on the support of the skill distribution (Lemma 1). This removes a technical assumption that recurs throughout the assignment literature.
2. **An endogenous, firm-specific promotion bar.** In place of the fixed economy-wide quantile rule common in the polarisation literature, the threshold $\bar{h}_{j,t}$ is pinned down by each firm's own profitability, productivity state, and capital stock, in the informational-monopsony tradition of Acemoglu and Pischke (1998). This generates *durable cross-firm heterogeneity* in who gets promoted, trained, and displaced.
3. **A routine core that survives automation.** We introduce a lower bound on the routine employment share—capturing operational, logistical, and maintenance functions that persist even in highly automated plants—and show (Proposition 2) that when it binds it acts exactly like a per-promotion tax whose distortion is *strictly decreasing in firm size*. Small constrained firms hold back promotions most.
4. **A signed automation–threshold relationship.** We show (Proposition 4) that whether automation raises or lowers the skill bar for promotion is governed by a single, interpretable elasticity comparison, $\eta \geq \sigma/(1 - s_R)$, linking within-bundle capital–labour substitutability to across-bundle substitutability and the routine revenue share.

Scope. To keep the focus on within-firm adjustment along the automation and training margins, this version abstracts from firm entry and exit and from external hiring; the workforce of each long-lived firm adjusts only through displacement and internal promotion, and all cross-firm heterogeneity originates from idiosyncratic productivity states and their stochastic evolution. The analysis is conducted in partial equilibrium, with a general-equilibrium closure developed separately (Remark 2).

2 Model Framework: Heterogeneous Firms

Firm Heterogeneity: States and Their Dynamics

The economy is populated by a unit-measure continuum of firms indexed by $j \in \mathcal{J}$. Each firm is characterised in period t by an idiosyncratic state variable $s_{j,t} \in \mathcal{S} \subset \mathbb{R}_{++}$ summarising firm-

specific production conditions—total factor productivity, product-market demand, managerial quality, and organisational efficiency. A higher $s_{j,t}$ raises the return to output and thereby shifts the firm's optimal automation and training policies.

Following Hopenhayn and Rogerson (1993), the idiosyncratic state evolves according to a first-order Markov process with a log-linear specification:

$$\log(s_{j,t}) = a + \rho_s \log(s_{j,t-1}) + \varepsilon_{j,t}, \qquad \varepsilon_{j,t} \sim N(0, \sigma_\varepsilon^2), \quad 0 \le \rho_s < 1, \tag{1}$$

with conditional distribution $F(s' \mid s) = \Phi\big([\log(s') - a - \rho_s \log(s)]/\sigma_\varepsilon\big)$, where $\Phi(\cdot)$ is the standard normal CDF. Persistence in $s_{j,t}$ makes current decisions forward-looking; the persistence parameter ρ_s is consequently central to policy-counterfactual analysis, since highly persistent shocks imply a long expected duration of elevated profitability and thus amplify the incentive to invest in automation and training.

Firm Problem

Each firm j employs workers assigned to either routine (R) or non-routine (N) tasks and simultaneously chooses automation capital $K_{R,j,t}$, training intensity $T_{j,t}$ (provided exclusively to routine-assigned workers), and the measure of routine workers to displace $D_{j,t} \ge 0$. It maximises the expected discounted present value of profits,

$$V_j(\mathcal{S}_{j,0}) = \max_{\{K_{R,j,t}, T_{j,t}, D_{j,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \Pi_{j,t} \right], \qquad 0 < \beta < 1, \tag{2}$$

subject to the production technology, the endogenous evolution of the firm's internal skill distribution, and the transition rules below.

Worker Heterogeneity and Task Assignment

A worker i in firm j is characterised by fixed exogenous characteristics $\chi_{j,i}$ (age, gender, educational background) and by human capital $h_{j,i,t}$, which accumulates endogenously through training, experience, and task assignment. At the opening of period t workers are partitioned into the routine group $R_{j,t}$ and the non-routine group $N_{j,t}$ according to their previous assignment; within the period, displacement and reassignment produce the assigned sets $\widetilde{R}_{j,t}, \widetilde{N}_{j,t}$ that enter production.

Assumption 1 (Absorbing non-routine assignment)**.** *Once a worker is assigned to a non-routine role, the worker remains in the non-routine pool in all subsequent periods.*

Assumption 1 rules out demotions, motivated by the task-specific organisational capital that accrues in non-routine roles and by the reputational and morale costs that make demotion an empirically rare margin.

Production Technology and Output Aggregation

Firm j produces a homogeneous good by combining routine and non-routine task bundles through a CES aggregator, with $s_{j,t}$ a multiplicative productivity shifter:

$$Y_{j,t} = s_{j,t} \cdot F(Y_{R,j,t}, Y_{N,j,t}), \quad F(Y_R, Y_N) \equiv \Big[\theta_t^{1/\sigma} Y_R^{\frac{\sigma-1}{\sigma}} + (1-\theta_t)^{1/\sigma} Y_N^{\frac{\sigma-1}{\sigma}} \Big]^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

with $\sigma > 0$, $\sigma \neq 1$, and $\theta_t \in (0,1)$ an exogenous, deterministic technology weight capturing secular automation trends. The routine bundle nests automation capital and quality-adjusted routine labour,

$$Y_{R,j,t} = \Big[\alpha K_{R,j,t}^{\frac{\eta-1}{\eta}} + (1-\alpha) A_{j,t}^{\frac{\eta-1}{\eta}} \Big]^{\frac{\eta}{\eta-1}}, \qquad A_{j,t} \equiv \int_{\widetilde{R}_{j,t}} h_{j,i,t}^{\phi_R} \, di, \tag{4}$$

with $\eta > 0$, $\eta \neq 1$; non-routine output is linear in skill-adjusted contributions, $Y_{N,j,t} = \int_{\widetilde{N}_{j,t}} h_{j,i,t}^{\phi_N} di$, with $\phi_N > \phi_R$ capturing the greater sensitivity of non-routine productivity to human capital. Writing $Z_{j,t} \equiv \alpha K_{R,j,t}^{(\eta-1)/\eta} + (1-\alpha) A_{j,t}^{(\eta-1)/\eta}$, the two marginal revenue products relevant for assignment are

$$\mathrm{MRP}_{j,t}^{N}(h) = P_t s_{j,t} \frac{\partial F}{\partial Y_{N,j,t}} h^{\phi_N}, \qquad \mathrm{MRP}_{j,t}^{R}(h) = P_t s_{j,t} \frac{\partial F}{\partial Y_{R,j,t}} Z_{j,t}^{\frac{1}{\eta-1}} (1-\alpha) A_{j,t}^{-\frac{1}{\eta}} h^{\phi_R}. \tag{5}$$

Automation capital $K_{R,j,t}$ enters $Z_{j,t}$ directly, so changes in the capital stock shift the effective marginal revenue product of every routine worker—a channel through which automation alters the assignment threshold that is absent from specifications omitting the inner CES layer.

Task Allocation and the Endogenous Firm-Specific Threshold

Assumption 2. *The skill elasticities satisfy $0 < \phi_R < \phi_N < 1$, and the equilibrium wage premium is strictly positive: $w_{N,t} > w_{R,t}$ for all t.*

The threshold $\bar{h}_{j,t}$ is the skill level at which the firm is exactly indifferent between assigning the marginal worker to routine and to non-routine work:

$$\mathrm{MRP}_{j,t}^{N}(\bar{h}_{j,t}) - w_{N,t} = \mathrm{MRP}_{j,t}^{R}(\bar{h}_{j,t}) - w_{R,t}. \tag{6}$$

Lemma 1 (Existence and uniqueness of the threshold)**.** *Fix t and j, and treat the composition-dependent constants $\kappa_{j,t}^{N} \equiv P_t s_{j,t}\, \partial F/\partial Y_{N,j,t} > 0$ and $\kappa_{j,t}^{R} \equiv P_t s_{j,t}\, (\partial F/\partial Y_{R,j,t}) Z_{j,t}^{1/(\eta-1)} (1-\alpha) A_{j,t}^{-1/\eta} > 0$ as given. Under Assumption 2 the net assignment gap*

$$\Delta_{j,t}(h) \equiv \big(\mathrm{MRP}_{j,t}^{N}(h) - w_{N,t}\big) - \big(\mathrm{MRP}_{j,t}^{R}(h) - w_{R,t}\big) = \kappa_{j,t}^{N} h^{\phi_N} - \kappa_{j,t}^{R} h^{\phi_R} - (w_{N,t} - w_{R,t})$$

has exactly one zero $\bar{h}_{j,t} \in (0, \infty)$, with $\Delta_{j,t}(h) < 0$ for $h < \bar{h}_{j,t}$ and $\Delta_{j,t}(h) > 0$ for $h > \bar{h}_{j,t}$. The rule "N if $h \geq \bar{h}_{j,t}$, R otherwise" is therefore well defined on the entire skill support.

Proof. As $h \to 0^+$ both power terms vanish, so $\Delta_{j,t}(0^+) = -(w_{N,t} - w_{R,t}) < 0$; as $h \to \infty$ the term $\kappa_{j,t}^{N} h^{\phi_N}$ dominates since $\phi_N > \phi_R$, so $\Delta_{j,t}(h) \to +\infty$. Differentiating, $\Delta'_{j,t}(h) = h^{\phi_R - 1}\big[\kappa_{j,t}^{N} \phi_N h^{\phi_N - \phi_R} - \kappa_{j,t}^{R} \phi_R\big]$; the bracket is strictly increasing in h, tends to $-\kappa_{j,t}^{R}\phi_R < 0$ as $h \to 0^+$, and diverges to $+\infty$, so it changes sign exactly once at $h_{j,t}^{\dagger} = (\kappa_{j,t}^{R}\phi_R / \kappa_{j,t}^{N}\phi_N)^{1/(\phi_N - \phi_R)}$. Thus $\Delta_{j,t}$ is strictly decreasing on $(0, h_{j,t}^{\dagger})$ and strictly increasing on $(h_{j,t}^{\dagger}, \infty)$. On $(0, h_{j,t}^{\dagger}]$ it stays strictly negative; on $[h_{j,t}^{\dagger}, \infty)$ it increases from a negative value to $+\infty$, so by the intermediate value theorem it has exactly one zero $\bar{h}_{j,t} > h_{j,t}^{\dagger}$ with $\Delta'_{j,t}(\bar{h}_{j,t}) > 0$. The sign pattern follows from the monotonicity structure. □

Remark 1 (Consistency as a fixed point)**.** *The constants $(\kappa_{j,t}^{N}, \kappa_{j,t}^{R})$ depend on the assigned aggregates $(A_{j,t}, Y_{N,j,t})$, which in turn depend on the assignment implied by $\bar{h}_{j,t}$. The period-t threshold is therefore the solution of a fixed-point problem; Lemma 1 guarantees the inner map is single-valued, and existence of a consistent assignment follows from Brouwer's theorem applied to the composite map. Uniqueness is verified numerically.*

Proposition 1 (Threshold comparative statics)**.** *Fix the assignment composition and $K_{R,j,t}$, and write $\bar{w}_t \equiv w_{N,t} - w_{R,t} > 0$. Then*

$$\frac{\partial \bar{h}_{j,t}}{\partial s_{j,t}} = -\frac{\bar{w}_t}{s_{j,t}\, \Delta'_{j,t}(\bar{h}_{j,t})} < 0, \qquad \frac{\partial \bar{h}_{j,t}}{\partial \bar{w}_t} = \frac{1}{\Delta'_{j,t}(\bar{h}_{j,t})} > 0.$$

Proof. Both constants are proportional to $s_{j,t}$, so $\Delta_{j,t}(h) = s_{j,t}(\gamma^N_{j,t} h^{\phi_N} - \gamma^R_{j,t} h^{\phi_R}) - \bar{w}_t$ with $\gamma^N_{j,t}, \gamma^R_{j,t}$ invariant to $s_{j,t}$. At the threshold $\partial\Delta_{j,t}/\partial s_{j,t} = \bar{w}_t/s_{j,t} > 0$; the premium enters only the constant term, so $\partial\Delta_{j,t}/\partial\bar{w}_t = -1$. Applying the implicit function theorem with $\Delta'_{j,t}(\bar{h}_{j,t}) > 0$ gives both results. □

Corollary 1 (Productivity and internal reallocation)**.** *For two firms with a common composition and capital stock, identical surviving workforces, and $s_{j,t} > s_{k,t}$, we have $\bar{h}_{j,t} < \bar{h}_{k,t}$ and the promoted measures satisfy $M_{j,t} \geq M_{k,t}$.*

The routine core constraint imposes a lower bound $\tilde{q}_{j,t} \equiv \widetilde{L}_{R,j,t}/(\widetilde{L}_{R,j,t} + \widetilde{L}_{N,j,t}) \geq q_{\min} > 0$, enforced through a multiplier $\lambda_{j,t} \geq 0$ with complementary slackness $\lambda_{j,t}(q_{\min} - \tilde{q}_{j,t}) = 0$. A promotion moves one unit of mass out of the routine pool, tightening the constraint at rate $\Gamma^M_{j,t} = 1/\widetilde{L}_{j,t}$; a displacement, which also shrinks total employment, tightens it at the strictly smaller rate $\Gamma^D_{j,t} = (1 - \tilde{q}_{j,t})\Gamma^M_{j,t}$.

Proposition 2 (A binding routine core raises the threshold)**.** *For any penalty $c \geq 0$, the constrained condition $\Delta_{j,t}(h) = c$ has exactly one solution $\bar{h}_{j,t}(c) \in [\bar{h}_{j,t}, \infty)$; the map $c \mapsto \bar{h}_{j,t}(c)$ is strictly increasing and continuous with $\bar{h}_{j,t}(0) = \bar{h}_{j,t}$. Hence whenever the constraint binds ($\lambda_{j,t} > 0$), the effective threshold strictly exceeds the unconstrained one, and because $\Gamma^M_{j,t} = 1/\widetilde{L}_{j,t}$ the distortion is strictly decreasing in firm size.*

Proof. Rearranging the binding indifference condition gives $\Delta_{j,t}(h) = \lambda_{j,t}\Gamma^M_{j,t} \equiv c \geq 0$. By Lemma 1, $\Delta_{j,t} < 0$ on $(0, \bar{h}_{j,t})$, so no solution lies below $\bar{h}_{j,t}$ when $c > 0$. On $[\bar{h}_{j,t}, \infty)$, $\Delta_{j,t}$ is a strictly increasing continuous bijection onto $[0, \infty)$, so $\Delta_{j,t}(h) = c$ has the unique solution $\bar{h}_{j,t}(c) = (\Delta_{j,t}|_{[\bar{h}_{j,t},\infty)})^{-1}(c)$, whose inverse is strictly increasing and continuous. Strict monotonicity delivers $\bar{h}_{j,t}(c) > \bar{h}_{j,t}(0)$ whenever $c > 0$. □

Displacement removes the lowest-skill routine workers. Letting $\mu_{j,t}(h)$ denote the retention value of a marginal routine worker of skill h, we have:

Proposition 3 (Bottom-truncation displacement)**.** *Under Assumptions 1–2 and the maintained assumptions on the skill-updating maps, $\mu_{j,t}(h)$ is strictly increasing in h. Since the wage saving from displacement is common across routine workers, the value-maximising displacement set is bottom-truncated: $\mathcal{D}_{j,t} = \{i \in R_{j,t} : h_{j,i,t} \leq h^*_{j,t}(D_{j,t})\}$ with $h^*_{j,t} < \bar{h}_{j,t}$ at an interior optimum.*

Remark 2 (Partial-equilibrium closure)**.** *The price sequences $\{P_t, w_{R,t}, w_{N,t}, r_t\}$ are exogenous to the firm. A general-equilibrium closure with a household block and market clearing is developed in a companion chapter; nothing in the within-firm analysis depends on the particular closure, provided firms remain price takers.*

3 Quantitative Illustration

This section illustrates the analytical results numerically. All objects are evaluated at a *fixed assignment composition*, so that they correspond to the direct (impact) effects characterised in Propositions 1–4, prior to re-equilibration of the fixed point of Remark 1 and to any general-equilibrium price adjustment; no dynamic programme is solved. Parameters are set to standard values ($\sigma = 1.5$, $\eta = 2.0$, $\theta = 0.6$, $\alpha = 0.4$, $\phi_R = 0.3$, $\phi_N = 0.7$, $w_N/w_R = 1.4$) and the baseline idiosyncratic state is normalised so that the baseline threshold coincides with the 60th percentile of a log-normal skill distribution; the implied routine revenue share is $s_R \approx 0.59$. Table 1 collects the resulting magnitudes.

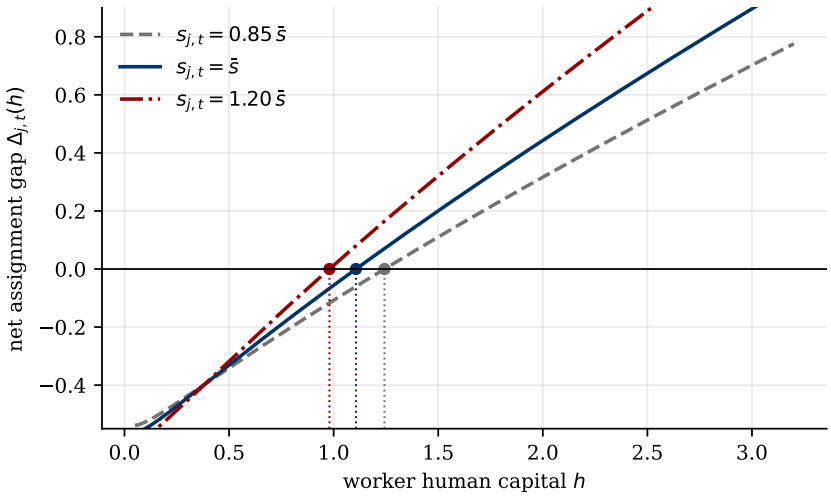


Figure 1: Single crossing of the net assignment gap (Lemma 1). For each productivity level the gap $\Delta_{j,t}(h)$ crosses zero exactly once; the crossing point is the promotion threshold $\bar{h}_{j,t}$. A higher idiosyncratic state $s_{j,t}$ shifts the whole gap upward and moves the crossing to the left, consistent with Proposition 1(i).

Figure 1 confirms the geometry behind Lemma 1: the gap is negative at low skill, crosses zero once, and is positive thereafter, so the assignment rule is well defined with no restriction on the skill support. Raising $s_{j,t}$ lifts the entire curve and pushes the threshold down.

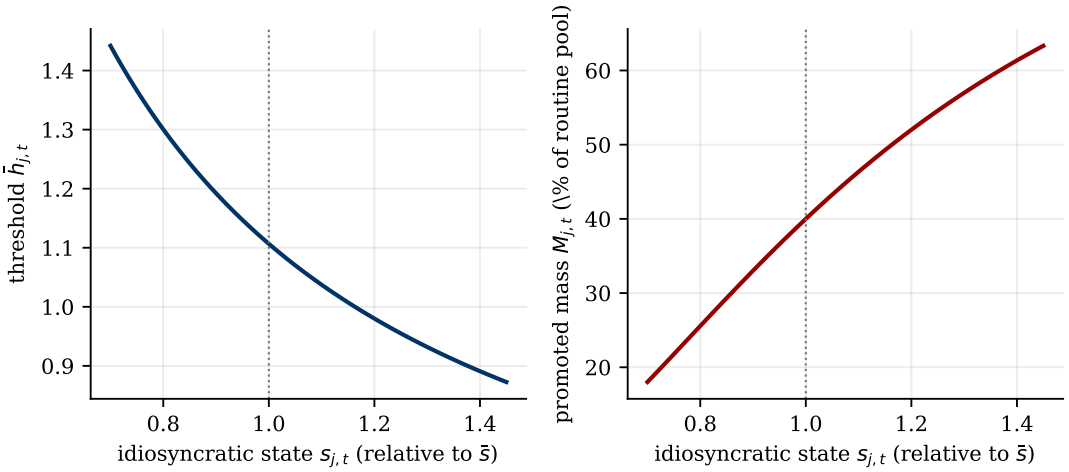


Figure 2: Productivity and internal reallocation (Proposition 1, Corollary 1). The threshold falls monotonically in the idiosyncratic state (left), and the promoted share of the routine workforce rises correspondingly (right). Between a firm at the mean and one 20% more productive, the promoted share climbs from 40% to 52%.

The elasticity of the threshold with respect to productivity is about -0.68 at the baseline. Because the comparison in Corollary 1 follows directly from the sign pattern of Lemma 1 rather than from integrating a local derivative, it holds for arbitrarily large productivity gaps: more productive firms always promote at lower skill and reallocate more workers into non-routine roles.

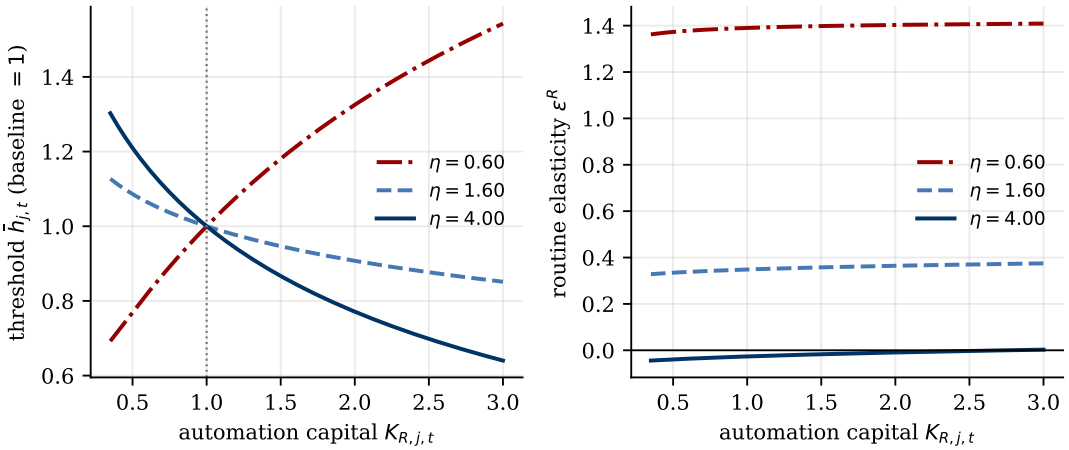


Figure 3: Automation capital and the threshold (Proposition 4). Left: the normalised threshold as automation capital varies, for three values of the within-bundle elasticity η. When capital and routine labour are highly substitutable ($\eta = 4$) automation *lowers* the bar and accelerates promotion; when they are complementary ($\eta = 0.6$) it *raises* the bar. Right: the routine elasticity $\epsilon^R = 1/\eta - (1 - s_R)/\sigma$, whose sign governs the direction of the effect.

Figure 3 makes Proposition 4 concrete. At the baseline routine share the elasticity cutoff is $\eta^\star = \sigma/(1 - s_R) \approx 3.6$: doubling automation capital lowers the threshold by roughly 23% when $\eta = 4$ (substitutable) but raises it by roughly 33% when $\eta = 0.6$ (complementary). The sign of the automation–reallocation relationship is thus an empirical question about a single elasticity, not a modelling assumption.

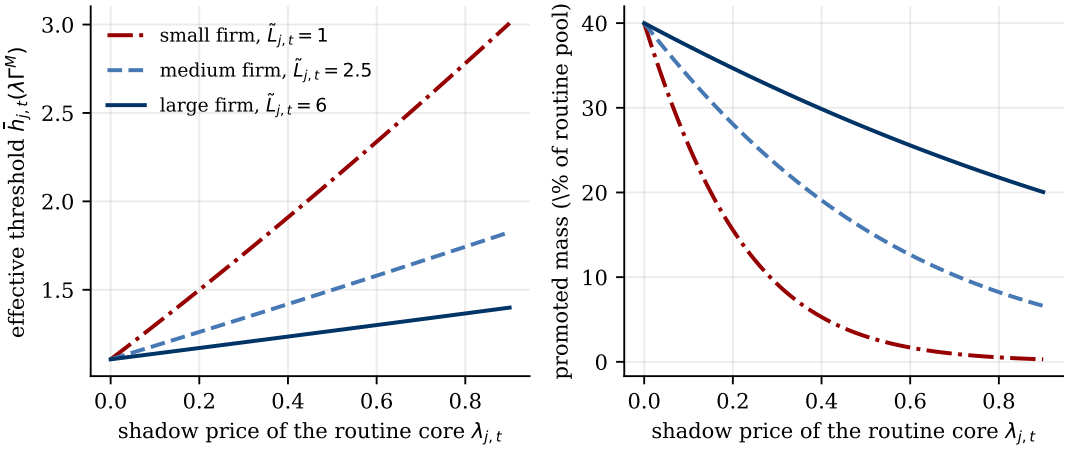


Figure 4: The routine core acts as a size-dependent promotion tax (Proposition 2). As the shadow price $\lambda_{j,t}$ rises, the effective threshold increases (left) and the promoted share collapses (right). Because $\Gamma^M_{j,t} = 1/\widetilde{L}_{j,t}$, the distortion is far larger for a small firm than for a large one: at $\lambda_{j,t} = 0.5$ the promoted share falls to about 3% in the small firm but remains near 28% in the large firm.

Table 1: Selected quantitative magnitudes at the calibrated baseline.

Object	Value
Routine revenue share s_R	0.585
Baseline threshold $\bar{h}_{j,t}$ (60th pctile)	1.107
Elasticity of $\bar{h}_{j,t}$ w.r.t. $s_{j,t}$	-0.68
Promoted share, baseline firm	40.0%
Promoted share, +20% productivity	52.0%
Automation elasticity cutoff $\eta^\star = \sigma/(1 - s_R)$	3.62
$\Delta\bar{h}_{j,t}$ from doubling K_R, $\eta = 4.0$	-22.9%
$\Delta\bar{h}_{j,t}$ from doubling K_R, $\eta = 0.6$	$+32.5\%$
Promoted share at $\lambda = 0.5$, small firm	3.0%
Promoted share at $\lambda = 0.5$, large firm	27.6%

Takeaway. The three margins the model isolates—productivity, automation, and the routine core—each move the promotion threshold by economically large amounts, and they do so in directions that depend on observable primitives (the productivity gap, the elasticity η relative to $\sigma/(1 - s_R)$, and firm size). The framework therefore delivers sharp, testable cross-sectional predictions about which firms retrain and reassign their routine workers and which displace them.

Acknowledgements and Disclosure Statement

This research received no specific grant or financial support from any funding agency in the public, commercial, or not-for-profit sectors. The work was carried out as part of the author's doctoral research at University College London and draws on collaboration with Onur Demiray (EPFL) on the broader research project, who received no compensation in connection with this paper. No proprietary or restricted-access data were used; the numerical illustration relies solely on simulated data generated from the model's stated parameters, and the code is available from the author on request.

The author declares that she has no relevant or material financial interests, paid or unpaid consulting or advisory relationships, board memberships, equity holdings, or other financial or personal relationships that relate to or could be perceived as relating to the research described in this paper. No party had the right to review the paper prior to its dissemination. This disclosure is made in accordance with the NBER Research Dissemination Disclosure Policy; any questions regarding appropriate disclosures may be directed to that policy.

References

Acemoglu, D. and J.-S. Pischke (1998). Why Do Firms Train? Theory and Evidence. *Quarterly Journal of Economics* 113(1), 79–119.

Acemoglu, D. and J.-S. Pischke (1999). Beyond Becker: Training in Imperfect Labour Markets. *Economic Journal* 109(453), F112–F142.

Hopenhayn, H. and R. Rogerson (1993). Job Turnover and Policy Evaluation: A General Equilibrium Analysis. *Journal of Political Economy* 101(5), 915–938.

A Threshold Comparative Statics: Automation Capital

Throughout this appendix the assignment composition is held fixed, so that $(A_{j,t}, Y_{N,j,t})$ and hence the shadow values are treated as parameters; the result is a direct (impact) effect in the same sense as Proposition 1. Firm and time subscripts are suppressed. An increase in K_R raises Z and hence Y_R; with the composition fixed, $\partial Y_R/\partial K_R = Z^{1/(\eta-1)}\alpha K_R^{-1/\eta} > 0$. Let $s_R \equiv (\partial F/\partial Y_R)Y_R/F \in (0,1)$ denote the routine bundle's share. Direct differentiation of the CES shadow values gives the elasticities with respect to Y_R:

$$\epsilon^N \equiv \frac{\partial \ln \kappa^N}{\partial \ln Y_R} = \frac{s_R}{\sigma} > 0, \qquad \epsilon^R \equiv \frac{\partial \ln \kappa^R}{\partial \ln Y_R} = \frac{1}{\eta} - \frac{1 - s_R}{\sigma}.$$

Proposition 4 (Automation and the threshold)**.** *Holding the assignment composition fixed,*

$$\operatorname{sign}\left(\frac{\partial \bar{h}}{\partial K_R}\right) = -\operatorname{sign}\left(\kappa^N \bar{h}^{\phi_N}\epsilon^N - \kappa^R \bar{h}^{\phi_R}\epsilon^R\right).$$

In particular automation strictly lowers the threshold whenever $\epsilon^R \leq 0$, i.e. whenever $\eta \geq \sigma/(1 - s_R)$, and more generally if and only if $\kappa^N \bar{h}^{\phi_N}\epsilon^N > \kappa^R \bar{h}^{\phi_R}\epsilon^R$.

Proof. With the composition fixed, K_R affects Δ only through Y_R and the wedge $Z^{1/(\eta-1)}$ in κ^R, both subsumed in the elasticities above. Then $\partial\Delta/\partial K_R|_{h=\bar{h}} = (\kappa^N \bar{h}^{\phi_N}\epsilon^N - \kappa^R \bar{h}^{\phi_R}\epsilon^R)\,\partial \ln Y_R/\partial K_R$ with $\partial \ln Y_R/\partial K_R > 0$. The implicit function theorem with $\Delta'(\bar{h}) > 0$ yields the stated sign. If $\epsilon^R \leq 0$ the bracket is a sum of a strictly positive and a non-negative term, giving $\partial \bar{h}/\partial K_R < 0$ unambiguously. □

When capital and routine labour are sufficiently substitutable within the routine bundle relative to across-bundle substitutability (adjusted by the routine revenue share), additional automation depresses the effective marginal revenue product of routine workers while raising that of non-routine workers, and the firm optimally promotes at lower skill levels. When η is small (strong within-bundle complementarity), automation can raise routine marginal revenue products enough that the direct effect on the threshold is positive.