

Article

Not peer-reviewed version

Elevate Before You Eliminate: Firms Should Redesign High-Risk Roles Before Any AI-Attributed Layoffs

[Chaoyue He](#)*, [Xin Zhou](#), Di Wang, Hong Xu, Wei Liu, [Chunyan Miao](#)

Posted Date: 20 April 2026

doi: 10.20944/preprints202604.1354.v1

Keywords: generative AI; AI-attributed layoffs; job redesign; workforce transition; AI governance; automation exposure; human capital reporting; social sustainability; elevation space; organizational resilience



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Elevate Before You Eliminate: Firms Should Redesign High-Risk Roles Before Any AI-Attributed Layoffs

Chaoyue He ^{1,*}, Xin Zhou ¹, Di Wang¹, Hong Xu ¹, Wei Liu ² and Chunyan Miao ¹

¹ Alibaba–NTU Global e-Sustainability CorpLab (ANGEL), Singapore

² Alibaba Group, Hangzhou, China

* Correspondence: cyhe@ntu.edu.sg

Abstract

This work argues for a rebuttable presumption against AI-attributed layoffs. The central claim is that many roles now labeled “highly exposed” still contain substantial *elevation space*: when AI compresses routine substrate, it often expands the human layer of judgment, exception handling, coordination, trust, and accountability. We sharpen that claim in a more dynamic way than a fixed-ceiling reading allows. What looks like a role’s “ceiling” is often only a local frontier under frozen job design. With workflow redesign, widened decision rights, new human-owned task creation, and paid skill investment, the frontier can move outward and elevated value can continue rising at higher automation intensities [1–5]. We apply this idea to role families frequently discussed as vulnerable—administration, HR and recruiting, marketing, sales operations, customer support, finance operations, legal and compliance, coding, and research. The argument is sharpened by an emerging managerial pattern. Several firms are no longer merely hinting that AI may eventually reduce headcount; they are explicitly using labor cuts or role reshuffling to self-fund AI investment or AI-linked growth, as shown by recent cases at Atlassian, Block, Workday, HP, SAP, Salesforce, Meta, and others [6–13]. Yet the broader evidence still points more strongly to redesign, complementarity, new-skill demand, and skills bottlenecks than to generalized labor redundancy [14–20]. We therefore propose an **elevate-first rule**: before attributing layoffs to AI, firms should disclose not only the role-specific elevation space they tried to capture, but also the frontier-moving redesigns they attempted, the financing alternatives they considered, the transition budget they allocated, the internal mobility pathways they opened, and the apprenticeship capacity they preserved. Beneficial AI should be evaluated not only by throughput or substitution rates, but by worker elevation, frontier movement, oversight burden, transition quality, and the long-run sustainability of deployment.

Keywords: generative AI; AI-attributed layoffs; job redesign; workforce transition; AI governance; automation exposure; human capital reporting; social sustainability; elevation space; organizational resilience

1. Introduction

Across the AI economy, a managerial default is hardening: once a role is shown to contain draftable, classifiable, retrievable, or otherwise model-addressable tasks, it is increasingly treated as *cuttable headcount*. Reuters updated a running factbox of AI-linked layoffs in April 2026, while Challenger, Gray & Christmas tracked AI as a stated reason in 54,836 announced U.S. job cuts during 2025 and 12,304 more by the end of February 2026 [13,21,22]. These counts do not prove that AI was the sole cause of every cut. They do show that exposure is increasingly being translated—publicly, strategically, and financially—into elimination.

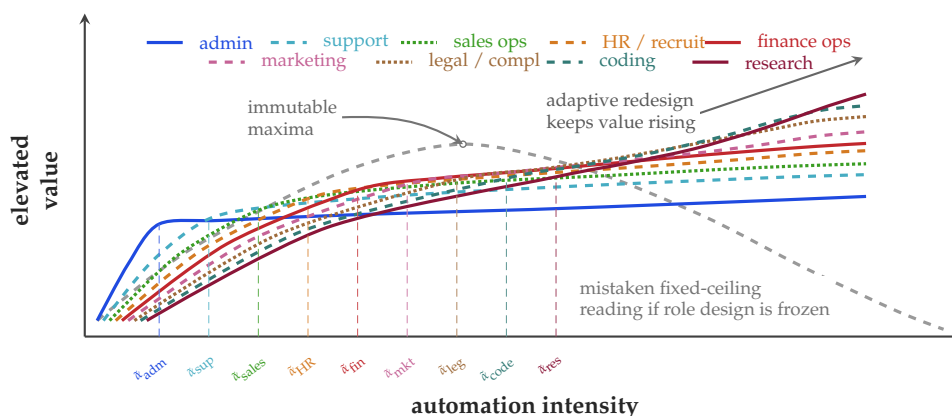


Figure 1. From fixed ceilings to moving frontiers. The gray dashed arc shows the mistaken frozen-role reading, with its peak marking the immutable-maximum interpretation. Colored curves show design-sensitive role paths, and the vertical \bar{a} markers are local redesign checkpoints rather than immutable maxima.

The core analytical mistake is to treat exposure as destiny. Exposure maps identify routinized, model-addressable task bundles. They do *not* tell us how much higher-order human work remains, expands, or newly appears once AI is embedded in a workflow. Yet public discussion often jumps directly from “this role is highly exposed” to “this role should be reduced.” That leap hides the part of the job that becomes more valuable after AI arrives: judgment, exception handling, coordination, trust repair, escalation ownership, and institutional accountability. It also hides a more precise question that firms rarely answer before cutting people: how much of the released routine substrate can be converted into elevated human work before the current job design stalls—and how much further could the frontier be moved with redesign?

We argue for a rebuttable presumption against AI-attributed layoffs: firms must elevate high-risk roles through targeted redesign before eliminating headcount. Rather than treating AI automation as a mandate for immediate workforce reduction, organizations must first systematically capture the role’s “elevation space”—the expanded human layer of orchestration and accountability that emerges when AI absorbs routine execution.

The paper advances four linked claims. Descriptively, exposure scores identify vulnerable task bundles rather than whole-job destiny. Analytically, a role’s first apparent ceiling is usually local to a particular design state and can shift with redesign. Normatively, AI-attributed layoffs should face a rebuttable presumption unless firms show serious redesign, training, mobility, and apprenticeship-preservation efforts. Evaluatively, deployment success should be reported with worker-transition metrics, not throughput alone.

The counterexamples matter because they show that the layoff-first reading is managerial, not technological. Reuters reported that Baidu began large-scale layoffs in late 2025 while AI- and cloud-related positions were largely protected [23]. ByteDance also cut moderation roles while expanding AI moderation [24]. But other firms in the same global race chose a different path. Alibaba said in March 2025 that it would “reboot and rehire” [25]; in March 2026 it launched the Wukong enterprise platform to automate documents, spreadsheets, transcription, and research tasks for DingTalk users [26]. Tencent simultaneously ramped AI talent acquisition, and Xiaomi committed at least 60 billion yuan to AI over three years [27,28]. SAP’s 2024 transformation program is an especially revealing middle case: it explicitly pursued Business AI and AI-driven efficiencies, but paired that move with voluntary programs, internal re-skilling, and a plan to end the year at roughly similar headcount [10,29]. Heavy automation, then, does not logically entail layoff-first restructuring.

A further mistake is to treat a first apparent peak in elevated value as final. Recent organizational evidence suggests that sustained AI value requires end-to-end workflow redesign rather than isolated tool use, while only a minority of organizations are yet fundamentally redesigning how work is performed [4,5]. Recent theory makes the same point from different directions: AI can augment workers without automating tasks, reorganize how steps are chained into jobs, and create new human-

owned task mass by compressing coordination costs [1–3]. We therefore argue not only against cut-first governance, but also against fixed-ceiling reasoning. A local bottleneck should trigger redesign pressure, not an automatic redundancy conclusion.

The claim is not that every role can be preserved indefinitely, nor that all layoffs are illegitimate. It is a sequencing rule and a burden-of-proof rule. Before firms declare a role redundant, they should redesign the job upward, fund paid upskilling, create internal mobility pathways with wage protection, and preserve apprenticeship ladders. The deeper claim is that many roles now described as vulnerable—administration, HR, recruiting, marketing, sales operations, customer support, finance operations, legal and compliance, coding, research, and adjacent knowledge work—still contain substantial elevation space, which is endogenous to managerial choices.

This framework holds particular weight for the AI research community, which shapes not only benchmarks but also what organizations believe counts as deployment success. If AI success is narrativized mainly as payroll compression, firms will optimize toward substitution. If successful deployment is instead understood as technology that raises the level at which more people can contribute, then workflow redesign, capability diffusion, oversight quality, and apprenticeship preservation become central research questions rather than afterthoughts.

2. The New Financing Logic: AI Is Increasingly Being Paid for Through Labor Reallocation

A stronger version of the layoff-first case has now become visible. For several firms, payroll savings and role reshuffling are being explicitly used right now to self-fund AI capacity and AI-linked go-to-market expansion. Recent restructuring cases at Atlassian, Block, Workday, HP, SAP, Salesforce, Oracle, and Meta demonstrate that labor cuts are increasingly functioning as a capital-reallocation mechanism rather than pure task automation [6–12,30,31]. To avoid conflating distinct mechanisms, Appendix I codes the company evidence into four classes: direct automation displacement, AI-financing or role-mix reallocation, mixed restructuring, and counterexamples with rehiring or re-skilling. Appendix J then provides the extended catalog of cases.

This distinction matters because it changes the burden of proof. When firms openly state that cuts are being used to finance AI or rebalance toward AI-linked growth, the relevant question is no longer whether AI played some background role. It clearly did. The relevant question becomes whether payroll compression was the least-destructive financing instrument, whether the newly funded roles could have been filled through internal mobility, and whether the firm attempted to convert released routine work into elevated ownership before eliminating the people who held that work. In other words, once AI becomes part of the financing story, AI governance must evaluate *transition design*, not only model performance.

We therefore do not claim that all AI-linked layoffs are fictitious, nor that every cut announced alongside AI spending is pure opportunism. Some cases are explicit self-funding moves; some are mixtures of real technological reallocation and ordinary restructuring; some combine large cuts with internal reskilling or renewed hiring in adjacent areas. But that heterogeneity strengthens rather than weakens the elevate-first position. When the financing and skill-mix logic is real, the need for role-specific frontier measurement, transition disclosure, and apprenticeship preservation becomes more urgent, not less. Recent theory sharpens the point: competition can produce an automation arms race in which firms rationally over-displace labor relative to the social optimum, even when they understand the demand-side harms of doing so [32]. A presumption against AI-attributed layoffs is therefore not an appeal to nostalgia. It is a governance response to a live strategic equilibrium.

3. From Exposure to Elevation Space

3.1. Exposure Indices Identify Vulnerable Task Bundles, not Whole-Job Destiny

The exposed-role evidence is real and should not be minimized. GPT-style systems affect a large share of white-collar task bundles across the labor market [20]. The ILO finds that clerical

occupations are the most exposed, while professional and technical work is increasingly exposed as systems improve; yet it also concludes that most exposed jobs are more likely to be transformed than eliminated [15,33]. OECD work similarly shows that highly exposed occupations include administrative assistants, accountants and financial analysts, software developers, managers, and HR professionals, but that the skill demand in those roles is already broader than raw automation narratives imply [16]. The World Economic Forum likewise reports that routine clerical roles, administrative assistants, and accountants are under pressure, even as employers simultaneously report rising demand for more analytical, managerial, and technology-complementary capabilities [34,35].

The problem is how these exposure signals are used. In both popular reporting and internal management rhetoric, exposed roles are often treated as if their economically meaningful content were exhausted by the routine slice that current models compress. But jobs are not single tasks. They are bundles of production, assurance, escalation, communication, coordination, and accountability. The task-based literature has made this point for years: technological change usually displaces some tasks while reinstating others and changing the boundary between them [36–38]. Current generative AI evidence strengthens rather than weakens that logic.

3.2. *The Overlooked Layer Is the Elevated Human Layer*

Once AI handles the first pass, the remaining human work does not disappear. It often changes level. Current AI can draft a response, generate variants, write boilerplate, summarize a dossier, propose code, or screen a set of candidates, but someone still has to own policy fit, exception handling, contextual judgment, stakeholder trust, and the consequences of being wrong. That “elevated” layer is not a moral decoration around the system. It is part of the production function.

Evidence from multiple domains points in this direction. Generative AI often lifts less experienced workers most on tasks within the model’s capability frontier [14,39–41]. Dell’Acqua et al. show that knowledge workers perform better with AI on tasks inside a jagged technological frontier but can do worse outside it, which implies that human value shifts toward recognizing frontier boundaries, validating outputs, and taking responsibility for outside-frontier cases [42]. Anthropic’s large-scale Claude usage evidence likewise suggests a mixed pattern of automation and augmentation rather than replacement [43]. In other words, the question is not whether AI touches a role, but how much judgment-heavy, exception-heavy, and coordination-heavy work surrounds the part that AI touches.

This reframing is especially important for roles that many observers now describe as disposable. Consider HR, marketing, sales, administration, legal, coding, and research. In each case, AI compresses a real substrate of search, drafting, formatting, scheduling, or boilerplate production. But in each case, the remaining work can be elevated upward into more valuable functions: calibration, orchestration, experimentation, architecture, synthesis, and institutional accountability. The fact that this elevated layer is less routinized is exactly why it is easy to miss in exposure scores and easy to destroy in layoff-first restructuring.

3.3. *The Frontier Itself Is Endogenous*

The larger correction is that the frontier is not fixed. Many discussions implicitly assume a one-shot hump: AI first raises the value of the remaining human layer, then eventually drains it. Recent organizational evidence points to a more dynamic process. The World Economic Forum argues that the larger gains from AI come from end-to-end workflow redesign rather than isolated use cases, and its 2026 organizational-transformation synthesis reports that only a small minority of organizations are yet redesigning work at that depth [4,5]. Recent theory points in the same direction. Agrawal et al. model AI that enhances worker productivity without automating tasks; Demirer et al. show that AI changes how steps are chained into tasks and jobs; Farach argues that coordination compression can generate endogenous task creation; and Loaiza and Rigobon find that tasks emerging in 2024 are more human-intensive on their EPOCH measure [1–3,19]. The IMF likewise stresses that technological change creates new tasks or occupations and that the diffusion of new skills requires worker movement across occupations and regions [18].

This matters for governance. A first observed plateau in a role is often a local organizational bottleneck—a sign that the role has not yet been redesigned far enough—not proof that the elevated layer must now shrink. Some roles will still face genuine displacement pressure. But the burden of proof should rest on the claim that the frontier cannot be moved further, not on the assumption that the first apparent ceiling is technologically final.

4. Local Ceilings, Moving Frontiers, and Design Capacity

We model a role family r as a bundle of four value-weighted task masses: routine substrate R_r , judgment and problem framing J_r , exception and oversight work E_r , and coordination or relational work C_r . Let automation intensity $\alpha \in [0, 1]$ index the degree to which a workflow's routine substrate has been delegated to AI, and let $d \in \mathcal{D}_r$ index feasible redesign choices and transition investments for role family r , including workflow redesign, widened decision rights, paid upskilling, service expansion, and apprenticeship preservation. For empirical use, we recommend measuring all task masses in either value-weighted weekly labor minutes or normalized shares of total weekly task mass, with the unit held fixed within a study. The purpose of the notation is not to imply universal parameter values. It is to define estimands that can be approximated from workflow logs, time-allocation audits, escalation records, quality reviews, and redesign documents over a stated observation window.

AI primarily compresses R_r , but deployment can also expand or reveal additional higher-level human work ΔJ_r , ΔE_r , and ΔC_r , while also creating new human-owned task mass G_r through model steering, quality assurance, customer recovery, knowledge curation, cross-functional coordination, and newly viable service scope [1–3,18,38]. We therefore define the role's *elevation space* (that is, the share of released routine substrate that can plausibly be reallocated upward into more accountable human work) and the corresponding *displacement pressure* as shown in Equation (1):

$$ES_r(\alpha, d) = \min\{R_r(\alpha), \Delta J_r(\alpha, d) + \Delta E_r(\alpha, d) + \Delta C_r(\alpha, d) + G_r(\alpha, d)\}, \quad (1)$$

$$DP_r(\alpha, d) = \max\{0, R_r(\alpha) - [\Delta J_r(\alpha, d) + \Delta E_r(\alpha, d) + \Delta C_r(\alpha, d) + G_r(\alpha, d)]\}.$$

If $ES_r(\alpha, d)$ is large, AI can shrink drudgery while preserving or deepening the human role through redesign. If $DP_r(\alpha, d)$ remains large even after serious redesign attempts, then genuine shrinkage pressure may exist. Under this interpretation, ES_r and DP_r are workflow-level estimands. A firm should therefore disclose the baseline window used to estimate pre-deployment task mass, the post-stabilization window used to estimate task reallocation after deployment, and the evidence sources used to classify routine, elevated, and newly created work. Appendix B provides a cross-role estimation template, and Appendix C works through one customer-support protocol.

A narrow, frozen job design d_0 may still generate a local ceiling. Let the role's local redesign checkpoint under design state d be

$$\tilde{\alpha}_r(d) = \sup\{\alpha \in [0, 1] : ES_r(\alpha, d) \geq DP_r(\alpha, d)\}. \quad (2)$$

This is the largest automation intensity for which the role can still be elevated fast enough *under that specific design*. But the relevant governance object is not just $\tilde{\alpha}_r(d_0)$ for today's narrow job description. It is the best frontier attainable over feasible redesigns:

$$\alpha_r^* = \sup_{d \in \mathcal{D}_r} \tilde{\alpha}_r(d).$$

Job cuts may indeed occur beyond $\tilde{\alpha}_r(d_0)$ for a frozen design d_0 . But that should be read as evidence that the role has reached a *local bottleneck*, not automatically as proof that the role has no further future. Under successful redesign, the realized elevated value need not peak at all over the relevant automation range; it can plateau, kink upward, or continue rising.

4.1. A Workflow-Level Estimation Protocol

The most important empirical question is not whether a role is exposed in the abstract, but whether redesign captures the released routine substrate in practice. A workable protocol has five steps. First, decompose the workflow into tasks and record current decision rights, escalation points, and quality gates. Second, estimate released routine substrate from the reduction in human time spent on draftable, classifiable, retrievable, or otherwise first-pass work after deployment. Third, estimate captured elevation from newly observed time in judgment, exception handling, coordination, and newly created human-owned tasks. Fourth, measure oversight burden directly rather than treating it as invisible residual labor. Fifth, re-measure service quality, error severity, novice progression, and internal mobility before drawing a redundancy conclusion. The point is not to force one universal metric on every workplace. It is to make the burden of proof concrete enough to audit.

Table 1 summarizes the frontier logic for the main role families discussed in the paper.

Table 1. High-exposure role families still contain an elevated human layer, and the key question is whether redesign keeps the frontier moving.

Role family	AI-compressed substrate	Elevated human layer	Local redesign checkpoint
Administration	Scheduling, note-taking, travel plans, template documents, inbox triage	Workflow orchestration, stakeholder routing, priority arbitration, meeting-to-execution follow-through	Stalls early only if the role is kept narrow and low-discretion
HR / recruiting	JD drafting, sourcing, screening, scheduling, FAQ handling	Interview calibration, talent advising, internal mobility, onboarding exceptions, employee relations	Moves outward with ambiguity, regulation, and people judgment
Marketing	First-draft copy, campaign variants, segmentation ideas, SEO baselines	Brand stewardship, experiment design, cross-channel learning, partner alignment, customer interpretation	Stalls in commodity content factories; rises in strategy-rich functions
Sales operations	CRM updates, proposal drafts, routine pipeline reporting	Deal orchestration, forecast integrity, strategic account planning, exception management	Stalls earlier when account complexity is low and highly transactional
Customer support	FAQ retrieval, routing, response drafts, transcript summaries	Escalations, emotion-rich cases, churn prevention, knowledge-base curation, root-cause analysis	Simple queues stall earlier; complex service layers retain more frontier movement
Finance ops	Invoice coding, reconciliations, reporting drafts, variance explanations	Controls, exception analysis, audit readiness, business partnering, scenario interpretation	Moves outward with compliance burden and exception intensity
Legal & compliance	Document review, first-pass drafting, policy retrieval, checklist generation	Precedent judgment, edge-case escalation, regulator interface, audit defense	Moves outward in regulated industries where accountability cannot be delegated
Coding	Boilerplate, unit tests, refactors, API glue, search-heavy debugging	Architecture, integration, security, eval design, code review, reliability ownership	Moves outward in complex systems; stalls sooner in repetitive maintenance
Research	Literature triage, summaries, baseline scripts, transcription, memo drafts	Question selection, evaluation design, source trust, synthesis, causal identification, significance judgment	High in frontier or ambiguous work; stalls sooner in repetitive desk research

5. High-Risk Role Families and Worked Elevation Paths

To keep the argument concrete, we use three running examples in the main paper—recruiting coordination, junior software engineering, and customer support—because they span people operations, technical production, and service work while remaining concrete enough to audit. Appendix A catalogs additional role families, Appendix D adds executive-assistance and research-analysis examples, and Appendix C gives one workflow-level protocol in more operational detail. The point is not that every exposed role can be preserved. It is that many roles cannot be evaluated honestly without first testing whether redesign captures the released routine substrate.

Worked example 1: Recruiting coordination.

A recruiting coordinator's role shrinks if leadership treats it as a throughput machine after AI compresses scheduling and outreach drafts; however, if treated as talent architecture, the role can move upward into interviewer calibration, candidate-experience design, and internal mobility support.

The elevated coordinator helps hiring managers refine job requirements dynamically, identifies edge cases that automated screening misses, and becomes partially responsible for whether the interview funnel is fair, interpretable, and decision-useful.

Worked example 2: Junior software engineering.

Reducing boilerplate and search-heavy debugging provides the capacity for juniors to shift toward supervised architecture mapping, system integration, code review participation, and reliability work [40]. Instead of spending hours writing redundant unit tests, junior developers can be mentored on evaluating AI-generated code for security vulnerabilities, on understanding how local patches affect system boundaries, and on learning the reliability consequences of seemingly small changes.

Worked example 3: Customer support.

Agent acceleration and generative routing [14] should purposefully redirect human time from simple queries toward complex escalations, retention saves, empathetic intervention, and knowledge-base curation. Freed from copying and pasting FAQ answers, the support agent can resolve multi-layered customer frustrations, repair broken institutional trust, and identify systemic product flaws that are invisible when support is treated as a narrow ticket-closing function. Appendix C shows how this case can be measured without relying on stylized percentages.

A useful boundary case is a narrow, highly standardized queue with low discretion and little service-recovery scope. These are the settings in which a rebuttal is most likely to succeed after redesign attempts, because the room to widen decision rights or create new human-owned task mass may truly be limited. Reported shifts toward AI-heavy moderation at ByteDance/TikTok are at least consistent with this lower-frontier class [24]. The point is not that such cases do not exist. It is that firms should have to show that a role belongs to this class rather than infer it directly from exposure.

What matters is not whether a role is broadly exposed, but whether the firm has honestly tested frontier-moving redesigns and adjacent transfer opportunities. A local ceiling appears not when AI can complete the clerical first pass, but when the organization has stopped expanding the higher-discretion work around it. Further detailed breakdowns across exposed role families are cataloged in Appendix A.

6. The Elevate-First Rule

By *AI-attributed layoffs*, we mean reductions justified by AI-driven efficiency or reallocations specifically meant to fund AI growth. Such layoffs should face a rebuttable presumption, requiring firms to pass an elevate-first test encompassing six core rules to demonstrate responsible transition governance. We propose that this test be instantiated as a compact *Workplace AI Transition Card*: a reporting artifact for deployment papers, field studies, board materials, and public disclosures. The card does not settle normative questions on its own, but it makes the transition legible enough to audit.

Rule 1: Task-level proof.

Firms must decompose automation claims into changed task boundaries, residual human accountability, and newly created oversight layers instead of relying on role-level slogans. This requires a precise accounting of what the AI system reliably handles, what exceptions it predictably generates, and which human workers are ultimately responsible for validating the outputs and assuming the liability of errors.

Rule 2: Paid upskilling.

Firms must offer paid upskilling by providing structured, on-the-clock domain and AI-literacy training [4,16,17,44,45]. Rather than merely offering optional after-hours course libraries, organizations must actively invest in teaching affected workers how to orchestrate these new systems, evaluate their outputs, and redesign their own daily workflows.

Rule 3: Internal mobility.

Wage-protected pathways to adjacent elevated work, such as quality assurance or AI-system supervision, must be guaranteed before executing layoffs. If the routine production in one department shrinks permanently, the firm must map the displaced talent to areas where human oversight is expanding, treating internal redeployment as a primary operational objective rather than an afterthought.

Rule 4: Apprenticeship preservation.

Organizations must redesign junior roles for supervised judgment instead of collapsing the talent pipelines through which future experts are formed [14,35,46–48]. If entry-level execution is automated, firms must deliberately construct new cognitive scaffolding, allowing junior employees to safely learn the underlying causal logic of their domain while managing AI tools under senior mentorship.

Rule 5: Worker consultation.

Active worker consultation is required to map failure modes, transition metrics, and required redesigns in direct dialogue with frontline staff. Frontline operators possess the tacit knowledge required to identify where AI models confidently fail in practice; their input is indispensable for designing realistic transition plans and ensuring that safety and service quality do not quietly degrade.

Rule 6: Financing transparency.

Firms must ensure financing transparency by disclosing the full financing chain and the alternatives considered if claiming that cuts are necessary to fund AI investments. When headcount reductions are explicitly framed as a mechanism to self-fund AI infrastructure or product teams, the company must demonstrate that it exhausted less destructive cost-saving measures—such as contractor rationalization or phased capital expenditure—first.

Table 2 synthesizes these requirements into a practical disclosure card.

Table 2. A Workplace AI Transition Card for AI-attributed layoffs.

Dimension	Minimum disclosure	Illustrative metrics or evidence
Task map	What tasks were automated, accelerated, or reallocated; what human-accountable work remains; and what baseline / post-stabilization windows were used.	Residual exception rates, escalation ownership, quality-control steps, customer handoff rules, work-flow scope, and rollout milestones.
Training	What protected, paid training affected workers received.	Hours per worker, completion and application rates, role-specific AI literacy versus supervisory skill [16,45].
Mobility	What elevated roles were opened before any layoffs.	Share of affected workers receiving internal offers, wage protection, transition duration, placement outcomes.
Apprenticeship	Whether junior, internship, or graduate-intake pathways were preserved or redesigned.	Intake numbers, junior-to-mid promotion rates, apprentice conversion, mentor ratios.
Consultation	Whether workers and managers were consulted and what design changes followed.	Documented consultations, appeal channels, incident reporting, redesign decisions.
Distribution of gains	How productivity improvements benefited workers and service quality.	Promotion pathways, reduced drudgery, better staffing for higher-order work, stability of quality metrics.
Financing logic	What AI investment or AI-role expansion the workforce change is funding, and what alternatives were considered first.	AI capex or product spend, AI-role openings, internal-fill rate, share of savings ring-fenced for training/wage protection, and non-layoff financing options evaluated.
Public reporting	Whether the transition was disclosed through human-capital and social-sustainability metrics.	ESRS S1, GRI labor training, ISO 30414, OECD responsible-business conduct alignment [45,49–51].

This standard is economically sound and aligns natively with core ESG and human-capital frameworks [45,49–53]. Replacing experienced local knowledge with external hires is a wasteful loop in tight talent markets [16,17,44]. Alternatively, utilizing staged capex, slower backfills, or transition reserves to fund redeployment acknowledges that AI-powered wages and revenue often rise concurrently [54]. Short-term cost discipline should not be allowed to obscure the long-run losses

in skill formation, trust, and adaptive capacity when firms bypass redesign for immediate payroll compression [6,7,9,31].

The position is also falsifiable in a concrete sense. It would weaken if firms could show, across broad role classes, that post-deployment workflows sustain quality, trust, safety, apprenticeship formation, and organizational adaptability even after substantial role deletion; or if elevate-first firms systematically underperform layoff-first firms on both productivity and resilience. Appendix O states the fuller conditions.

7. Alternative Views

These alternative views are not peripheral objections. They identify the precise conditions under which the rebuttable presumption should fail, narrow, or be qualified.

View 1: What if adaptive frontiers are genuinely low?

In hyper-optimized, low-margin sectors, the cost of retaining staff for marginal oversight may exceed the economic value of the elevated tasks. Some transactional clerical layers and heavily standardized queues may hit their adaptive frontier quickly. **Response:** Elevate-first does not deny this; it requires firms to show, with task-level evidence, that redesign and adjacent transfer were seriously attempted before elimination, and that they are pointing to a genuinely low adaptive frontier rather than to a frozen local bottleneck. When the frontier is genuinely low, the emphasis should shift to structured severance and external transition support rather than artificial job preservation. Narrow moderation or FAQ queues are the clearest candidates for this lower-frontier class [24].

View 2: Does role redesign concede first-mover advantage?

Organizations argue that delaying workforce restructuring to execute complex role redesigns cedes first-mover advantage to more aggressive rivals. Markets move fast, and firms may argue they cannot afford friction. **Response:** Rapid AI integration without capability retention can yield brittle systems. Durable advantage often relies on the contextual knowledge embedded in the retained workforce, especially when models require ongoing policy adaptation, exception handling, and service recovery. Recent theory also suggests that competition itself can produce an automation arms race that overshoots the socially efficient level of displacement [4,32].

View 3: Doesn't this simply delay the hollowing out of middle management?

Critics might contend that elevate-first merely delays the inevitable hollowing out of middle management. If routine reporting and coordination are automated, the remaining oversight roles may be too few to absorb the displaced middle layer. **Response:** The framework addresses this by redefining middle management: if a middle layer is automated, the relevant question is what higher-order ownership remains—cross-functional risk orchestration, exception governance, and continuous AI alignment—effectively shifting the value proposition from operational bottlenecks to strategic control.

View 4: Does this mandate the retention of low-value jobs?

Skeptics warn that mandating role elevation could institutionalize low-value monitoring roles where humans are kept merely to watch capable agents out of compliance rather than necessity. Some firms may also use AI as a convenient narrative for ordinary restructuring. **Response:** Both points reinforce the need for rigorous disclosure. The framework is not a blanket jobs guarantee; it asks firms to distinguish high-value oversight and exception handling from needless bureaucracy, ensuring that preserved roles genuinely contribute to institutional accountability.

View 5: Should AI-enabled flexible work be the default mode instead?

Another perspective suggests that rather than full-time redesign, the residual elevated tasks should facilitate a transition toward alternative employment models. As routine tasks compress, AI-

enabled flexible work could naturally become the default mode for knowledge work [55]. **Response:** This allows firms to retain essential human oversight, institutional memory, and complex reasoning capabilities on a fractional or decentralized basis rather than executing binary layoffs or forced full-time retention.

View 6: What if exposed workers cannot or do not want to elevate?

A fundamental assumption of the elevate-first rule is that workers in highly exposed roles possess the baseline capacity or desire to transition into judgment-heavy, orchestrating roles. Detractors argue this ignores the reality of varied worker aptitudes and preferences; not every administrative assistant wants to become a workflow architect. **Response:** Transition frameworks should therefore account for horizontal mobility into high-touch, AI-resistant, or differently skilled roles rather than assuming that every worker should move vertically into the same kind of elevated work.

8. Broader Implications for AI Evaluation and Deployment

If the elevate-first position holds, the ecosystem must rethink how it defines “beneficial AI.” Currently, AI is evaluated largely by automation rates, benchmark scores, and raw throughput. For the broader scientific and industry ecosystem, this implies a necessary shift: workplace AI should be evaluated by *worker elevation*, *frontier movement*, and *long-term sustainability*.

First, evaluations should capture *oversight burden*: the monitoring, verification, escalation, and accountability work that safe deployment inevitably creates. A benchmark that ignores these residual costs systematically overstates substitution potential. We therefore recommend a reporting profile rather than an unqualified single-number story. For a deployment d , let the core profile be

$$P_d = (\Delta N_d, \Delta Q_d, M_d, F_d, O_d, L_d),$$

where ΔN_d is novice lift, ΔQ_d is quality improvement, M_d is redeployment capacity, F_d is role-frontier movement, O_d is oversight burden, and L_d is entry-ladder loss. Where a single internal comparison score is useful, that profile can be scalarized as an optional *Elevation Impact Factor* (Equation (3)):

$$\text{EIF}_d = \lambda_1 \Delta N_d + \lambda_2 \Delta Q_d + \lambda_3 M_d + \lambda_4 F_d - \lambda_5 O_d - \lambda_6 L_d, \quad (3)$$

with the important caveat that the substantive recommendation is the disclosed profile and its components, not the claim that a universal scalar can settle redundancy decisions. The weights should therefore be pre-registered or at least justified, and the component metrics should always be reported alongside the scalar.

Second, reporting should capture the *financing topology* of deployment. Did productivity gains fund service improvement, worker elevation, and broader access, or were they immediately translated into payroll compression used to finance additional AI investment? By analogy to model cards, datasheets, and internal audit documentation in ML [56–58], we propose a *Workplace AI Transition Card* as the concrete NeurIPS-facing artifact for workplace deployments. Deployment papers and industry case studies should therefore report which role bundles are accelerated, which human responsibilities remain, whether the deployment moved the frontier outward, what oversight burdens were introduced, whether junior pathways were preserved, and whether the gains support redesign or merely justify substitution.

9. A Specific Research Agenda for Workforce Elevation

To operationalize the elevate-first framework, the scientific and organizational research communities must pursue a targeted agenda focused on capability measurement, longitudinal evaluation, interface design, and epistemic resilience. For NeurIPS-facing empirical work, this means reporting oversight burden, novice lift, apprenticeship effects, redeployment pathways, and whether measured performance gains came from redesign or from role deletion elsewhere.

First, the technical community should prioritize the development of **frontier-aware metrics** that estimate elevation space, residual displacement pressure, frontier movement, and supervision burden at a granular, role-specific level. Rather than evaluating AI purely on benchmark completion of isolated tasks, researchers should design evaluation protocols that distinguish local bottlenecks from genuinely low adaptive frontiers and that measure the specific human capital required to oversee, correct, and orchestrate these models in production.

Second, empirical deployment studies should shift from short-term throughput analyses to **longitudinal evaluations of workforce dynamics**. This involves tracking redeployment flows, the preservation of entry-level apprenticeship ladders, the evolution of decision rights over time, and whether the observed frontier keeps moving after initial deployment. By observing self-evolving evaluation and deployment protocols within live agentic workflows, longitudinal studies can map how human expertise evolves alongside continuous model updates.

Third, there should be tighter integration between **human-computer interaction system design** and labor economics. Interfaces should be explicitly designed to scaffold the transition from routine execution to supervised judgment, providing workers with the telemetry needed to act as effective operators. Concurrently, economic analyses of AI transitions should incorporate comprehensive transition accounting, comparing the true cost of immediate layoffs against alternatives like phased capital expenditure or transition reserves.

Fourth, the field should investigate the **long-term epistemic resilience** of organizations that aggressively compress their cognitive substrate. As AI absorbs the foundational layers of drafting and analysis, firms risk eroding the uncodified, tacit knowledge that workers traditionally acquire through routine practice. Future research should model the capability friction necessary to sustain human mastery, exploring how deliberate cognitive scaffolding can be integrated into AI architectures and deployment routines to prevent institutional forgetting at the frontier.

10. Conclusions

The wrong question for an AI transition is not “how many workers can now be removed?” but “how much of the released routine substrate can be converted into more accountable human work, how far can the frontier still be moved, and how is the transition being financed?” Firms should elevate high-risk roles before executing AI-attributed layoffs by redesigning jobs upward, funding paid upskilling, preserving apprenticeship ladders, opening wage-protected internal mobility, and disclosing when labor cuts are being used to finance AI investment. The deeper choice facing industry and the research community is between narrow productivity accounting—which converts AI gains quickly into headcount reduction while obscuring transition costs—and sustainable productivity, which uses those gains to build robust organizational capability. The necessary correction is therefore conceptual as much as policy-oriented. A first apparent ceiling is often a local redesign checkpoint, not a final technological verdict. Layoff-first AI strategy is a managerial choice, not a technological inevitability. An elevate-first presumption helps steer deployment toward expanding reliable human expertise rather than simply compressing payroll, and it keeps organizations focused on the harder but more valuable task: moving the frontier outward instead of declaring it closed.

Acknowledgments: Funding in direct support of this work: RIE2025 Industry Alignment Fund (Award I2301E0026) and the Alibaba–NTU Global e-Sustainability CorpLab.

Appendix A. Extended Role Catalog

The roles in Table A1 are drawn from occupations commonly flagged as exposed in ILO, OECD, WEF, and GPT-based exposure analyses, combined with our elevation-frontier lens [15,16,20,33–35]. The point is not that every role can be preserved intact. It is that many exposed roles still contain a meaningful elevated layer that is ignored when exposure is treated as destiny.

Table A1. Broader catalog of exposed roles and their potential elevation space.

Role family	AI-compressed substrate	Elevated human layer	Ceiling / caveat
Administrative assistant	Scheduling, note-taking, document formatting, inbox triage	Workflow orchestration, stakeholder routing, deadline management, executive context memory	Lower when role remains purely transactional
Executive assistant	Travel planning, brief drafting, follow-up drafting	Priority arbitration, confidential escalation, cross-team orchestration, principal leverage	Higher when discretion and trust are central
Reception / front desk	FAQ answering, routine routing, visitor logging	Exception handling, service recovery, facilities coordination, relationship continuity	Lower in fully standardized sites
HR coordinator	JD drafting, scheduling, first-pass screening	Candidate calibration, onboarding exceptions, people analytics interpretation, internal mobility support	Rises with compliance and ambiguity
Recruiting coordinator	Scheduling, sourcing outreach drafts, FAQ handling	Interview quality control, role calibration, candidate experience, funnel redesign	Lower in high-volume commodity hiring
Learning & development	Content drafting, quiz generation, scheduling	Skill diagnosis, coaching design, curriculum adaptation, manager enablement	Higher when reskilling is continuous
Marketing coordinator	Copy drafts, content variants, baseline analytics summaries	Experiment design, audience learning, campaign integration, brand stewardship	Lower in low-end content mills
Performance marketer	Ad variant generation, keyword ideation, routine reporting	Budget allocation, causal inference, incrementality judgment, channel portfolio decisions	Depends on measurement sophistication
Sales operations	CRM notes, proposal drafts, forecast writeups	Deal orchestration, forecast integrity, exception management, customer signal synthesis	Lower when account complexity is low
Customer support agent	FAQ retrieval, response drafts, routing, summaries	Escalations, empathy-heavy cases, churn prevention, root-cause identification, knowledge-base updates	Simpler queues have lower adaptive frontiers
Customer success associate	Routine check-ins, report drafting	Adoption strategy, stakeholder mapping, renewal risk judgment, product feedback synthesis	Higher in enterprise accounts
Bookkeeping / AP-AR	Coding invoices, matching, routine reconciliations	Controls, exception triage, vendor issue resolution, audit readiness	Heavily standardized environments may shrink more
Payroll specialist	Routine calculations, document collection, templated explanations	Exception adjudication, compliance interpretation, employee communication on edge cases	Compliance burden preserves some elevated layer
Procurement / vendor ops	RFP drafts, spend classification, comparison tables	Negotiation, supplier risk judgment, relationship management, cross-functional alignment	Lower in simple catalog procurement
Legal assistant / paralegal	Document review, first-pass drafting, retrieval	Precedent judgment, edge-case escalation, case strategy support, regulator interface	High accountability can keep adaptive frontiers relatively high
Compliance analyst	Policy retrieval, monitoring summaries, checklist generation	Incident response, control design, governance interpretation, audit defense	Higher in regulated industries
Trust & safety reviewer	Routine triage, policy lookup, simple moderation	Edge-case adjudication, policy refinement, escalation patterns, harm synthesis	Commodity queues may shrink; policy-heavy layers remain
Business analyst / BI	SQL drafting, dashboard drafts, memo summaries	Metric governance, problem framing, decision support, causal interpretation	Lower if the role is restricted to reporting only
Project / program coordinator	Status reports, action logs, scheduling	Dependency resolution, stakeholder alignment, execution risk management	Ceiling depends on authority to intervene
Technical writer	Draft documentation, changelog summaries, formatting	Information architecture, release coordination, audience adaptation, accuracy ownership	Lower for rote internal docs
QA / test engineer	Test-case drafting, regression scripting, bug summaries	Reliability ownership, failure analysis, release risk judgment, coverage strategy	Higher in safety- or uptime-critical systems
Junior software engineer	Boilerplate, tests, refactors, search-heavy debugging	System understanding, integration, code review participation, evaluation discipline	Lower in repetitive maintenance-only teams
Data engineer / MLOps	Template pipelines, config scaffolding, routine transformations	Reliability, governance, observability, infra trade-off ownership	High in production environments
Research analyst	Literature triage, transcription, baseline scripts, memo drafts	Question formulation, evidence evaluation, synthesis, significance judgment	Lower in repetitive desk research
Research assistant	Search, annotation aids, first-pass coding	Experimental execution, error analysis, documentation discipline, replication support	Depends on whether mentors redesign the role upward
Writer / journalist	First-draft copy, summaries, headline variants	Source trust, editorial judgment, investigation, interviewing, narrative accountability	Lower in commodity SEO writing

Appendix B. Workflow-Level Estimation Template

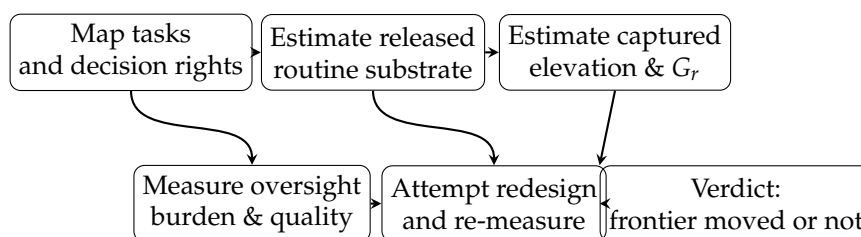
The earlier stylized simulation table is replaced here with an operational template. The point is not to assign universal percentages to roles. It is to show what should be observed, logged, and compared before a firm concludes that a workflow has reached a genuine adaptive frontier. Table A2 translates the main-text notation into observable workflow evidence.

Table A2. Workflow-level estimation template for the main role families. The point is to operationalize what should be measured before claiming a role has reached a genuine adaptive frontier.

Role family	Released routine substrate: observable signals	Captured elevation: observable signals	Oversight / quality checks	Typical redesign levers
Administration	Drop in weekly human minutes on scheduling, note formatting, inbox triage, travel booking, template production	Increase in stakeholder routing, priority arbitration, follow-through tracking, meeting-to-execution coordination	Missed-handoff rate, executive follow-up failures, deadline slippage, stakeholder satisfaction	Wider decision rights, owner-of-record status for follow-through, cross-team routing authority
HR / recruiting	Reduction in manual sourcing, scheduling, JD drafting, FAQ handling, status emails	Increase in interviewer calibration, candidate-experience repair, onboarding exceptions, internal mobility advising	Candidate-dropout rate, override rate on AI screens, hiring-manager satisfaction, time-to-acceptance	Role calibration rights, bias review, interview debrief ownership, mobility desk integration
Marketing	Reduction in first-draft copy, baseline analytics summaries, content variants, SEO scaffolds	Increase in experiment design, audience interpretation, cross-channel synthesis, partner coordination	Brand-consistency checks, campaign error severity, experiment-learning velocity, downstream conversion quality	Broader experimentation mandate, insight ownership, tighter links to product and sales
Sales operations	Reduction in CRM updating, proposal scaffolding, routine forecast writeups	Increase in deal orchestration, exception handling, forecast integrity work, strategic account support	Forecast-bias rate, exception aging, account-escalation outcomes, seller satisfaction	Access to account strategy reviews, exception queues, customer-signal synthesis responsibilities
Customer support	Reduction in human time on FAQ-only tickets, routing, templated response drafting, transcript summarization	Increase in escalations, churn prevention, service recovery, knowledge-base curation, root-cause analysis	Reopen rate, severe-incident rate, CSAT/NPS, escalation quality, hallucination cleanup time	Escalation authority, KB ownership, retention-save remit, product-feedback loop participation
Finance ops	Reduction in invoice coding, matching, reconciliations, first-draft variance explanations	Increase in controls work, audit readiness, exception analysis, business-partnering interpretation	Control failures, audit findings, exception aging, rework burden, sign-off latency	Control ownership, exception-routing authority, business-partner access
Legal & compliance	Reduction in first-pass review, policy retrieval, checklist generation, document comparison	Increase in edge-case adjudication, regulator interface, precedent judgment, audit defense preparation	Override rate, severe-compliance incidents, regulator response time, quality of legal sign-off	Escalation ownership, regulator-facing responsibilities, review of edge-case classes
Coding	Reduction in boilerplate writing, test scaffolding, repetitive refactors, search-heavy debugging	Increase in architecture mapping, integration work, code review, security reasoning, reliability ownership	Defect severity, rework time, post-release incidents, review burden, junior progression	Broader review rights, ownership of service boundaries, reliability and eval tasks
Research	Reduction in literature triage, baseline scripts, memo drafts, transcription and formatting	Increase in question framing, source trust evaluation, methodology choice, synthesis under ambiguity	Evidence-quality audits, replication errors, supervisor review load, novelty of insights	Research-design ownership, synthesis tasks, evaluation planning, ambiguity-heavy assignments

Appendix C. Illustrative Workflow Audit: Customer Support

Customer support is a useful anchor workflow because it combines routinized first-pass work with clearly auditable escalation, retention, and service-recovery layers. The protocol below is deliberately operational: it asks what a firm would have to measure before it could credibly claim that a support layer had reached a genuine adaptive frontier.

**Figure A1.** Workflow-transition decision pipeline. A credible redundancy claim should come only after the workflow has been mapped, measured, redesigned, and re-measured.

A support deployment moves the frontier outward only if released routine time is matched by captured elevation and new human-owned task mass *without* hidden deterioration in quality, safety, trust, or apprenticeship formation. In practice, that means using a disclosed baseline window, a disclosed post-stabilization window, and a redesign log that records what the organization actually changed between measurement rounds. The same logic can be adapted to recruiting, finance operations,

coding, and other role families; the point of the support example is simply to show that the paper's notation can be turned into an auditable protocol.

Table A3. One workflow-level protocol for operationalizing the main-text constructs. The key move is to estimate released routine substrate, captured elevation, and oversight burden from auditable traces rather than from stylized role percentages.

Construct	Observable definition in customer support	Typical evidence source
Released routine R_r	Reduction in weekly human minutes on FAQ-only tickets, simple routing, templated replies, and transcript summarization after rollout	Ticket timestamps, handle-time logs, queue tags, time-motion audit
Captured judgment ΔJ_r	Increase in human time on retention saves, exception resolution, ambiguous policy decisions, and root-cause diagnosis	CRM outcomes, escalation tags, save-rate workflows, manager review notes
Captured oversight ΔE_r	Increase in AI-output review, policy overrides, escalation validation, and severe-case ownership	QA review logs, override records, incident trackers
Captured coordination ΔC_r	Increase in cross-team handoffs, knowledge-base updates, callbacks to product or operations, and follow-through on recurring defects	Knowledge-base edit history, issue tracker, cross-functional tickets
New human-owned tasks G_r	Tasks created by deployment itself, such as prompt/eval maintenance, gap analysis, escalation taxonomy upkeep, or service-recovery playbook maintenance	Evaluation logs, documentation repos, quality-program records
Oversight burden O_d	Review time, hallucination cleanup, false escalations, duplicated contacts, and rework introduced by the system	QA queue, reopen rate, duplicate-contact logs, postmortem records
Quality guardrails	Changes in reopen rate, severe-incident rate, CSAT/NPS, first-contact resolution, and complaint severity	Support analytics, trust-and-safety logs, customer-survey systems
Entry-ladder effects L_d	Whether novice agents progress to higher-complexity queues, what supervisor load changes, and whether promotion pathways remain intact	Training logs, queue-allocation rules, mentor ratios, promotion records

Appendix D. Additional Worked Examples

Executive assistance.

Automating baseline scheduling, note-taking, and travel planning allows the role to be redesigned into complex cross-team orchestration, priority routing, and context memory. A local ceiling only appears if the firm refuses to widen the role's discretion. By delegating transactional scheduling to agents, the assistant becomes an information router who anticipates executive bottlenecks, manages stakeholder relationships, and ensures strategic follow-through on meeting outcomes.

Research analysis.

As AI absorbs literature triage and baseline drafting, the elevated human layer can expand into problem framing, methodological choice, source trust evaluation, and synthesis under ambiguity [42]. The analyst moves from data gathering toward epistemic gatekeeping, taking ownership of causal validity, source quality, and the interpretation of conflicting evidence.

Appendix E. Supplementary Discussion: How Firms Can Finance AI Without Defaulting to Layoffs

The strongest practical objection to elevate-first is budgetary: firms may say that models, compute, and product integration are expensive, so labor cuts are the only realistic source of funding. Recent cases from Atlassian, Block, HP, and Oracle make this objection concrete, because executives have explicitly or implicitly used labor cuts as a financing mechanism for AI expansion [6,7,9,31]. That logic is weaker than it first appears. The question is not whether AI investment is costly. It is whether layoffs are the *best* financing instrument once one accounts for hiring frictions, organizational memory, transition risk, and the value of preserving complementary capabilities. In labor markets where AI-relevant skills are scarce, firing experienced employees only to rehire adjacent talent later can be an unusually wasteful loop [16,17,44]. Even when a firm truly needs to reduce cost growth, it has more instruments available than immediate job deletion, and those instruments can buy time to move the frontier outward.

One instrument is time. Many firms can phase AI investment over several quarters and use natural attrition, slower backfills, and tighter external hiring to create budget room while redesign proceeds. Another is composition. Companies often spend heavily on contractors, duplicated software,

consulting, real estate, or low-value coordination layers that can be reduced before eliminating workers whose local process knowledge is hard to replace. A third is allocation of gains. If AI reduces cycle time or expands service capacity, part of the resulting surplus can be booked into a transition reserve that funds paid training, wage protection, and temporary redeployment pools. That reserve is not charity; it is a way of converting short-run efficiency into long-run absorptive capacity.

The sequencing matters. When firms cut first and redesign later, they remove exactly the people most able to explain which steps can safely be automated, which customers or cases are atypical, where compliance risk sits, and what new roles should exist after deployment. By contrast, when firms keep workers in the loop during transition, they gain a richer source of local knowledge and a more credible path to adoption. This is one reason public-sector and sectoral responses are focusing on mobility and job redesign rather than only on replacement. The aim is not to freeze every role in place, but to avoid paying for AI by destroying the organizational complements that make AI useful.

Firms also have choices about what they do with the productivity that AI creates. One option is extractive: convert most gains into thinner payroll. Another is expansive: use those gains to raise service quality, shorten response times, widen product coverage, improve compliance, or move workers upward into judgment-heavy tasks. The expansive path is often under-appreciated because managers can count headcount savings immediately, whereas the value of better quality, stronger resilience, and broader capability arrives less theatrically. But that does not make the latter less real. Klarna's public recalibration is instructive here: after aggressively leaning into cost-cutting narratives, its CEO later said the company may have gone too far, too soon [59]. That is exactly the kind of reversal we should expect when firms discover that customer trust, service quality, and organizational learning are complements rather than leftovers.

In other words, elevate-first is not a demand that firms ignore cost. It is a claim about *cost accounting*. The relevant comparison is not "layoffs versus no discipline." It is "layoffs versus alternative financing and redesign strategies once transition costs are counted honestly." On that comparison, layoff-first often looks less like necessity and more like the easiest item to justify on a spreadsheet because its savings are immediate while its long-run losses are organizationally diffuse. A stronger governance rule is therefore warranted precisely because the short-run accounting is biased toward visible payroll savings and against invisible losses in skill formation, trust, and adaptive capacity.

Appendix F. Supplementary Discussion: What Would Count as a Genuine Rebuttal?

Our claim is intentionally strong, so it should also be clear about what would count against it. An employer could legitimately argue for AI-attributed layoffs if it could show, with evidence rather than slogans, that a role family's core responsibilities had truly contracted after accounting for oversight, exception handling, quality control, and accountability; that structured paid training and internal mobility had been offered and measured; that apprenticeship effects had been considered; and that service quality, trust, safety, and legal compliance were not being quietly offloaded onto a smaller residual workforce. In other words, the rebuttal standard is demanding but not impossible.

This standard is stricter than today's public discourse, where it is often enough to announce that AI has improved productivity and then infer that labor is now dispensable. Productivity alone is not enough. A firm can become more productive and still owe stronger transition support to workers precisely because the productivity gain gives it more room to finance redesign responsibly. Nor is it enough to point to isolated successful automations. The relevant unit is the post-deployment workflow as a whole: what work disappeared, what work intensified, who now bears risk, and whether the organization has preserved the capacity to learn from mistakes.

A serious rebuttal would also need to reckon with distribution. If a firm's AI deployment raises throughput but reduces junior hiring, weakens promotion ladders, concentrates remaining work into higher-stress exception handling, or shifts monitoring burdens onto a smaller set of employees, then the gain is not well described as simple efficiency. It is a redistribution of risk and opportunity

inside the organization. That redistribution may sometimes be justified, but it should be named and defended rather than hidden behind a generic story that “AI can do more now.” This is exactly why human-capital and social-sustainability disclosure belongs near the center of the debate: it makes the transition legible enough to evaluate rather than merely admire.

The same logic clarifies why we frame elevate-first as a *rebuttable presumption* rather than an absolute prohibition. There will be cases in which a product line is closed, a workflow genuinely disappears, or a firm’s financial condition is too weak to support a long transition. But those are arguments for transparent exception handling, not for treating payroll compression as the default definition of AI success. The more visible and normalized AI-attributed layoffs become, the more important it is to distinguish a genuine rebuttal from a convenient managerial narrative.

Appendix G. Supplementary Discussion: Why Entry-Ladder Preservation Is a Technical Issue, Not Only a Social One

One reason the elevate-first rule may sound more radical than it is that many firms still treat entry-level work as expendable by default. But in AI-intensive organizations, the entry ladder is not merely a social obligation or a recruiting nicety. It is part of the technical capability stack. Junior roles are where firms build future evaluators, operators, managers, policy owners, and domain experts. If AI removes a large share of routine production from those roles, the right response is not automatically to delete them. The better response is often to redesign them so that juniors spend less time on mechanical production and more time on supervised verification, synthesis, escalation, and customer or stakeholder context. That is exactly the kind of redesign that lets AI broaden expertise rather than concentrate it [14,40,41,60].

The alternative is a brittle organizational structure in which firms keep a thinner layer of already-experienced workers while starving the pipeline that would produce the next cohort. That is risky even on narrow business grounds. Many AI deployments require persistent local adaptation: prompts change, interfaces evolve, policies shift, customer expectations move, and model behavior is non-stationary. Organizations need people who can learn these systems from the inside and gradually take on more judgment-heavy tasks. If the junior layer disappears, the firm can become dependent on external hiring for roles that used to be developed internally. In markets where AI-relevant talent is already scarce, that dependence is expensive and unstable [16,17,44].

Entry-ladder erosion is also a measurement problem. Public evidence already suggests that the pressure is real: the World Economic Forum reports a substantial global decline in entry-level postings, and the World Bank finds especially sharp reductions in substitutable entry-level white-collar roles alongside growth in AI-related postings [35,46,48]. Yet most discussions of AI productivity still stop at the individual task or worker. They rarely ask whether the deployment preserves a credible training path into expert work. That omission matters. A system that makes one experienced worker faster while reducing the opportunities through which future experienced workers are formed may look efficient in the short run while degrading capability over a longer horizon.

This is why entry-ladder preservation belongs inside the technical evaluation agenda. A well-designed workplace-AI system should ideally let organizations trust juniors with more responsibility *under appropriate oversight*, not render juniors unnecessary altogether. That is a directly testable proposition. Researchers and deploying firms can examine whether the system raises novice quality, reduces time-to-proficiency, improves supervisor leverage without hollowing out supervision, and sustains promotion pipelines over time. Once framed this way, preserving apprenticeships and junior roles is not peripheral to beneficial AI. It is one of the clearest ways to tell whether a deployment is actually expanding reliable human capability or merely compressing payroll.

Appendix H. Related Work

While not an exhaustive review, this framework builds on several adjacent bodies of scholarship. The first is the task-based view of technological change. Classic work by Autor, Levy, and Murnane,

followed by Autor and by Acemoglu and Restrepo, argues that technology changes task composition and the division of labor more often than it cleanly deletes occupations [36–38]. Recent AI-specific exposure work from the ILO and large-scale usage evidence from Claude conversations reinforce the importance of distinguishing task exposure from full-job elimination [15,33,43].

A second adjacent literature studies AI productivity, complementarity, and organizational redesign. Field and experimental evidence shows that generative AI can lift worker productivity, often with especially large gains for less experienced workers or for narrower workflow components [14,39–41]. More recent theory makes the dynamic point explicit: AI can enhance worker productivity without automating tasks, change how steps are chained into jobs, and create new human-owned work by compressing coordination costs [1–3]. Loaiza and Rigobon add complementary empirical evidence that newly emerging tasks are more human-intensive on their EPOCH measure [19]. That literature is often read as support for substitution. Our argument is that it more naturally supports redesign and elevation when the complementary oversight, exception handling, coordination, and new-task channels are made visible.

An adjacent methodological stream asks what workplace-AI evidence should count as evidence in the first place. One contribution argues that human–AI productivity claims should be reported as time-to-acceptance under explicit acceptance tests rather than by raw draft speed alone [61]. Related work on harness engineering makes a complementary point: measured language-agent performance depends heavily on the surrounding harness layer of control, agency, and runtime rather than on the base model alone [62]. In our terms, both contributions reinforce the idea that apparent substitution is workflow-dependent and that residual verification and oversight labor must be measured directly.

Another adjacent stream examines how agentic systems reshape research work and public AI ecosystems. Recent work on automated research argues that AI can move human contribution upward from direct experiment execution toward question selection, evaluation, and research direction [63]. OpenClaw surveys a public agent ecosystem in the wild, while *Let Papers Flow* considers how autonomous review pipelines could reshape scientific throughput and the organization of scholarly labor [64,65]. These works are more infrastructural than our argument here, but they are consistent with the broader claim that AI often changes the locus of human work rather than simply deleting it.

A further adjacent literature develops auditable knowledge and reporting infrastructure for sustainability-related AI systems. ESGenius and MMESGBench benchmark language and multimodal models on ESG tasks [66,67]. SSKG Hub and KG4ESG construct expert-guided knowledge-graph infrastructure for sustainability standards and ESG semantics, and ESGlass argues for more glass-box, provenance-aware sustainability reports [68–70]. At a more general ML-governance level, model cards, datasheets, and internal algorithmic-audit documentation show how compact reporting artifacts can improve transparency around deployment assumptions and failure modes [56–58]. Although these works are not labor-market papers, they support our emphasis on inspectability, expert oversight, and governance-bearing workflow artifacts in high-stakes deployment.

A third body of work concerns organizational complements and human-capital formation. Prior evidence on information technology and workplace organization shows that returns from digital tools depend strongly on complementary practices and skilled labor [71]. Deming’s work on the growing importance of social skills helps explain why communication, coordination, and judgment remain valuable complements even when routine production is partly automated [72]. Autor’s recent middle-class-jobs argument can be read in a similar spirit: AI can be deployed to broaden expertise rather than merely to compress payroll [60].

Finally, there is a growing policy and practice literature on AI, skills, and responsible transition. OECD’s 2025 skills brief, IMF’s 2026 analysis, the UK DSIT labour-market survey, and related public-sector materials all frame the AI transition as a workforce-development challenge as much as an automation challenge [16,17,44,73]. OECD’s broader synthesis work on AI and work points in the same direction [74]. Our contribution is to convert that emerging descriptive literature into a stronger prescriptive claim: firms should adopt a rebuttable presumption against AI-attributed layoffs.

Appendix I. Case-Coding Taxonomy for Company Evidence

The company cases used in the main text do not all instantiate the same mechanism. Table A4 therefore codes them into four classes: (A) direct automation displacement, (B) AI-financing or role-mix reallocation, (C) mixed restructuring with re-skilling or renewed hiring, and (D) counterexample or capability-building case. The coding is interpretive rather than mechanical, but it makes explicit what claim each case is doing in the argument. Evidence-tier labels distinguish cases grounded mainly in official filings, official memos, earnings materials, or government documents (Tier 1) from those relying primarily on high-quality secondary reporting (Tier 2).

Table A4. Case-coding taxonomy for the company evidence used in the paper. The goal is to distinguish direct automation stories from financing, mixed restructuring, and counterexample cases rather than treating all AI-linked workforce changes as identical.

Organization	Primary code	Evidence tier	How the paper uses the case
Atlassian	B	1	Explicit self-funding of AI and enterprise-sales investment through cuts elsewhere [6,75].
Block	B	1	Public conversion of AI-enabled productivity and intelligence-tools rhetoric into workforce compression [7,76].
Workday	C	1	Portfolio reallocation toward AI, product, cybersecurity, and sales talent rather than a pure "AI replaced the job" story [8,77].
HP	B	1	AI adoption initiative tied to large expected savings and future headcount reduction [9,78].
SAP	C	1	AI-driven restructuring paired with voluntary programs and internal re-skilling rather than simple net deletion [10,29].
Salesforce	B/C	1	Repeated role-mix reallocation around Agentforce growth, with commercial conditions also relevant [11,30,79,80].
Oracle	B	2	Capital-allocation case in which cuts help finance AI infrastructure expansion [31].
Meta	B/C	2	Large layoffs paired with AI-first reorganization in a highly profitable firm [12,81].
Microsoft	B	2	AI savings made legible through labor reduction, especially in call-centre reporting [82,83].
Amazon	B/C	1	Cascading corporate cuts in an AI-efficiency context, but with broader reorganization pressures also in play [84,85].
Baidu	C	2	Uneven internal redistribution protecting AI and cloud roles while cutting elsewhere [23].
ByteDance/TikTok	A	2	Stronger substitution-pressure case, especially for standardized moderation layers [24].
Klarna	A/C	2	Early AI cost-cutting rhetoric followed by a later public recalibration [59,86].
Alibaba	D	2	Counterexample in which aggressive enterprise automation coexists with renewed hiring [25,26].
Tencent	D	2	Capability-building case centered on AI talent acquisition and capex expansion [27].
Xiaomi	D	2	Capability-building case centered on large AI investment rather than labor retrenchment [28].

Appendix J. Selected 2024–2026 Company Evidence on Explicit AI-Financed Cuts, Reallocation, Rehiring, and Capability-Building

The cases in Table A5 are used to distinguish at least four patterns: explicit self-funding of AI through layoffs; selective cuts combined with AI hiring; AI-driven restructuring paired with internal re-skilling; and counterexamples in which firms expand talent while automating workflows. Wherever possible, the table prioritizes company filings, shareholder letters, earnings releases, or official memos, using Reuters when primary materials are not public.

Table A5. Selected company cases used in the paper.

Date	Organization	Concrete fact	Why it matters
Jan. 2024	SAP [10,29]	Officially announced a company-wide restructuring affecting about 8,000 positions to focus on Business AI and AI-driven efficiencies, while expecting roughly similar year-end headcount through re-investment and internal re-skilling.	Strong case that even when AI-driven restructuring is real, firms can still pair it with mobility and re-skilling rather than treating net labor deletion as the only option.
Aug. 2024	Klarna [86]	CEO highlighted AI chatbots as helping shrink headcount.	Early public example of AI-based labor-saving rhetoric being used as a success signal.
Dec. 2024	Salesforce [87]	Reuters reported that Salesforce had closed 1,000 paid Agentforce deals and was leaning hard into an agentic-sales narrative.	Important because later workforce changes were framed against a fast-growing AI-product strategy.
Feb. 2025	Salesforce [11,30]	Cut more than 1,000 roles while simultaneously hiring workers to sell AI products; by FY2026 results, Agentforce ARR had reached \$800 million and 29,000 deals.	Clear role-mix reallocation: cuts did not occur in the absence of genuine AI business expansion.
Sep. 2025	Salesforce [80]	Later Reuters coverage reported weak current-quarter revenue guidance.	Useful reminder that AI-centered restructuring narratives often unfold alongside ordinary commercial pressure rather than pure automation logic alone.
Feb. 2026	Salesforce [30,79]	Reuters reported another round of cuts affecting fewer than 1,000 roles even as Salesforce posted record fourth-quarter fiscal 2026 results.	Suggests that AI-linked workforce reallocation can persist even when near-term financial performance remains strong.
Feb. 2025	Workday [8,77]	Official filings describe an approximately 8% workforce reduction meant to prioritize investments and durable growth while the firm continues competing for AI, product, cybersecurity, and sales talent.	Shows how AI-era cuts are often about portfolio reallocation, not a declaration that people have become broadly unnecessary.
Mar. 2025	Alibaba [25]	Chairman said the firm had reached the bottom and would “reboot and rehire” after prolonged headcount decline.	High-profile counterexample to the claim that AI competition naturally forces ongoing payroll compression.
Jul. 2025	Microsoft [82,83]	Nearly 4% layoffs; Reuters later reported over \$500 million in AI savings in call centres.	Clear case of AI productivity gains being made legible through labor reduction.
Sep. 2025	Klarna [59]	CEO said the company may have gone too far, too soon in using AI for cost cuts.	Important reversal signal against naive substitution narratives.
Oct. 2025	Amazon [84]	About 14,000 corporate roles cut in a shakeup driven in part by AI adoption.	Large white-collar reduction linked to AI adoption and efficiency logic.
Nov. 2025	Baidu [23]	Reuters reported large-scale layoffs affecting multiple units, with some teams facing cuts as high as 40%, while AI- and cloud-related positions were largely protected.	Key case of uneven internal redistribution toward AI-linked lines.
Nov. 2025	HP [9,78]	Announced a company-wide AI adoption initiative expected to generate \$1 billion in run-rate savings while reducing gross headcount by 4,000–6,000 by fiscal 2028.	Evidence that explicit AI-financing logic has spread beyond software platforms into broader tech operations.
Jan. 2026	Amazon [85]	Another 16,000 roles cut, taking the two rounds to roughly 30,000 corporate reductions.	Shows how AI-era restructuring can cascade across multiple rounds.
Feb. 2026	Block [7,76]	Shareholder letter said a significantly smaller team using intelligence tools could “do more and do it better”; Block then cut over 4,000 jobs, nearly half its workforce.	One of the clearest public cases of AI-enabled productivity being converted directly into payroll compression.
Mar. 2026	Atlassian [6,75]	CEO told staff the company was restructuring to “self-fund further investment in AI and enterprise sales,” affecting about 1,600 roles.	Especially strong evidence for the claim that some companies are explicitly financing AI-related roles through cuts elsewhere.
Mar. 2026	Oracle [31]	Reuters reported planned thousands of job cuts as Oracle faced a cash crunch from a major AI data-centre expansion effort and slowed hiring in parts of its cloud division.	Highlights the capital-allocation channel: layoffs can finance AI infrastructure as much as direct task automation.
Mar. 2026	Alibaba [26]	Wukong enterprise platform automates document, spreadsheet, transcription, and research tasks via an agent interface.	Shows that aggressive automation and rehiring can coexist.
Mar. 2026	Tencent [27]	About 79 billion yuan capex in 2025 and explicit ramp-up in AI talent acquisition.	AI transition can mean expansion of talent and infrastructure rather than labor deletion.
Mar. 2026	Xiaomi [28]	At least 60 billion yuan pledged for AI over three years.	Another case of capability-building rather than retrenchment.

Table A5. Cont.

Date	Organization	Concrete fact	Why it matters
Mar.–Apr. 2026	Meta [12,81]	Reuters first reported in March on sizeable planned layoffs and later reported a first wave of roughly 8,000 layoffs with further cuts possible later in 2026, alongside AI-driven reorganization into Applied AI and related AI-first units.	Shows that even highly profitable firms may pair large AI investment with deep workforce reduction.
2024–2025	ByteDance / TikTok [24]	Hundreds of moderation jobs cut as more AI moderation was introduced.	Illustrates substitution pressure in trust-and-safety work.

Appendix K. Labor-Market, Productivity, and Skills Evidence

Table A6 summarizes key labor-market, productivity, and skills evidence showing that the transition bottleneck is redesign rather than broad redundancy.

Table A6. Institutional evidence that the current bottleneck is redesign and skills transition more than broad labor redundancy.

Source	Concrete figure or finding	Relevance to the position
BLS productivity [88]	Annual average nonfarm business productivity increased 2.1% from 2024 to 2025.	Rising productivity does not by itself justify layoffs; it sharpens the distributional question of who captures AI-era gains.
BLS JOLTS [89]	Revised January 2026 data showed roughly 7.2 million U.S. job openings and 5.3 million hires.	Broad labor demand remained substantial; the story is not one of generalized labor redundancy.
ILO [15,33]	One in four workers globally have some GenAI exposure; most exposed jobs are more likely to be transformed than made redundant.	Supports the task-transformation view over the full-job-elimination view.
OECD [16,74]	One in three vacancies in OECD economies have high AI exposure, but only about 1% require complex AI skills; only a small fraction of analyzed training courses include AI content.	Most workers need AI literacy and transition capacity, not elite model-building skills.
IMF [17]	About one in ten vacancies in advanced economies now demands at least one new skill.	The transition is as much about new-skill creation as about task automation.
WEF [34,35,46]	Firms expect both reductions in AI-exposed roles and substantial hiring of AI-skilled workers; entry-level postings fell 29% worldwide since January 2024.	Consistent with a redesign bottleneck and an apprenticeship problem rather than simple redundancy.
World Bank [48]	AI-related postings in South Asia grew faster than other postings and carried a wage premium, while the most substitutable entry-level white-collar listings fell by around 20%.	Highlights the entry-ladder paradox at the center of our argument.
Singapore MOM [90]	1.58 vacancies per job seeker in Dec. 2025; 49.3% of vacancies newly created.	Concrete country case showing that an AI-intensive labor market can stay tight while job content changes.
UK DSIT [44]	97% reported at least one AI skills gap; 35% struggled to fill AI roles.	Reinforces that the present scarcity is transition capacity, not too many workers.
China policy / Reuters [91–93]	12.7 million university graduates in 2026; AI-related training, internships, and more than 12 million targeted urban jobs were paired with a weak youth labor market.	Shows why entry-ladder preservation matters when graduate absorption is already fragile.

Appendix L. Government and Public-Institution Cases Showing an Alternative Path

Table A7 outlines concrete measures from governments and public institutions treating AI transition as a redesign, literacy, and mobility problem.

Table A7. Examples of governments and public institutions treating AI transition as a redesign, literacy, and mobility problem.

Jurisdiction / source	Concrete measure	Why it matters
Singapore IMDA [94]	Official position that AI should “not replace jobs, but create better, safer and more rewarding jobs”; more than two-thirds of AI-using firms planned to prioritize training and upskilling existing workers.	Provides a direct policy articulation of the people-first norm advocated here.
Singapore MOM [90]	TalentTrack+ and related workforce tools are aimed at internal mobility, skills-based hiring, and a “plug-train-play” approach.	Shows how public tools can encourage re-deployment instead of ready-made external hiring.
Singapore WSG / MAS / IBF [95]	The Finance Jobs Transformation Map expects most roles to be augmented and names new posts such as AI/GenAI Policy and Ethics Officer.	Concrete sectoral framework for redesign instead of deletion.
Thailand Ministry of Labour + Microsoft [96]	Partnership to deliver AI skills and certifications to 150,000 workers nationwide.	Large-scale public-private attempt to expand transition capacity rather than shrink labor.
China State Council / ministries [92]	March 2026 package links graduate employment to AI training, internship recruitment, social-insurance subsidies, and more than 12 million new urban jobs.	Official attempt to pair AI-led upgrading with labor absorption and human development.
China Central Economic Work Conference [97]	Called for combining investment in physical assets with investment in human capital while advancing the AI Plus Initiative.	Explicitly links AI-led development with human-capital formation.
China policy / Reuters [91]	Reuters reported a Beijing labor ruling that dismissing employees solely to replace them with AI is illegal.	Indicates that substitution pressures are already meeting worker-protection limits.
European Union [49,98]	AI literacy obligations and workforce disclosure rules connect deployment to staff capability and employability.	Suggests that AI readiness can be treated as a governance and compliance issue.
United Kingdom [44,73]	Skills England published employer-facing AI-skills tools; DSIT documented acute shortages and expanding apprenticeship efforts.	Shows a government treating the AI transition as a workforce-development challenge.

Appendix M. Elevate-First as a Social Sustainability and Human-Capital Reporting Standard

Table A8 maps existing ESG, sustainability, and human-capital frameworks to workforce transitions, supporting an elevate-first rule.

Table A8. How existing ESG, sustainability, and human-capital frameworks support an elevate-first rule.

Framework / source	Concrete requirement or principle	Implication for AI workforce transitions
ESRS S1 [49]	Requires disclosure of training and skills-development metrics and continued employability.	AI transitions should be reported in terms of employability, training intensity, and career development rather than only headcount savings.
GRI labor training project [45]	Treats training as an investment in employees and the organization’s future; emphasizes career transitions and paid time for training.	Supports the view that AI-related re-skilling should be protected, funded, and linked to real career mobility.
ISO 30414 [50]	Human-capital reporting should make transparent the contribution of people to the organization and support workforce sustainability.	Reinforces that AI transition is a human-capital reporting issue, not only a technology strategy issue.
OECD Guidelines [51]	Ask firms to provide training, cooperate with worker representatives, and mitigate major employment effects.	Strongly supports notice, consultation, re-training, and mitigation before AI-attributed layoffs.
UN Global Compact [99]	Social sustainability concerns business impacts on people, stakeholder relationships, and decent jobs.	Places AI workforce transition squarely inside the “S” of ESG.
ILO just transition [52]	Frames transition policy around more and better jobs, decent work, rights, and social protection.	Provides a transferable logic: productivity gains should not come from dumping social costs onto workers.
SDG 8 [53]	Links technological upgrading to productive employment, decent work, and labor rights.	Suggests that a sustainability-consistent AI transition must improve productivity <i>and</i> protect work quality and labor rights.

Appendix N. Illustrative Chinese Company Human-Capital and Upskilling Disclosures

Table A9 highlights reported metrics from recent Chinese technology-company ESG reports, illustrating how such transitions can be made auditable.

Table A9. Examples from recent Chinese technology-company ESG reports showing that worker-development metrics are already auditable.

Company / report	Reported metric	Relevance to elevate-first AI transition
Tencent 2024 ESG [100]	98.7% of male employees and 99.0% of female employees received training; average training hours were 37.3 and 39.9, respectively.	Illustrates that firms can report the training breadth and intensity needed to test whether transition is capability-building.
Baidu 2024 ESG [101]	General-staff training coverage reported as 100%; average training hours reported as 31.6; the report also described a large-scale "Second-Skill Learning Platform."	Shows that AI-literacy and second-skill infrastructure can be scaled and publicly reported.
JD.com 2024 ESG [102]	Talent-development training coverage reported as 100%; average training hours were 78.8 for technical employees, 35.3 for management, and 41.9 for non-management employees.	Demonstrates differentiated capability-building by role rather than elite-only training.
Alibaba public disclosures [25,26]	Publicly signaled both enterprise automation and renewed hiring.	Even without a single AI-transition metric, public disclosures already reveal whether firms pair automation with labor expansion.

Appendix O. Why the Position Is Falsifiable

This framework would be invalidated, or at least weakened, under several conditions. If future evidence showed that broad classes of firms can reliably delete large role families after AI deployment *without* meaningful losses in quality, trust, safety, learning, entry-ladder formation, or long-run organizational adaptability, the elevate-first presumption would weaken. If firms that pursue layoff-first strategies systematically outperform elevate-first firms on both productivity and resilience over time, that would also count against the thesis. Conversely, if job redesign, retraining, and internal mobility repeatedly prove cheaper and more durable than external hiring plus layoffs, the case for elevate-first grows stronger.

References

1. Agrawal, A.K.; McHale, J.; Oettl, A. Enhancing Worker Productivity Without Automating Tasks: A Different Approach to AI and the Task-Based Model. Working Paper 34781, National Bureau of Economic Research, 2026. Issued January 2026; <https://www.nber.org/papers/w34781>, <https://doi.org/10.3386/w34781>.
2. Demirel, M.; Horton, J.J.; Immorlica, N.; Lucier, B.; Shahidi, P. Chaining Tasks, Redefining Work: A Theory of AI Automation. Working Paper 34859, National Bureau of Economic Research, 2026. Issued February 2026; <https://www.nber.org/papers/w34859>, <https://doi.org/10.3386/w34859>.
3. Farach, A. AI as Coordination-Compressing Capital: Task Reallocation, Organizational Redesign, and the Regime Fork. *arXiv preprint arXiv:2602.16078* 2026. <https://arxiv.org/abs/2602.16078>.
4. World Economic Forum. AI at Work: From Productivity Hacks to Organizational Transformation. <https://www.weforum.org/stories/2026/01/ai-at-work-insights/>, 2026. World Economic Forum story summarizing the community paper; published January 15, 2026; accessed April 18, 2026.
5. World Economic Forum. Organizational Transformation in the Age of AI: How Organizations Maximize AI's Potential. <https://www.weforum.org/publications/organizational-transformation-in-the-age-of-ai-how-organizations-maximize-ais-potential/>, 2026. Published March 16, 2026; accessed April 18, 2026.
6. Atlassian. An Important Update on Our Team. <https://www.atlassian.com/blog/announcements/atlassian-team-update-march-2026>, 2026. Published March 11, 2026; accessed April 18, 2026.
7. Block, Inc.. Q4 2025 Shareholder Letter. <https://investors.block.xyz/financials/quarterly-earnings-reports/default.aspx>, 2026. Q4 2025 materials page; accessed April 18, 2026.

8. Workday, Inc.. Form 10-K for the Fiscal Year Ended January 31, 2025. <https://investor.workday.com/financials/sec-filings/default.aspx>, 2025. Official filing page; accessed April 18, 2026.
9. HP Inc.. HP Inc. Reports Fiscal 2025 Full Year and Fourth Quarter Results. <https://investor.hp.com/news-events/news/news-details/2025/HP-Inc--Reports-Fiscal-2025-Full-Year-and-Fourth-Quarter-Results/default.aspx>, 2025. Published November 25, 2025; accessed April 18, 2026.
10. SAP. SAP Updates Its Ambition 2025 and Announces Transformation Program for 2024. <https://news.sap.com/2024/01/sap-updates-its-ambition-2025-and-announces-transformation-program-for-2024/>, 2024. Published January 23, 2024; accessed April 18, 2026.
11. Reuters. Salesforce to Cut 1,000 Roles, Bloomberg News Reports. <https://www.reuters.com/technology/salesforce-cut-1000-roles-bloomberg-news-reports-2025-02-04/>, 2025. Published February 4, 2025; accessed March 21, 2026.
12. Reuters. Exclusive: Meta Targets May 20 for First Wave of Layoffs; Additional Cuts Later in 2026. <https://www.reuters.com/world/meta-targets-may-20-first-wave-layoffs-additional-cuts-later-2026-2026-04-17/>, 2026. Published April 17, 2026; accessed April 18, 2026.
13. Reuters. Companies Cutting Jobs as Investments Shift Toward AI. <https://www.reuters.com/business/world-at-work/companies-cutting-jobs-investments-shift-toward-ai-2026-04-15/>, 2026. Updated April 15, 2026; accessed April 18, 2026.
14. Brynjolfsson, E.; Li, D.; Raymond, L. Generative AI at Work. *The Quarterly Journal of Economics* **2025**, *140*, 889–942.
15. International Labour Organization. Generative AI and Jobs: A 2025 Update, 2025. Research Brief.
16. Organisation for Economic Co-operation and Development. Bridging the AI Skills Gap: Is Training Keeping Up? https://www.oecd.org/en/publications/bridging-the-ai-skills-gap_66d0702e-en.html, 2025. OECD policy brief; publication page; accessed April 18, 2026, <https://doi.org/10.1787/66d0702e-en>.
17. International Monetary Fund. New Skills and AI Are Reshaping the Future of Work. <https://www.imf.org/en/blogs/articles/2026/01/14/new-skills-and-ai-are-reshaping-the-future-of-work>, 2026. IMF Blog, January 14, 2026; accessed March 21, 2026.
18. Jaumotte, F.; Kim, J.; Koll, D.; Li, E.Z.; Li, L.; Melina, G.; Song, A.; Tavares, M.M. Bridging Skill Gaps for the Future: New Jobs Creation in the AI Age. <https://www.imf.org/en/publications/staff-discussion-notes/issues/2026/01/09/bridging-skill-gaps-for-the-future-new-jobs-creation-in-the-ai-age-572136>, 2026. IMF Staff Discussion Note SDN/2026/001, January 2026; accessed April 18, 2026.
19. Loaiza, I.; Rigobon, R. The EPOCH of AI: Human-Machine Complementarities at Work. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5028371, 2024. MIT Sloan Research Paper No. 7236-24; available at SSRN; posted November 21, 2024; accessed April 18, 2026, <https://doi.org/10.2139/ssrn.5028371>.
20. Eloundou, T.; Manning, S.; Mishkin, P.; Rock, D. GPTs are GPTs: Labor Market Impact Potential of LLMs. *Science* **2024**, *384*, 1306–1308. <https://doi.org/10.1126/science.adj0998>.
21. Challenger, Gray & Christmas. Challenger Report: January Job Cuts Surge; Lowest January Hiring on Record. <https://www.challengergray.com/blog/challenger-report-january-job-cuts-surge-lowest-january-hiring-on-record/>, 2026. Published February 6, 2026; accessed March 22, 2026.
22. Challenger, Gray & Christmas. Challenger Report: February Cuts Plunge, YTD Hiring Falls 56%. <https://www.challengergray.com/blog/challenger-report-february-cuts-plunge-hiring-falls-56-percent/>, 2026. Published March 5, 2026; accessed April 18, 2026.
23. Reuters. China's Baidu Starts Layoffs After Reporting Third-Quarter Loss - Sources. <https://www.reuters.com/business/world-at-work/chinas-baidu-starts-layoffs-after-reporting-third-quarter-loss-sources-2025-11-28/>, 2025. Published November 28, 2025; accessed March 22, 2026.
24. Reuters. ByteDance's TikTok Cuts Hundreds of Jobs in Shift Towards AI Content Moderation. <https://www.reuters.com/technology/bytedance-cuts-over-700-jobs-malaysia-shift-towards-ai-moderation-sources-say-2024-10-11/>, 2024. Published October 11, 2024; accessed March 22, 2026.
25. Reuters. Alibaba Says to Restart Hiring, Sees Signs of Start of AI Bubble in the US. <https://www.reuters.com/technology/alibaba-chairman-says-china-business-more-confident-since-xis-tech-summit-2025-03-25/>, 2025. Published March 25, 2025; accessed March 22, 2026.
26. Reuters. Alibaba Launches AI Platform for Enterprises as Agent Craze Sweeps China. <https://www.reuters.com/world/asia-pacific/alibaba-launches-new-ai-agent-platform-enterprises-2026-03-17/>, 2026. Published March 17, 2026; accessed March 22, 2026.

27. Reuters. Tencent Pledges Higher AI Investment in 2026 After Chip Curbs Hit Capex Plans. <https://www.reuters.com/world/asia-pacific/tencent-books-13-rise-quarterly-revenue-gaming-ai-demand-2026-03-18/>, 2026. Published March 18, 2026; accessed March 22, 2026.
28. Reuters. Xiaomi to Invest at Least \$8.7 Billion in AI Over Next Three Years, CEO Says. <https://www.reuters.com/world/asia-pacific/xiaomi-invest-least-87-billion-ai-over-next-three-years-ceo-says-2026-03-19/>, 2026. Published March 19, 2026; accessed March 22, 2026.
29. Reuters. SAP to Restructure 8,000 Jobs in Push Towards AI, Shares Hit Record. <https://www.reuters.com/technology/sap-announces-company-wide-restructuring-updates-2025-outlook-2024-01-23/>, 2024. Published January 24, 2024; accessed April 18, 2026.
30. Salesforce. Salesforce Delivers Record Fourth Quarter Fiscal 2026 Results. <https://investor.salesforce.com/news/news-details/2026/Salesforce-Delivers-Record-Fourth-Quarter-Fiscal-2026-Results/default.aspx>, 2026. Published February 25, 2026; accessed April 18, 2026.
31. Reuters. Oracle Plans Thousands of Job Cuts as Data Center Costs Rise, Bloomberg News Reports. <https://www.reuters.com/business/oracle-plans-thousands-job-cuts-data-center-costs-rise-bloomberg-news-reports-2026-03-05/>, 2026. Published March 5, 2026; accessed April 18, 2026.
32. Falk, B.H.; Tsoukalas, G. The AI Layoff Trap. *arXiv preprint arXiv:2603.20617* 2026. <https://arxiv.org/abs/2603.20617>.
33. International Labour Organization. Generative AI and Jobs: A Refined Global Index of Occupational Exposure. ILO Working Paper 140, International Labour Organization, 2025. <https://doi.org/10.54394/HETP0387>.
34. World Economic Forum. The Future of Jobs Report 2025. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>, 2025. Published January 7, 2025; accessed March 21, 2026.
35. World Economic Forum. The Top Jobs and Labour Market Stories of 2025. <https://www.weforum.org/stories/2026/01/top-jobs-and-labour-market-stories-2025/>, 2026. Published January 15, 2026; accessed March 21, 2026.
36. Autor, D.H.; Levy, F.; Murnane, R.J. The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics* **2003**, *118*, 1279–1333.
37. Autor, D.H. Why are there still so many jobs? The history and future of workplace automation. *Journal of economic perspectives* **2015**, *29*, 3–30.
38. Acemoglu, D.; Restrepo, P. Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives* **2019**, *33*, 3–30.
39. Noy, S.; Zhang, W. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **2023**, *381*, 187–192.
40. Peng, S.; Kalliamvakou, E.; Cihon, P.; Demirer, M. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. *arXiv preprint arXiv:2302.06590* 2023. Microsoft Research article with link to arXiv publication: <https://www.microsoft.com/en-us/research/publication/the-impact-of-ai-on-developer-productivity-evidence-from-github-copilot/>.
41. Cruces, G.; Fernández Mejjide, D.; Galiani, S.; Gálvez, R.H.; Lombardi, M. Does Generative AI Narrow Education-Based Productivity Gaps? Evidence from a Randomized Experiment. Working Paper 34851, National Bureau of Economic Research, 2026. Issued February 2026; <https://www.nber.org/papers/w34851>, <https://doi.org/10.3386/w34851>.
42. Dell’Acqua, F.; McFowland III, E.; Mollick, E.R.; Lifshitz, H.; Kellogg, K.C.; Rajendran, S.; Kraymer, L.; Candelon, F.; Lakhani, K.R. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *Organization Science* **2026**, *37*, 403–423. <https://pubsonline.informs.org/doi/10.1287/orsc.2025.21838>, <https://doi.org/10.1287/orsc.2025.21838>.
43. Handa, K.; Tamkin, A.; McCain, M.; Huang, S.; Durmus, E.; Heck, S.; Mueller, J.; Hong, J.; Ritchie, S.; Belonax, T.; et al. Which Economic Tasks Are Performed with AI? Evidence from Millions of Claude Conversations. *arXiv preprint arXiv:2503.04761* 2025. <https://arxiv.org/abs/2503.04761>.
44. Department for Science, Innovation and Technology, UK. AI Labour Market Survey 2025 Report: Executive Summary. <https://www.gov.uk/government/publications/ai-labour-market-survey-2025-report/ai-labour-market-survey-2025-report-executive-summary>, 2026. Published January 28, 2026; accessed March 22, 2026.
45. Global Reporting Initiative. GRI Topic Standard Project for Labor: Training and Education. <https://www.globalreporting.org/standards/standards-development/topic-standards-project-for-labor/>, 2025. Topic Standards Project for Labor page listing the Training and Education exposure draft; accessed April 18, 2026.

46. World Economic Forum. How AI Is Reshaping the Career Ladder, and Other Trends in Jobs and Skills on Labour Day. <https://www.weforum.org/stories/2025/04/ai-jobs-international-workers-day/>, 2025. Published April 30, 2025; accessed March 21, 2026.
47. World Economic Forum. How AI Is Changing Early Careers: A View from Entry. <https://initiatives.weforum.org/reskilling-revolution/learning-to-earning>, 2026. World Economic Forum initiative page referencing the briefing; accessed April 18, 2026.
48. World Bank. South Asia Development Update – Chapter 2 Highlights: Artificial Intelligence, Real Impact: Labor Market Implications of AI Adoption in South Asia. <https://openknowledge.worldbank.org/entities/publication/94725cf6-f681-42c9-acc7-63a076e0fc2b>, 2025. South Asia Development Update publication page; accessed April 18, 2026.
49. EFRAG. ESRS S1: Own Workforce. <https://www.efrag.org/en/sustainability-reporting/esrs-workstreams/sector-agnostic-standards-set-1-esrs>, 2023. ESRS Set 1 landing page; accessed April 18, 2026.
50. International Organization for Standardization. ISO 30414:2018 – Human Resource Management — Guidelines for Internal and External Human Capital Reporting. <https://www.iso.org/standard/69338.html>, 2018. Standard page; accessed March 22, 2026.
51. Organisation for Economic Co-operation and Development. OECD Guidelines for Multinational Enterprises on Responsible Business Conduct. https://www.oecd.org/en/publications/oecd-guidelines-for-multinational-enterprises-on-responsible-business-conduct_81f92357-en.html, 2023. Publication page; accessed April 18, 2026.
52. International Labour Organization. Just Transition Towards Environmentally Sustainable Economies and Societies. <https://www.ilo.org/topics-and-sectors/just-transition-towards-environmentally-sustainable-economies-and-societies>, 2025. Page accessed March 22, 2026.
53. United Nations Department of Economic and Social Affairs. Goal 8: Promote sustained, inclusive and sustainable economic growth, full and productive employment and decent work for all. <https://sdgs.un.org/goals/goal8>, 2025. Page accessed March 22, 2026.
54. PwC. The Fearless Future: PwC's 2025 Global AI Jobs Barometer. <https://www.pwc.com/gx/en/services/ai/ai-jobs-barometer.html>, 2025. Insight page; published June 3, 2025; accessed April 18, 2026.
55. He, C.; Zhou, X.; Wang, D.; Xu, H.; Liu, W.; Miao, C. Remote-Capable Knowledge Work Should Default to AI-Enabled Flexibility, 2026. Preprint.
56. Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I.D.; Gebru, T. Model Cards for Model Reporting. In Proceedings of the Proceedings of the Conference on Fairness, Accountability, and Transparency. Association for Computing Machinery, 2019, pp. 220–229. <https://doi.org/10.1145/3287560.3287596>.
57. Gebru, T.; Morgenstern, J.; Vecchione, B.; Wortman Vaughan, J.; Wallach, H.; Daumé III, H.; Crawford, K. Datasheets for Datasets. *Communications of the ACM* **2021**, *64*, 86–92. <https://doi.org/10.1145/3458723>.
58. Raji, I.D.; Smart, A.; White, R.N.; Mitchell, M.; Gebru, T.; Hutchinson, B.; Smith-Loud, J.; Theron, D.; Barnes, P. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. In Proceedings of the Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. Association for Computing Machinery, 2020, pp. 33–44. <https://doi.org/10.1145/3351095.3372873>.
59. Reuters. Sweden's Klarna Shifts AI Focus from Cost Cuts to Growth. <https://www.reuters.com/business/swedens-klarna-shifts-ai-focus-cost-cuts-growth-2025-09-10/>, 2025. Published September 10, 2025; accessed March 21, 2026.
60. Autor, D. Applying AI to Rebuild Middle Class Jobs. Technical report, National Bureau of Economic Research, 2024.
61. He, C.; Zhou, X.; Wang, D.; Xu, H.; Liu, W.; Miao, C. Human-AI Productivity Claims Should Be Reported as Time-to-Acceptance under Explicit Acceptance Tests, 2026. TechRxiv preprint.
62. He, C.; Zhou, X.; Wang, D.; Xu, H.; Liu, W.; Miao, C. Harness Engineering for Language Agents: The Harness Layer as Control, Agency, and Runtime, 2026. Preprint.
63. He, C.; Zhou, X.; Wang, D.; Xu, H.; Liu, W.; Miao, C. The AutoResearch Moment: From Experimenter to Research Director, 2026. Preprint.
64. He, C.; Zhou, X.; Wang, D.; Xu, H.; Liu, W.; Miao, C. OpenClaw as Language Infrastructure: A Case-Centered Survey of a Public Agent Ecosystem in the Wild, 2026. Preprint.
65. He, C.; Zhou, X.; Wang, D.; Xu, H.; Liu, W.; Miao, C. Let Papers Flow: AI Conferences Should Embrace Submission Explosion via Autonomous Review Pipelines, 2026. Preprint.

66. He, C.; Zhou, X.; Wu, Y.; Yu, X.; Zhang, Y.; Zhang, L.; Wang, D.; Lyu, S.; Xu, H.; Wang, X.; et al. ESGenius: Benchmarking LLMs on Environmental, Social, and Governance (ESG) and Sustainability Knowledge. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, 2025, pp. 14623–14664.
67. Zhang, L.; Zhou, X.; He, C.; Wang, D.; Wu, Y.; Xu, H.; Liu, W.; Miao, C. MMESGBench: Pioneering Multimodal Understanding and Complex Reasoning Benchmark for ESG Tasks. In Proceedings of the Proceedings of the 33rd ACM International Conference on Multimedia, 2025, pp. 12829–12836.
68. He, C.; Zhou, X.; Yu, X.; Zhang, L.; Zhang, Y.; Wu, Y.; Xiao, L.; Li, L.; Wang, D.; Xu, H.; et al. SSKG Hub: An Expert-Guided Platform for LLM-Empowered Sustainability Standards Knowledge Graphs. *arXiv preprint arXiv:2603.00669* 2026.
69. He, C.; Zhou, X.; Wang, D.; Yu, X.; Xiao, L.; Li, L.; Xu, H.; Liu, W.; Miao, C. KG4ESG: The ESG Knowledge Graph Atlas, 2026. Preprint.
70. He, C.; Zhou, X.; Wang, D.; Xu, H.; Liu, W.; Miao, C. ESGlass: Glass-Box ESG and Sustainability Reports, 2026. Preprint.
71. Bresnahan, T.F.; Brynjolfsson, E.; Hitt, L.M. Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The quarterly journal of economics* 2002, 117, 339–376.
72. Deming, D.J. The growing importance of social skills in the labor market. *The quarterly journal of economics* 2017, 132, 1593–1640.
73. Skills England. AI Skills for the UK Workforce. <https://www.gov.uk/government/publications/ai-skills-for-the-uk-workforce>, 2025. Published October 29, 2025; accessed March 22, 2026.
74. Organisation for Economic Co-operation and Development. AI and Work. <https://www.oecd.org/en/topics/sub-issues/ai-and-work.html>, 2026. Topic page; accessed March 21, 2026.
75. Reuters. Atlassian to Cut Roughly 10% Jobs in Pivot to AI. <https://www.reuters.com/technology/atlassian-lay-off-about-1600-people-pivot-ai-2026-03-11/>, 2026. Published March 11, 2026; accessed April 18, 2026.
76. Reuters. Block Shares Soar as Dorsey Leans on AI to Trim Workforce. <https://www.reuters.com/sustainability/sustainable-finance-reporting/block-shares-soar-dorsey-leans-ai-trim-workforce-2026-02-27/>, 2026. Published February 27, 2026; accessed April 18, 2026.
77. Reuters. Workday to Cut 1,750 Jobs in AI Push. <https://www.reuters.com/business/workday-cut-85-its-workforce-2025-02-05/>, 2025. Published February 5, 2025; accessed March 21, 2026.
78. Reuters. HP to Cut About 6,000 Jobs by 2028, Ramps Up AI Efforts. <https://www.reuters.com/business/hp-cut-about-6000-jobs-by-2028-ramps-up-ai-efforts-2025-11-25/>, 2025. Published November 25, 2025; accessed April 18, 2026.
79. Reuters. Salesforce Cuts Fewer Than 1,000 Jobs, Business Insider Reports. <https://www.reuters.com/business/world-at-work/salesforce-cuts-less-than-1000-jobs-business-insider-reports-2026-02-10/>, 2026. Published February 10, 2026; accessed March 21, 2026.
80. Reuters. Salesforce Forecasts Weak Current-Quarter Revenue, Shares Fall. <https://www.reuters.com/sustainability/sustainable-finance-reporting/salesforce-forecasts-weak-current-quarter-revenue-shares-fall-2025-09-03/>, 2025. Published September 3, 2025; accessed March 22, 2026.
81. Reuters. Meta Shares Jump After Reuters Report on Plans for Layoffs of 20% or More. <https://www.reuters.com/business/meta-shares-jump-after-reuters-report-plans-layoffs-20-or-more-2026-03-16/>, 2026. Published March 16, 2026; accessed March 21, 2026.
82. Reuters. Microsoft to Cut About 4% of Jobs Amid Hefty AI Bets. <https://www.reuters.com/business/world-at-work/microsoft-lay-off-many-9000-employees-seattle-times-reports-2025-07-02/>, 2025. Published July 2, 2025; accessed March 21, 2026.
83. Reuters. Microsoft Racks Up over \$500 Million in AI Savings While Slashing Jobs, Bloomberg News Reports. <https://www.reuters.com/business/microsoft-racks-up-over-500-million-ai-savings-while-slashing-jobs-bloomberg-2025-07-09/>, 2025. Published July 9, 2025; accessed March 22, 2026.
84. Reuters. Amazon to Cut About 14,000 Corporate Jobs in AI Push. <https://www.reuters.com/sustainability/amazon-lay-off-about-14000-roles-2025-10-28/>, 2025. Published October 28, 2025; accessed March 21, 2026.
85. Amazon. Update on Our Organization. <https://www.aboutamazon.com/news/company-news/amazon-layoffs-corporate-jan-2026>, 2026. Statement by Beth Galetti, January 28, 2026; accessed March 22, 2026.
86. Reuters. Sweden's Klarna Says AI Chatbots Help Shrink Headcount. <https://www.reuters.com/technology/artificial-intelligence/swedens-klarna-says-ai-chatbots-help-shrink-headcount-2024-08-27/>, 2024. Published August 27, 2024; accessed March 21, 2026.

87. Reuters. Salesforce Closes 1,000 Paid ‘Agentforce’ Deals, Looks to Robot Future. <https://www.reuters.com/technology/artificial-intelligence/salesforce-closes-1000-paid-agentforce-deals-looks-robot-future-2024-12-17/>, 2024. Published December 17, 2024; accessed March 22, 2026.
88. U.S. Bureau of Labor Statistics. Productivity and Costs, Fourth Quarter and Annual Averages 2025. <https://www.bls.gov/bls/news-release/prod.htm>, 2026. Archived BLS productivity news releases page; see the 2026 fourth-quarter 2025 release (HTML); accessed April 18, 2026.
89. U.S. Bureau of Labor Statistics. Job Openings and Labor Turnover – January 2026. <https://www.bls.gov/bls/news-release/jolts.htm>, 2026. Archived BLS JOLTS news releases page; see the 2026 January 2026 release (HTML); accessed April 18, 2026.
90. Ministry of Manpower, Singapore. Job Vacancies 2025: Labour Demand Gradually Shifting to Growth Areas as Firms Adjust Hiring Plans. <https://www.mom.gov.sg/newsroom/press-releases/2026/0320-job-vacancies-report-2025>, 2026. Press release, March 20, 2026; accessed March 22, 2026.
91. Reuters. China Pins Hopes on Society-Wide AI Push to Add Jobs, Rejuvenate Economy. <https://www.reuters.com/business/world-at-work/china-pins-hopes-society-wide-ai-push-add-jobs-rejuvenate-economy-2026-03-10/>, 2026. Published March 10, 2026; accessed March 22, 2026.
92. State Council of the People’s Republic of China. Chinese Authorities Roll Out Measures to Boost Employment for Graduate Job Seekers. https://english.www.gov.cn/news/202603/20/content_WS69bd46c3c6d00ca5f9a0a070.html, 2026. Published March 20, 2026; accessed March 22, 2026.
93. Reuters. China’s February Youth Jobless Rate Dips to 16.1%. <https://www.reuters.com/world/asia-pacific/chinas-february-youth-jobless-rate-dips-161-2026-03-19/>, 2026. Published March 19, 2026; accessed March 22, 2026.
94. Infocomm Media Development Authority. Singapore to Build AI-Fluent Workforce to Accelerate National AI Ambition. <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/press-releases/2025/sg-to-build-ai-fluent-workforce-to-accelerate-national-ai-ambition>, 2025. Press release, August 29, 2025; accessed March 21, 2026.
95. Workforce Singapore. Jobs Transformation Map – Generative AI in Finance. <https://www.wsg.gov.sg/home/employers-industry-partners/jobs-transformation-maps/jobs-transformation-map-generative-ai>, 2026. Page updated January 2, 2026; accessed March 22, 2026.
96. Microsoft Source Asia. Ministry of Labour and Microsoft Join Forces to Accelerate AI Skill Development for 150,000 Thai Workers, Driving Thailand Towards Becoming “Creator” Nation in Digital Economy. <https://news.microsoft.com/source/asia/2025/11/25/ministry-of-labour-and-microsoft-join-forces-to-accelerate-ai-skill-development-for-150000-thai-workers-driving-thailand-towards-becoming-creator-nation-in-digital-economy/>, 2025. Published November 25, 2025; accessed March 21, 2026.
97. State Council of the People’s Republic of China. China Holds Central Economic Work Conference to Plan for 2026. https://english.www.gov.cn/news/202512/11/content_WS693a9c0dc6d00ca5f9a08098.html, 2025. Published December 11, 2025; accessed March 22, 2026.
98. European Commission. AI Talent, Skills and Literacy. <https://digital-strategy.ec.europa.eu/en/policies/ai-talent-skills-and-literacy>, 2025. Page updated January 16, 2026; accessed March 22, 2026.
99. United Nations Global Compact. Social Sustainability. <https://unglobalcompact.org/what-is-gc/our-work/social>, 2025. Page accessed March 22, 2026.
100. Tencent Holdings. 2024 Environmental, Social and Governance Report. <https://www.tencent.com/en-us/esg/esg-reports.html>, 2025. ESG reports page; accessed April 18, 2026.
101. Baidu. 2024 Environmental, Social and Governance Report. https://esg.baidu.com/en_reports.html, 2025. ESG reports page; accessed April 18, 2026.
102. JD.com. 2024 Environmental, Social and Governance Report. <https://ir.jd.com/esgcsr>, 2025. ESG / Sustainability page; accessed April 18, 2026.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.