WHITE PAPER WORKFORCE ECONOMICS FOR ENTERPRISES

Reskill, Redeploy, or Replace: A Cost Model for the AI Workforce Transition

A defensible four-variable cost framework for the AI workforce transition, populated with named public benchmarks and made auditable by a reproducible, per-occupation exposure map.

JobRoute Research · June 3, 2026 · 18 min read

The reskill, redeploy, or replace choice is an economic decision, yet most enterprises still make it on instinct or inside a vendor model the board cannot inspect. Replacing a worker whose role is exposed to AI frequently costs more than reskilling and redeploying that same worker once recruiting, ramp-to-productivity, lost institutional knowledge, and severance are all counted. The useful part is that this is no longer a matter of opinion. Named public benchmarks from Gallup, the World Economic Forum, the Boston Consulting Group, SHRM, McKinsey, and Lightcast make the comparison quantifiable rather than asserted, which is what lets a CFO or audit committee defend the spend line by line.

This paper sets out a defensible four-variable cost framework, populates it with those public benchmarks, and works one honestly labeled illustrative example so a reader can rebuild the calculation with their own organization's numbers. JobRoute's role is narrow and specific: it is the measurement engine that turns a reproducible, source-traceable, per-occupation exposure map into the prioritized redeployment plan that feeds this model. It does not publish a competing headline statistic.

KEY TAKEAWAYS

- Replacement is rarely the cheap option. Gallup puts the cost of replacing one employee at one-half to two times annual salary and calls that conservative (Gallup, 2019), and SHRM's 2025 benchmark adds an average cost-per-hire of $5,475 for nonexecutive roles and $35,879 for executives, before you count ramp time and lost institutional knowledge (SHRM, 2025 Benchmarking Reports).

- Reskilling and redeployment have a defensible business case at scale. The World Economic Forum and BCG pegged the average cost to reskill a displaced US worker at about $24,800 (Towards a Reskilling Revolution, 2019), and WEF's 2025 report projects 19 of every 100 workers will be reskilled and redeployed inside their organization by 2030, with about half of employers planning such transitions (WEF, Future of Jobs Report 2025).

- Most AI exposure is augmentation, not replacement. Anthropic's Economic Index, built on US Department of Labor O*NET tasks, finds roughly 36% of occupations use AI across a quarter or more of their tasks, with usage skewing toward augmentation rather than wholesale automation (about 57% augmentation and 43% automation) (Anthropic, February 2025), which tilts the math toward reskilling for many exposed roles.

- The skills you build carry measurable market value. Lightcast finds AI skills command a 28% salary premium, about $18,000 a year, across analysis of more than 1.3 billion job postings (Lightcast, 2025), so a retraining investment lands against a documented wage signal rather than a guess.

- Board defensibility depends on traceable inputs. A model built on the dated 2013 Frey and Osborne 47% figure cannot be defended; one built on reconciled, version-locked, per-occupation exposure can be audited

line by line, which is what converts a workforce decision into a fundable one.

Why is reskill, redeploy, or replace an economic decision and not an HR judgment call?

Start with clean definitions, because the three paths are often blurred. To **reskill** is to upskill a worker inside their current role so they can do it alongside the AI tools now embedded in it. To **redeploy** is to reskill that worker and move them to an adjacent, growing role inside the organization. To **replace** is to separate the worker and hire externally for the new capability. These are not moral positions. They are three different routes to the same destination, which is the capability the business needs next.

At the scale of a single role, choosing between them looks like an HR judgment call. At the scale enterprises now face, it is a capital-allocation decision with a measurable cost on each path. The World Economic Forum estimates that 39% of workers' existing core skills will be transformed or become outdated between 2025 and 2030, and that of every 100 workers, 59 will need training over that period (World Economic Forum, Future of Jobs Report 2025). That is not a handful of edge cases for a manager to weigh case by case. It is a structural reallocation that touches the majority of a workforce, and decisions at that scale get funded, audited, and reported.

The reason this reaches the board is the value at stake on the other side of the ledger. McKinsey estimates that generative AI could add the equivalent of $2.6 trillion to $4.4 trillion annually across 63 use cases spanning 16 business functions, and that current generative AI could automate work activities that today absorb 60% to 70% of employees' time (McKinsey & Company, 2023). The right way to read that 60% to 70% figure is as the size of the reallocation problem, not as a layoff forecast. Time freed from routine tasks is time that has to be redirected somewhere, and the choice of where is exactly the reskill, redeploy, or replace question.

WHY THIS MATTERS

The question is not whether a role is exposed. Almost every role is exposed to something. The question is which path returns the capability the business needs at the lowest fully-loaded cost and risk. That reframe is the whole point: it moves the decision from sentiment to arithmetic. For the strategic framing behind this model, see the strategic overview for CHROs.

Answering that per role requires a model with explicit variables and traceable inputs. The rest of this paper builds one.

What are the four variables in a defensible cost model?

A defensible model rests on four variables, and each can be tied to a named public benchmark.

Replacement cost is the fully-loaded cost of separating the exposed worker and bringing in someone new. It includes recruiting and cost-per-hire, severance for the exiting worker, ramp-to-productivity for the new hire, and the value of lost institutional knowledge. Gallup frames the total at one-half to two times annual salary and explicitly calls that a conservative estimate (Gallup, 2019).

Retraining cost is the cost of building the new capability in a worker you keep. It includes direct training spend, wages paid during training, the opportunity cost of time off task, and an allowance for program failure rate. The canonical public benchmark is about $24,800 to reskill a displaced US worker (World Economic Forum and BCG, Towards a Reskilling Revolution, 2019).

Productivity ramp is the time-to-full-productivity curve that any worker, external hire or redeployed internal worker, has to climb. It is the variable most often underweighted. A redeployed worker frequently starts higher on the curve,

because firm-specific knowledge, system access, and relationships transfer with them rather than being rebuilt from zero.

Exposure-adjusted risk is the probability that the role's tasks are genuinely displaced, weighted by whether the exposure is augmentation or automation. This is where task-level evidence matters. Anthropic's Economic Index, which maps Claude usage to O*NET occupational tasks, finds roughly 36% of occupations show AI use across at least a quarter of their tasks, with usage skewing toward augmentation rather than wholesale replacement (Anthropic, February 2025). Augmentation-weighted exposure lowers the displacement probability and tilts the model toward keeping the worker.

VARIABLE	WHAT IT CAPTURES	PUBLIC BENCHMARK	PUBLISHER, YEAR
Replacement cost	Severance, recruiting, ramp, lost knowledge	One-half to two times annual salary (called conservative)	Gallup, 2019
Retraining cost	Tuition, wages during training, time off task, failure rate	About $24,800 average per displaced US worker	WEF and BCG, 2019
Productivity ramp	Time-to-full-productivity for new hire vs redeployed worker	Recruiting line benchmarked at $5,475 nonexec / $35,879 exec cost-per-hire	SHRM, 2025
Exposure-adjusted risk	Displacement probability, weighted by augmentation vs automation	About 36% of occupations use AI for 25%+ of tasks; usage skews to augmentation	Anthropic, 2025

The four cost-model variables, each tied to a named public benchmark

Source: Multiple; see source column

The discipline that makes this auditable is simple: every input traces to a named, dated public source. A finance committee does not have to trust a vendor's black

box. It can follow each number back to its publisher.

What does it actually cost to replace an exposed worker?

Walk the replacement path in order and the components stack up. First comes **separation**: severance and the administrative cost of an exit. Then **vacancy and lost output** while the seat is empty. Then **recruiting and hiring** to fill it. Then **onboarding and ramp** as the new hire climbs the productivity curve. And finally the hardest to recover, **lost institutional knowledge**, which walks out the door with the departing worker.

The headline benchmark ties it together. Gallup estimates that replacing an individual employee costs one-half to two times that employee's annual salary, and describes that range as a conservative estimate (Gallup, 2019). At the macro level, Gallup estimates that US businesses lose approximately $1 trillion every year to voluntary turnover, which is the scale of avoidable replacement cost when retention and redeployment are not on the table.

The recruiting line alone is well measured. SHRM's 2025 Benchmarking Report puts average cost-per-hire at $5,475 for nonexecutive roles and $35,879 for executive roles, from a survey fielded January to March 2025 (SHRM, 2025 Benchmarking Reports). The important framing is that this cost-per-hire figure is one line item inside the broader Gallup multiple, not the whole cost. A model that counts only cost-per-hire dramatically understates the true cost of replacement.

50% to 200%

of annual salary to replace one employee, called a conservative estimate

$5,475

average cost-per-hire for nonexecutive roles (one line item, not the whole cost)

SHRM, 2025 BENCHMARKING REPORTS

~$1 trillion

lost annually by US businesses to voluntary turnover

GALLUP, 2019

GALLUP, 2019

Lost institutional knowledge deserves its own emphasis, because it is the component most often left out of instinct-based decisions and the one that most favors keeping the worker. Years of context on customers, systems, exceptions, and informal process do not appear on an invoice, which is precisely why they get omitted. There is no single public dollar benchmark that captures it cleanly, so it is best treated qualitatively in the model and estimated from an organization's own knowledge-transfer experience. Leaving it at zero is not conservative. It is a thumb on the scale in favor of replacement.

What does it cost to reskill and redeploy that same worker?

Now itemize the other path. Retraining and redeployment cost includes program tuition or platform fees, wages paid during training, partial productivity during the transition period, and a residual allowance for program failure.

The canonical public reskilling benchmark is about $24,800 on average to reskill a displaced US worker, and the World Economic Forum and BCG modeled that 77% of at-risk US workers could be reskilled with a positive cost-benefit balance (WEF and BCG, Towards a Reskilling Revolution, 2019). Two honest caveats belong with that number. It is a displaced-worker average, and in-role reskilling is often cheaper because the worker keeps producing while they learn. And it is a 2019 figure that should be inflation-adjusted and re-sourced before it enters a live model.

The skills built are not only a cost line. They carry measurable market value. Lightcast finds that job postings requiring AI skills pay a 28% salary premium, roughly $18,000 more per year, based on analysis of more than 1.3 billion job postings (Lightcast, 2025). That premium is the wage signal a retraining investment lands against. A program that moves workers toward AI-

complementary skills is investing against documented market demand, not a hunch.

BY THE NUMBERS

Set the two anchors side by side and the comparison comes into focus. Replacing one worker runs one-half to two times annual salary, called conservative by Gallup (Gallup, 2019). Reskilling a displaced US worker averaged about $24,800 (WEF and BCG, 2019). For a worker earning, say, $70,000, the Gallup range alone spans $35,000 to $140,000, which brackets the reskilling benchmark on both sides and exceeds it badly at the upper end. The operational machinery to act on this comparison is covered in JobRoute for enterprises.

Redeployment has a specific structural advantage over an external hire. A redeployed worker retains firm-specific knowledge, security clearance to internal systems, and existing relationships, so the productivity ramp is shorter and the early months are more productive. The WEF demand signal confirms this is where the macro trend is heading: 19 of every 100 workers are projected to be reskilled and redeployed inside their organization by 2030, and about half of employers plan such transitions (WEF, Future of Jobs Report 2025). Redeployment is not a fringe tactic. It is a planned movement across a meaningful slice of the workforce.

How do the two paths compare for a representative exposed role?

This worked example is illustrative. It uses public benchmark inputs to show the structure of the calculation. It is not a JobRoute-specific finding, not a real client, and not a forecast. The purpose is to let a reader see the shape of the math and rebuild it with their own organization's numbers.

Take a hypothetical back-office customer-operations analyst at an illustrative salary of $70,000, a role chosen because McKinsey flags customer operations among the highest-value areas for generative AI (McKinsey, 2023). Every assumption below is stated on purpose.

- **Salary:** $70,000 (illustrative).
- **Gallup multiple chosen:** 1.0x within the published 0.5x to 2.0x range (Gallup, 2019), a deliberately mid-conservative pick.
- **Recruiting line:** $5,475, the SHRM 2025 nonexecutive cost-per-hire (SHRM, 2025).
- **Reskilling cost:** $24,800, the WEF and BCG displaced-worker average (WEF and BCG, 2019).
- **Ramp:** the external hire is assumed to ramp more slowly than the redeployed worker, who retains firm-specific knowledge.

COST COMPONENT	REPLACE PATH	RESKILL-AND-REDEPLOY PATH
Separation / severance	$11,700 (illustrative)	$0
Recruiting and hiring	$5,475 (SHRM, 2025)	$0
Training / reskilling	$0	$24,800 (WEF and BCG, 2019)
Ramp-to-productivity loss	~$45,000 (illustrative, slower external ramp)	~$18,000 (illustrative, faster ramp from retained knowledge)
Lost institutional knowledge	Material, treated qualitatively	$0 (retained)
Illustrative total (ex. knowledge loss)	**~$62,000+**	**~$42,800**

Illustrative side-by-side comparison for a hypothetical $70,000 customer-operations analyst whose exposure is augmentation-weighted

Source: Illustrative model; inputs from Gallup 2019, SHRM 2025, WEF and BCG 2019

Under these clearly labeled public-benchmark assumptions, the reskill-and-redeploy path comes in below the replace path for a role whose exposure is augmentation-weighted, and the gap widens further once the qualitative cost of lost institutional knowledge is added on the replace side. The honest caveat matters as much as the result: the comparison can flip. For a role whose tasks are genuinely automated rather than augmented, the ramp advantage shrinks and replacement can win. The single variable that decides which way it tips is exposure type, augmentation versus automation, which is exactly why the model has to measure it per occupation rather than assume it.

> ***Exposure is the start of a plan, not the end of a career.***

Why does reproducibility and source traceability matter for board defensibility?

A finance or audit committee cannot approve material workforce spend on inputs it cannot inspect. That is the governance problem at the heart of this decision. Most proprietary talent-intelligence platforms hand back a score, not a derivation. A board asked to fund a multi-million-dollar reskilling program on the strength of a number it cannot trace back to a source has been asked to approve a black box, and a diligent committee will decline.

The dated consumer baseline shows what indefensible looks like. The widely cited claim that 47% of US employment is at high risk of computerization comes from Frey and Osborne's 2013 Oxford working paper, which classified 702 occupations at the occupation level (Frey and Osborne, Oxford Martin School, 2013). That was important early work, but it is an occupation-level probability that task-level measurement has moved well beyond, because most occupations are bundles of tasks where some are exposed and others are not. A workforce model built on that single 2013 figure cannot be defended in 2026.

Reproducible and version-locked means something concrete in practice. Every input traces to a named, dated public source. The version of each source is fixed, so the analysis can be re-run and reconciled rather than silently drifting. And a third party can reproduce the exposure ranking from the same inputs. This is the JobRoute approach: reconciling O*NET 30.2, ESCO v1.2.1, Lightcast Open Skills, the Anthropic Economic Index, the WEF Future of Jobs Report 2025, and BLS data into one per-occupation measurement across 1,016 O*NET occupations.

The differentiator is worth stating plainly and honestly. JobRoute is the measurement engine, not a competing headline statistic. It does not publish a claim that "JobRoute found X% of jobs are at risk." It operationalizes the public numbers per occupation, reconciling sources that disagree and exposing the workings. That restraint is not modesty. It is exactly what keeps the model citable and audit-ready, because the inputs remain the public record rather than a proprietary assertion.

WHY THIS MATTERS

The difference between a workforce decision and a fundable one is whether the inputs can be audited line by line. A score you cannot trace is a position. A reconciled, version-locked, per-occupation exposure map is evidence. The version-locking detail sits in the JobRoute methodology.

How does a CHRO turn an exposure map into a prioritized redeployment plan?

The model becomes operational as a sequence of steps a CHRO can run:

1. **Map exposure at the task level** per occupation, rather than labeling whole jobs as safe or doomed.
2. **Classify each exposed role** as augmentation-weighted or automation-weighted, since that classification decides which way the cost comparison tips.
3. **Score each role on the four cost variables**: replacement cost, retraining cost, productivity ramp, and exposure-adjusted risk.
4. **Rank roles by the gap** between replace cost and reskill-and-redeploy cost, so the largest savings and the clearest wins surface first.
5. **Identify the specific adjacent role** each exposed worker can move toward, named concretely rather than left as a vague promise.

That fifth step is the operational payoff. Naming durable skills and mapping adjacent roles is what turns an exposure score into an actionable target. A model that can point a worker toward a specific adjacent role is delivering a plan; a model that only flags risk is delivering a warning. The difference shows up directly in employee trust and in whether a redeployment program actually retains people.

The plan should match the macro trend rather than fight it. The World Economic Forum reports that 85% of surveyed employers plan to upskill their workforce and about half plan to transition staff into growing roles by 2030 (WEF, Future of Jobs Report 2025). A CHRO holding a per-occupation map can act on the 19-of-100 redeployment opportunity deliberately, role by role, rather than reacting after the exposure has already become attrition.

One continuity is worth noting. The same reconciled exposure layer that drives this enterprise cost model also powers the free individual AI Ready Score and the public-sector case for WIOA-fundable program measurement. The methodology does not change across audiences; only the view onto it does. A worker can see how the AI Ready Score works or take the free AI Ready Score, and it rests on the same per-occupation engine the CHRO is using.

> NOTE
>
> The prioritized plan is a living artifact, not a one-time deliverable. The exposure inputs are version-locked for auditability, but they should be re-run as sources update, so the ranking reflects the current O*NET, ESCO, Lightcast, Anthropic, WEF, and BLS releases rather than a frozen snapshot.

Methodology, assumptions, and limitations

The framework combines four variables, each populated from named, dated, publicly verifiable benchmarks: Gallup (2019) for replacement cost, SHRM (2025) for cost-per-hire, the World Economic Forum and BCG (2019) for reskilling cost, McKinsey (2023) for the value at stake, Lightcast (2025) for the AI-skill wage premium, and the Anthropic Economic Index (February 2025) for task-level exposure. The worked example is illustrative. Two horizons matter for accuracy: WEF Future of Jobs figures are to 2030, not 2028, and the Frey and Osborne (2013) figure appears only as a contrast point, never as a current estimate.

Four limitations should travel with this model whenever it is used.

Public benchmarks are averages and ranges, not your organization's numbers. The Gallup 0.5x to 2.0x replacement range is wide by design, and an enterprise must calibrate it with its own cost-per-hire data, severance policy, and observed ramp curves before treating any output as a real budget figure.

The reskilling benchmark is dated and is a displaced-worker average. The roughly $24,800 figure (WEF and BCG, 2019) reflects displaced-worker reskilling in 2019. In-role reskilling and platform-based training may cost considerably less, and the figure should be inflation-adjusted and re-sourced before use.

Exposure type drives the result. Anthropic's evidence that exposure skews toward augmentation (Anthropic, February 2025) supports reskilling for many roles, but automation-weighted roles can favor replacement. The model must not be applied as a blanket conclusion in either direction; the augmentation-versus-automation classification is doing real work in every comparison.

Institutional-knowledge loss resists a single public dollar figure. It is treated qualitatively here, and enterprises should estimate it from their own attrition and knowledge-transfer data rather than borrowing an external number that does not fit their context.

A CAVEAT WORTH NAMING

The worked example in this paper is illustrative and built entirely from named public benchmarks. No JobRoute-specific findings are claimed anywhere in it, and no number in the side-by-side table should be read as a forecast for any real organization or role. For the version-locking and source-reconciliation detail that makes the exposure layer reproducible, see the JobRoute methodology.

Sources and further reading

1. This Fixable Problem Costs U.S. Businesses $1 Trillion (cost to replace one employee is one-half to two times annual salary, called conservative; about $1 trillion annual US voluntary-turnover loss) Gallup, 2019
2. SHRM Releases 2025 Benchmarking Reports (average cost-per-hire $5,475 nonexecutive and $35,879 executive; survey fielded January to March 2025) SHRM (Society for Human Resource Management), 2025
3. Towards a Reskilling Revolution: Industry-Led Action for the Future of Work (about $24,800 average reskilling cost per displaced US worker; 77% reskillable with a positive cost-benefit balance) World Economic Forum and Boston Consulting Group, 2019
4. Future of Jobs Report 2025 (39% of core skills transformed or outdated by 2030; 59 of 100 workers need training, 29/19/11 breakdown; 85% plan to upskill; about half plan to transition staff into growing roles) World Economic Forum, 2025
5. Introducing the Anthropic Economic Index (about 36% of occupations use AI across a quarter or more of tasks; about 57% augmentation and 43% automation; built on US Department of Labor O*NET tasks) Anthropic, 2025
6. New Lightcast Report: AI Skills Command 28% Salary Premium (28% premium, roughly $18,000 more per year, across analysis of more than 1.3 billion job postings) Lightcast, 2025
7. The economic potential of generative AI: The next productivity frontier ($2.6 trillion to $4.4 trillion across 63 use cases and 16 business functions; 60% to 70% of employee time automatable; customer operations among the four highest-value areas) McKinsey & Company, 2023
8. The Future of Employment: How Susceptible Are Jobs to Computerisation? (the dated 2013 estimate that 47% of US employment is at high risk, classifying 702 occupations at the occupation level; used only as a contrast point) Carl Benedikt Frey and Michael A. Osborne, Oxford Martin School, 2013

Frequently asked questions

Is it cheaper to reskill or replace an employee whose role is exposed to AI? +

What is the average cost to replace an employee? +

How much does it cost to reskill a worker for the AI economy? +

How many workers will need reskilling or redeployment by 2030? +

Is most AI exposure automation or augmentation? +

Why is the Frey and Osborne 47% figure the wrong basis for an AI workforce decision? +

Does reskilling produce measurable financial value, or only avoid cost? +

How can a CHRO make the reskill-versus-replace case auditable for the board? +