INDEPENDENT WORKING PAPER

Human Capacity Under Coupled Constraints

An Evidence-Based Compact for Learning, Work, Public Infrastructure, and AI Governance in an AI-Intensive Economy

Darrell Morris

Version 0.6 | September 8, 2026

> *This working paper advances a policy framework. It is not a fiscal score, macroeconomic forecast, or legislative text. Publication revision: September 8, 2026.*

Abstract

Public debate about artificial intelligence often collapses four different questions into one: whether a task is technically exposed, whether an organization adopts an AI system, whether the system substitutes for human work, and whether paid employment is ultimately eliminated. This paper argues that those distinctions are not merely semantic. They determine whether society can measure an automation gain, identify a public burden, or design a fair response. Drawing on primary research, official statistics, technical reports, and governance instruments, the paper develops an AI Transition Balance Sheet: linked, non-additive ledgers of private gains, household and fiscal incidence, physical resource costs, and rights. Distributional transfers are not treated as resource losses or mechanically subtracted to produce a social-surplus total.

The principal empirical finding is a disciplined non-result: the public data assembled here do not identify a ten-year forecast of AI labor hours by sector. A valid AI labor hour must be a human-equivalent task-hour: the gross counterfactual paid human time required to produce accepted output at an agreed level of quality, risk, and timeliness. Actual human setup, oversight, validation, rework, and exception handling enter the subsequent net-substitution calculation once. Existing measures of paid hours, occupational exposure, firm adoption, task-use incidence, substitution, labor share, and margins use incompatible denominators. They should not be multiplied into a headline forecast.

The paper nevertheless finds enough evidence to justify a conditional social compact. If audited AI deployments reduce covered payroll, concentrate economic surplus, or impose uncompensated electricity, water, grid, and subsidy burdens, a bounded share of the realized gain should support portable human capacity. The proposed compact includes a modular Universal Learning Guarantee, separately financed health-coverage continuity, completion and placement accountability, site-specific data-center impact accounting, and layered AI governance combining law, independent oversight, and auditable voluntary practice. The proposal is designed to scale with measured conditions, not with technological fear.

Keywords

Artificial intelligence; automation; human-equivalent task-hours; education finance; apprenticeship; health continuity; data centers; AI governance; public reciprocity; economic inequality.

> **CORE PROPOSITION: When verified AI-related gains are accompanied by displaced compensation or uncompensated public burdens, a bounded portion of the realized surplus should finance portable human capacity and the costs created by the deployment. If those conditions do not occur, the contribution should remain correspondingly small.**

Research status and scope

This is a policy and research working paper, not a fiscal score, macroeconomic forecast, or legislative text. Version 0.6 is a manuscript-and-matrix reconciliation dated September 8, 2026; the underlying statistical vintages remain those identified beside each measure. The United States is the primary jurisdiction; global labor evidence and European law are used as comparators. Broader proposals involving universal food, utilities, or income support remain outside the costed scope because the present evidence matrix does not identify their program designs or incremental costs. They should be evaluated as separate modules rather than attached to an education estimate.

Publication and version information

Suggested citation: Morris, D. (2026). Human capacity under coupled constraints: An evidence-based compact for learning, work, public infrastructure, and AI governance in an AI-intensive economy (Version 0.6) [Independent working paper].

Companion evidence: Morris, D. (2026). AI Transition Evidence Matrix (Version 0.6) [Companion research workbook]. The paper and workbook should be deposited and distributed together.

Repository record: Persistent URL/DOI pending public deposit. Until then, cite and distribute the versioned paper and workbook together.

Version history: Version 0.5 reconciled the manuscript and matrix. Version 0.6 removes the speculative career-span section, aligns evidence classes, freezes release-specific source links, reports correlation uncertainty, and adds focused task-economics and fiscal-policy context.

Executive Summary

The central question is not whether AI will be good or bad in the abstract. The practical question is whether we can measure who receives the gains, who bears the costs, and which institutions retain authority when AI becomes part of work, education, public infrastructure, and high-impact decisions.

This paper reaches six conclusions.

1. A ten-year sector forecast of AI labor hours is not currently identifiable from public data. The evidence measures paid employee hours, occupational exposure, firm adoption, and the incidence of selected task effects, but not the same unit of task time. The appropriate response is a transparent scenario model with disclosed assumptions and blank fields where evidence is missing.
2. Exposure is not elimination, and its distribution matters. Globally, the ILO's model-derived index places one in four workers in an occupation with some generative-AI exposure and 3.3 percent in the highest exposure gradient, comprising 4.7 percent of female and 2.4 percent of male employment. The gap widens with country income: within high-income countries the highest-gradient shares are 9.6 and 3.5 percent, respectively. The global and high-income figures use different denominators and should not be read as a single series. This supports gender-disaggregated transition monitoring, not a prediction of job loss (Gmyrek et al., 2025). Acemoglu projects modest productivity effects and less inequality amplification than earlier automation, but not a reduction in labor-income inequality (Acemoglu, 2025).
3. A Universal Learning Guarantee is financially assessable only when it is costed in modules. Replacing current net tuition revenue at public colleges and universities is approximately $80.7 billion annually. A mechanical preschool scenario that assumes full participation and unchanged average spending adds about $57.0 billion, producing a partial gross envelope of approximately $137.7 billion. That total is not the cost of universal completion. It excludes living support, new institutional capacity, advising and accessibility, devices and connectivity, graduate and professional capacity, and treatment of existing student debt (State Higher Education Executive Officers Association [SHEEO], 2026; National Institute for Early Education Research [NIEER], 2026; U.S. Census Bureau, 2026c).

4. Healthcare continuity is a separate policy and accounting problem. National health expenditures totaled $5.3 trillion in 2024, but that is the current system's total financing and resource base, not the incremental cost of coverage continuity. In the Congressional Budget Office's illustrative single-payer options, federal subsidies rise by $1.5 trillion to $3.0 trillion in 2030 while total national health expenditures change by between a $0.7 trillion decrease and a $0.3 trillion increase. A public-budget shift and a national resource-cost change are different measures (Centers for Medicare & Medicaid Services, 2026; Congressional Budget Office, 2022).
5. Education-to-work pathways must be evaluated as a funnel. Registered Apprenticeship reports 93 percent employment retention among completers. A DOL-sponsored study estimated median employer benefits of $1.44 per dollar invested, but the estimate includes indirect benefits and a partly projected post-program period; returns varied substantially. Access, start, persistence, completion, placement, earnings, and later employment must be measured separately (U.S. Department of Labor, n.d.; U.S. Department of Labor, 2022; U.S. Department of Labor, Office of Inspector General, 2026).
6. AI is physically coupled to public systems. Lawrence Berkeley National Laboratory estimated that U.S. data centers consumed 176 terawatt-hours of electricity, or 4.4 percent of national use, in 2023 and projected 325 to 580 terawatt-hours, or 6.7 to 12.0 percent, in 2028. Water, grid, ratepayer, and subsidy burdens vary substantially by location. They therefore require site-specific ledgers and community impact agreements, not only national averages (Shehabi et al., 2024).

These findings support a compact built around measurement and reciprocity. Firms should report the task-hour, payroll, quality, redeployment, and infrastructure effects necessary to distinguish augmentation from substitution. Public education should guarantee access through the first occupationally appropriate credential, with graduate education available when public need, occupational requirements, or demonstrated ability justify it. Employers should support apprenticeships, curriculum advising, and pooled financing without owning students, controlling admissions, or repurposing learner data. Health coverage should remain portable across school, employment, caregiving, retraining, and displacement. Governance should combine NIST's risk-management process, auditable industry practices such as Anthropic's Responsible Scaling Policy, binding law, independent review, and rights of appeal and redress.

1. Introduction: From Prediction to Public Reciprocity

Artificial intelligence is frequently described as if it were either an immaterial productivity tool or an inevitable replacement for human work. Neither description is sufficient. AI systems are embedded in organizations, labor processes, electrical grids, water systems, tax structures, educational institutions, and governance decisions. Their social value cannot be inferred from model capability alone.

The most credible policy position is therefore conditional. AI does not, by itself, justify an unlimited welfare state, a punitive tax on every automated tool, or the assumption that mass unemployment is inevitable. It does justify redesigning the social contract if audited evidence shows that adoption is reducing covered payroll, concentrating income in capital, displacing workers faster than they can adapt, or imposing costs on communities that are not reflected in private prices.

That distinction matters because augmentation and substitution can occur at the same time. A system may shorten one task, expand demand for another, create rework, move risk to a reviewer, or change the quality standard expected of the entire workforce. A firm may report AI use without any reduction in paid hours. Conversely, a firm may reduce payroll through a workflow change that is only partly attributable to AI. Policy that taxes a robot, token, or software license cannot resolve those differences. The unit of analysis must be the human task and the economic consequences around it.

This paper asks five questions:

7. What would count as an AI labor hour, and can it be measured with current public data?
8. What is the defensible cost boundary for a public learning guarantee and separately financed health continuity?
9. Which education-to-work outcomes demonstrate that access leads to completion, placement, and earnings?
10. How should electricity, water, grid, ratepayer, and subsidy burdens be attributed to data-center development?
11. What institutional role belongs to voluntary risk frameworks, corporate policies, independent review, and binding law?

The proposed answer is an evidence-based compact. It does not promise that every person will earn a master's degree, that a specific age will become the universal end of paid work, or that one new tax can finance every public guarantee. It proposes a measurable structure through which society can expand human capacity when verified technological gains and public burdens make that expansion both necessary and supportable.

2. Method, Evidence Classes, and Claim Boundaries

The research supporting this paper was organized as a claim-to-source matrix rather than as a narrative literature review alone. Each consequential claim was assigned a year, geography, denominator, evidence class, source, caveat, and policy use. Primary research, official statistics, statutes or public frameworks, and first-party policy documents were preferred over secondary summaries.

Table 1. Evidence classes used in this paper

Evidence class	Definition	Permitted use	Prohibited inference
Observed	Measured or administratively reported for a defined population, year, and denominator.	Baseline and descriptive comparison.	Treating a source's modeled projection as observation.
Derived	Reproducible arithmetic using stated inputs, with no behavioral or future-period assumption added.	Transparent envelope or accounting transformation.	Treating arithmetic precision as measurement accuracy.
Scenario	Result dependent on behavioral, valuation, or future-period assumptions, including a published model.	Sensitivity testing and policy design.	Reclassification as observed merely because a source reports it.
Gap	A decision-relevant value not identified by current compatible evidence.	Research agenda and disclosure.	Filling the gap with an unsupported midpoint or probability distribution.

Source: Author's synthesis and companion AI Transition Evidence Matrix.

This hierarchy is especially important where datasets appear numerically compatible but are not conceptually compatible. An occupational exposure share, a firm adoption rate, and an incidence of substitution among adopting firms may all be expressed as percentages. Their common format does not give them a common denominator. Multiplying them can create false precision.

The evidence also has temporal and coverage limits. Census firm-use measures changed wording in late 2025; current and earlier estimates should not be spliced without adjustment. The BEA/BLS labor-compensation share includes benefits and is measured against value added, not sales. Census Quarterly Financial Report margins cover selected large corporations rather than every firm in a sector. National data-center totals do not identify local water scarcity or rate allocation. The governance comparison examines Anthropic's RSP v3.0 as a documented shift in voluntary commitments, and traces subsequent revisions through v3.4, effective July 8, 2026.

The reconciled companion workbook preserves the evidence matrix, live formulas, standard errors, and source ledger used here. Version 0.6 of the paper and workbook use the same HETH definitions, evidence classes, pathway classifications, claim boundaries, and release-specific source links. The paper presents the findings needed for the argument; the workbook retains the fuller audit trail and all 19 sector rows.

3. The AI Transition Balance Sheet

The paper's first proposed contribution is an AI Transition Balance Sheet. It is a distributional incidence ledger, not a formal balance-sheet identity or a monetized social-welfare function. It prevents private productivity gains from being reported without the corresponding household, fiscal, physical, and governance consequences. The four ledgers are linked but are not additive.

Table 2. The four ledgers of the AI Transition Balance Sheet

Ledger	Measures	Primary decision use
Automation gain	Output, quality, cycle time, human task-hours augmented or substituted, payroll changes, margins, and attribution confidence.	Determines whether a deployment generated a verified and potentially shareable surplus.
Human capacity	Wages, hours, displacement, new jobs, education, completion, placement, health continuity, retraining, and worker participation.	Identifies transition needs and whether gains are broadly shared.
Public infrastructure	Electricity, water, land, grid upgrades, subsidies, ratepayer exposure, public revenue, and local value capture.	Allocates site-specific impact costs and community benefits.
Governance and rights	Risk tier, autonomy, sensitive data, incidents, appeals, human review, reversibility, and evidence integrity.	Determines whether deployment should be allowed, restricted, reviewed, suspended, or prohibited.

Source: Author's framework, informed by the evidence matrix and NIST AI RMF.

The ledgers answer different questions. The firm account records revenue, compensation, purchased inputs, capital costs, and profit. The household account records earnings, benefits, prices, taxes, transfers, and time use. The fiscal account records receipts, outlays, financing shifts, and contingent liabilities. The resource account records changes in real production, human effort, energy, water, capital use, transition services, and external harm relative to a stated counterfactual. Governance and rights impose additional constraints that cannot be waived by a favorable monetary result.

Lost wages matter greatly to affected households, but they are not automatically an aggregate resource loss: the corresponding reduction in a firm's payroll expense may already appear in its gain. Taxes, subsidies, and cash benefits likewise move purchasing power among actors. Adding both sides as separate social gains or losses would double count the same flow. Retraining services use real resources; a training grant is a financing flow, not an additional resource cost on top of the services it purchases.

Any later welfare analysis must consolidate and net transfers, value output and resource use consistently, specify the counterfactual and time horizon, and state distributional weights. Reduced hours may represent involuntary unemployment, valued leisure, or redeployment; the ledger does not assume which. Until those choices and inputs are established, this paper reports incidence and resource effects separately and does not calculate a single Net AI Social Surplus.

Three principles follow.

12. Measurement should precede taxation or subsidy. A contribution should respond to realized conditions, not to an assumption that every AI use replaces labor.
13. Gains and burdens should use compatible units and time periods. A one-time capital cost cannot be compared with a recurring wage reduction without an explicit amortization period.
14. Missing data should remain visible. A blank cell is more credible than an invented midpoint when the missing value determines the policy conclusion.

4. Measuring AI Labor Without False Precision

4.1 Human-equivalent task-hours

An AI labor hour should never mean machine uptime, model latency, token volume, or the number of simultaneous automated processes. Those units may be useful for engineering or billing, but they do not measure labor substitution.

This paper defines one human-equivalent task-hour, or HETH, as one hour of gross counterfactual paid human time required to produce accepted output at an agreed level of quality, risk, and timeliness. HETH is an output-equivalence unit, not a claim that one paid human hour has been removed. Human overhead is deducted only when calculating net substitution.

For sector i:

$$\text{HETH assisted}_i = H_i \times E_i \times A_i \times U_i \times Q_i$$

$$\text{Net task-hours saved}_i = \text{HETH assisted}_i \times S_i$$

$$\text{Paid hours displaced}_i = \max(0, \text{Net task-hours saved}_i) \times (1 - R_i)$$

H is annual paid human employee hours. E is time-weighted task exposure. A is hour-weighted adoption among the exposed workflows. U is the share of adopted, exposed task time routed through AI. Q is the fraction of that counterfactual time represented by outputs meeting the acceptance threshold. It adjusts for acceptable output only; it does not deduct human hours. Each fraction is conditional on the preceding workflow population, not an independent sector-wide percentage.

S is the net fraction of accepted-output counterfactual hours saved after counting all remaining human execution, setup, oversight, validation, rework, and exception time once, including effort spent on failed attempts. A negative S records a net increase in human work and is not converted into positive displacement. R is the fraction of positive savings absorbed by redeployment, rebound demand, or other covered human work. New output and new tasks are recorded separately rather than credited as removal of baseline work.

This operationalization sits within an established task-based literature. Acemoglu and Restrepo (2019) distinguish a displacement effect, as capital assumes tasks formerly performed by labor, from a reinstatement effect, as new tasks restore labor demand. HETH does not replace that framework or claim a new theory of task allocation. It translates a narrower part of it into an auditable reporting unit by separating assistance, net task-time saving, redeployment, and actual paid-hour displacement.

For example, accepted output equivalent to 100 baseline human hours, requiring 40 retained execution hours and 10 additional oversight/rework hours, saves 50 hours: $S = 0.50$. If 20 of those hours are redeployed, $R = 0.40$ and 30 paid hours are removed. Q does not deduct the 10 overhead hours a second time. This is a worked accounting example, not an empirical estimate.

The distinction between assisted, substituted, and displaced hours is deliberate. A task can be assisted without being removed from a job. A task can be removed while the worker is redeployed. A job can change without disappearing. The fiscal and distributional consequences arise only after those transitions are measured.

4.2 What current evidence actually measures

BLS July 2026 seasonally adjusted data report 135.588 million private nonfarm payroll jobs and 34.3 average weekly hours. Multiplying those inputs by 52 weeks gives an annualized snapshot, not observed annual hours or a ten-year projection (U.S. Bureau of Labor Statistics, 2026a):

$$\mathrm{H} = 135{,}588{,}000 \text{ payroll jobs} \times 34.3 \text{ hours/week} \times 52 \text{ weeks} \approx 241.8 \text{ billion hours/year}$$

The unrounded product, 241,834,756,800, is retained in Table 3A only so that the illustration can be reproduced exactly. It is not a claim to that level of accuracy. The measure excludes government, most agriculture, and the self-employed; multiple jobholders may appear on more than one payroll. It is also provisional. BLS published a preliminary Current Employment Statistics benchmark estimate for March 2026 of −178,000 for total private employment on August 28, 2026. That figure is not incorporated into published estimates; the final benchmark revision is scheduled for February 2027 with the January 2027 Employment Situation, and the preliminary benchmark is reported not seasonally adjusted, whereas the inputs above are seasonally adjusted. Treated only as an order-of-magnitude indication, a revision of that size applied at the same hours and weeks would shift H by roughly 0.3 billion hours—approximately ten times the entire displacement result of the lower illustration in Table 3A. The displayed precision of H is therefore an artifact of reproducibility, not measurement accuracy. H supplies an illustrative baseline, not the remaining task-level multipliers (U.S. Bureau of Labor Statistics, 2026a, 2026c, 2026d).

The ILO's refined index places 3.3 percent of global employment in its highest exposure gradient: 4.7 percent of female employment and 2.4 percent of male employment. In high-income countries, the corresponding rates are 9.6 and 3.5 percent, about a 2.7-fold difference. Across all exposure gradients, 34 percent of high-income employment is exposed. These occupational patterns warrant gender-disaggregated access and transition measures; they do not identify displacement. The index combines task data, expert input, and AI-model predictions, and it equally weights tasks because comparable frequency and importance data are unavailable. The companion matrix therefore classifies these exposure estimates as Scenario evidence under Table 1, not as observed or time-weighted HETH (Gmyrek et al., 2025).

Acemoglu's task-based model estimates no more than a 0.66 percent increase in total factor productivity over ten years, falling below 0.53 percent after accounting for harder-to-learn tasks. These are assumption-dependent model results, not observed productivity gains. The analysis anticipates less inequality amplification than earlier automation because exposure is more broadly distributed, while finding no evidence of reduced labor-income inequality and projecting a wider capital-labor income gap. Both conclusions matter: limited macroeconomic gains do not establish catastrophic inequality, and greater task-level productivity does not guarantee equal gains (Acemoglu, 2025).

Census Business Trends and Outlook Survey evidence adds adoption and task-use information. In late 2025 and early 2026, 17.9 percent of firms reported AI use on a firm-weighted basis and 32 percent on an employment-weighted basis. Among AI-using firms, 10.1 percent reported using AI for a task previously performed by an employee, 43.7 percent reported augmentation, 10.6 percent reported a new task, and approximately 2 percent reported lower employment. These are nonexclusive incidence rates among firms. They do not reveal the share of workers, task-hours, or wages affected (Bonney et al., 2026).

4.3 Sector evidence and denominator conflicts

Table 3. Selected sector indicators from different statistical vintages

Sector	Top-exposure-quintile employment	Firms using AI	Adoption SE	Prior employee task incidence	Augmentation incidence	Labor compensation share of value added
Information	71.1%	37.6%	1.0 pp	16.8%	59.5%	31.5%
Finance and insurance	100.0%	30.4%	0.8 pp	13.7%	54.9%	52.1%
Professional, scientific, and technical services	93.9%	34.2%	0.2 pp	15.5%	60.5%	76.4%
Educational services	3.1%	30.7%	1.1 pp	12.8%	57.0%	83.7%

Sector	Top-exposure-quintile employment	Firms using AI	Adoption SE	Prior employee task incidence	Augmentation incidence	Labor compensation share of value added
Health care and social assistance	0.3%	20.7%	0.5 pp	9.5%	44.8%	84.1%
Manufacturing	6.2%	12.4%	0.3 pp	7.5%	43.8%	48.3%
Retail trade	3.9%	12.6%	0.2 pp	8.7%	33.7%	58.5%

Column sources: Exposure = Tucker (2026), Table 2, all-ages share of sector employment in top-quintile detailed industry-states; adoption, its standard error, and task incidence = Census BTOS AI Supplement (2026a); labor share = BEA/BLS Integrated Industry-Level Production Account (2026), 2024 data. Adoption is firm-weighted; its standard error is reported in percentage points; task effects are conditional on AI users and nonexclusive. The table is intentionally a cross-source diagnostic with mixed vintages and incompatible denominators, not a comparable panel or model input. These columns must not be multiplied.

Table 3 is a purposive seven-sector illustration drawn from the companion matrix's 19 rows; it is neither ranked nor exhaustive. The displayed sectors show contrasting profiles with complete fields. Management of companies and enterprises, which has 98.1 percent of sector employment in Tucker's top exposure quintile, is omitted because several BTOS fields are suppressed and its firm-adoption estimate has an 8.09-percentage-point standard error. The full matrix retains it. Finance and insurance's 100.0 percent means every detailed industry-state unit assigned to that sector falls in Tucker's top quintile; it does not mean every finance worker or task is exposed.

The ILO gradients use a different classification. Information, finance, and professional services combine high exposure proxies with high adoption, whereas education and health care illustrate how adoption and labor intensity can diverge from that proxy. These patterns do not establish how many paid hours will disappear.

Margins add another boundary. BEA/BLS capital cost share is the complementary production-account envelope, not profit. Census Quarterly Financial Report after-tax margins apply only to selected covered large corporations. In the cited quarter, the information-sector margin was 32.6 percent, manufacturing 13.8 percent, mining 12.3 percent, retail 5.8 percent, wholesale 2.5 percent, and professional, scientific, and technical services negative 0.5 percent within the QFR coverage. Those figures should not be generalized to every firm or interpreted as AI-attributable profit (U.S. Census Bureau, 2026d).

4.4 The defensible conclusion

The evidence assembled here does not identify E, A, U, Q, S, and R in one compatible sector frame. A ten-year HETH result based on assumed values is therefore a scenario, not an empirically identified forecast. The companion model intentionally leaves these inputs blank. Table 3A fills them only to demonstrate sensitivity, retaining the unrounded 241.8347568-billion-hour H input solely so the arithmetic can be reproduced rather than to imply precision or project economic growth.

Table 3A. Reproducible HETH sensitivity illustration, not a forecast

Parameter or output	Lower-input illustration	Higher-input illustration	Own-range swing
E: time-weighted exposure	20%	50%	2.500×
A: adoption conditional on exposed hours	20%	60%	3.000×
U: AI routing conditional on adopted exposed hours	20%	50%	2.500×
Q: accepted-output share, time weighted	80%	90%	1.125×
S: net task-time saving, including all human overhead once	10%	40%	4.000×
R: share of saved time absorbed by other covered work	80%	50%	2.500×
Displaced share: E × A × U × Q × S × (1 - R)	0.0128%	2.7000%	210.9375×
Annual paid hours displaced	0.03095 billion	6.52954 billion	n/a
Annualized FTE equivalents at 2,080 hours each	About 14,882	About 3,139,201	n/a

Source: Author's scenario calculations. H = 135,588,000 × 34.3 × 52; every other parameter is assumed, not estimated from Table 3. Own-range swing is the factor by which displaced share changes when that parameter alone moves across its stated range. For R, the swing is computed on (1 - R), so a higher R produces a lower displaced share. Outputs use unrounded inputs; displayed results are rounded. FTE conversion is arithmetic, not a count of jobs or people.

For the lower-input illustration, 241.8347568 billion × 0.20 × 0.20 × 0.20 × 0.80 × 0.10 × 0.20 = 0.030954849 billion displaced hours. For the higher-input illustration, 241.8347568 billion × 0.50 × 0.60 × 0.50 × 0.90 × 0.40 × 0.50 = 6.529538434 billion hours. The outputs differ by 210.9375 times.

That 211-fold figure requires interpretation because it is not a measurement. Because the chain is a pure product, the ratio between the two illustrations is exactly the product of the six individual own-range swings: 4.000 × 3.000 × 2.500 × 2.500 × 2.500 × 1.125 = 210.9375. The spread is therefore an arithmetic consequence of the six ranges chosen and of moving every parameter to the same end simultaneously. It is not evidence about the world or a plausible interval. The two illustrations are opposite corners of a six-dimensional box, and corners are not representative of a box.

The own-range swings are more informative. In a multiplicative chain, moving one parameter from its lower to higher endpoint changes the result by the ratio of those endpoints, regardless of the other five values. Net task-time saving (S) dominates at 4.000×, followed by adoption (A) at 3.000×, with exposure, AI routing, and redeployment (E, U, and R) each at 2.500×. This ranking directly informs the reporting standard below: workflow-level measurement of S and A closes more of the gap than refinement of any other input.

Accepted-output share (Q) moves the result by only 1.125× across its stated range. Q therefore functions here as a definitional guardrail, not as a driver of the range. Its narrow assumed range should not be read as evidence that acceptance rates are well established.

Method note: This paper does not convert the range into a probability interval. Doing so would require marginal distributions for all six parameters and an assumption about their joint dependence. Neither is available, and Section 4.3 gives direct reason to expect that exposure and adoption are positively correlated across sectors, which independent sampling would misrepresent. Reporting the swings and declining to report an interval is consistent with the Gap class in Table 1.

Neither FTE equivalent counts eliminated jobs: hours can be spread across workers, absorbed through vacancies, or reflect shorter schedules. No year-specific 2036 outcome or AI-tax revenue follows from this illustration. A defensible projection would also require future output, employment, working-time, and adoption paths.

Closing the gap requires firm and public reporting at the workflow level. At minimum, a reporting standard should capture:

15. The defined output and quality threshold.
16. Paid human time before deployment.
17. Human setup, supervision, validation, rework, and exception time after deployment.
18. The share of workflow handled by AI and the failure or escalation rate.
19. Whether time was substituted, redeployed, or converted into increased output.
20. Changes in covered payroll, benefits, contractor spending, employment, and hours.
21. Changes in price, output, revenue, and operating cost.
22. The attribution method and confidence interval.

Only after those fields are observed should a firm claim a verified automation saving or government calculate a displacement-triggered contribution.

5. A Universal Learning Guarantee

5.1 From credential promise to human-capacity infrastructure

The policy goal should not be a master's degree for every person. That formulation risks credential inflation, unnecessary cost, and the assumption that every occupation requires the same educational pathway. The stronger guarantee is access through the first occupationally appropriate credential, with graduate or professional study available when occupational requirements, public need, or demonstrated ability justify it.

The Universal Learning Guarantee should include:

23. Publicly financed pre-K through grade 12.
24. A connected device, broadband access, accessible learning tools, and human support.
25. Tuition-free public pathways through an apprenticeship, certificate, associate degree, bachelor's degree, or first appropriate graduate credential.
26. Paid work-based learning and employer participation without employer ownership of admissions, curriculum, credentials, or student data.
27. Completion support, advising, accessibility, and placement accountability.
28. Lifelong learning accounts that remain available after initial entry into work.

Education in this model is a public talent commons, not a corporate minor league. Employers may help identify changing skills, sponsor paid learning, mentor students, and contribute to pooled funds. Public institutions, educators, learners, and democratic authorities must retain control of academic standards and individual rights.

5.2 What current costs can and cannot establish

Table 4. Learning-guarantee cost boundary

Module	Current evidence or proxy	Annual or per-pathway value	Evidence class	Boundary
Existing K-12 system	Current spending per pupil, FY2024	$17,619 per pupil	Observed	Existing public baseline, not an incremental guarantee cost.
Public higher-education tuition replacement	FY2025 net tuition revenue	$80.698 billion per year	Observed	Gross financing shift for current FTE enrollment; not universal-completion cost.
Preschool expansion	Population gap x current all-source spending	About $57.0 billion per year	Scenario	Full-population slot illustration at fixed unit spending; not a fiscal upper bound.
Partial learning envelope	Tuition replacement + preschool envelope	About $137.7 billion per year	Scenario	Excludes living support, capacity, completion, devices, graduate slots, and debt.
Certificate proxy	One two-year-sector FTE-year	$13,722	Scenario	Revenue envelope, not credential marginal cost.
Associate proxy	Two two-year-sector FTE-years	$27,444	Scenario	Ignores attrition, transfer, field, and capacity.
Bachelor's proxy	Four four-year-sector FTE-years	$86,032	Scenario	Not a causal estimate of completion cost.
Master's increment proxy	Two additional four-year-sector FTE-years	$43,016	Scenario	Public accounts do not isolate graduate marginal cost.

Sources: U.S. Census Bureau School System Finances; SHEEO FY2025; NIEER State of Preschool 2025; U.S. Census Bureau Vintage 2025 population estimates. Credential rows are FTE-year revenue envelopes, not marginal cost estimates.

SHEEO reports 10,819,079 public full-time-equivalent students in FY2025, $7,459 in net tuition per FTE, and $80,698 million in aggregate net tuition revenue. Its published aggregate, $80.698 billion, is the tuition-replacement baseline. Multiplying the separately rounded FTE and per-FTE values gives $80.699510261 billion; that small rounding difference is not a new cost estimate. Enrollment should not be back-calculated from those rounded dollar figures. Total education revenue is $19,443 per FTE, with two-year and four-year sector values of $13,722 and $21,508. These revenue envelopes do not isolate cost by credential or field (SHEEO, 2026).

SHEEO's all-public and sector tables are separately reported series, not additive components. The sector FTE values sum to 10,819,084, five above the all-public count. Reconstructing net tuition from the rounded sector inputs produces $80.986 billion, $0.288 billion above the published all-public aggregate. Reconstructing total education revenue produces $200.241 billion versus a $210.355 billion all-public proxy, a 4.81 percent difference. SHEEO explains that sector reporting excludes agency funding and includes only allocated federal stimulus, while some support and aid cannot be allocated by sector. These differences should be disclosed, not treated as arithmetic errors or used to overwrite either source-defined series (SHEEO, 2026).

The preschool illustration uses the Census Vintage 2025 population of 7,504,421 children ages three and four, approximately 1,800,000 children in state-funded preschool, and NIEER's $9,988 all-reported annual spending per enrolled child. The enrollment input is rounded, so the resulting envelope must also be reported approximately (NIEER, 2026; U.S. Census Bureau, 2026c):

$$\text{Additional slots} = 7{,}504{,}421 - 1{,}800{,}000 = \text{approximately } 5{,}704{,}421$$

$$\text{Preschool envelope} = 5{,}704{,}421 \times \$9{,}988 = \text{approximately } \$56.976 \text{ billion/year}$$

$$\text{Partial learning envelope} = \$80.698 \text{ billion} + \$56.976 \text{ billion} = \text{approximately } \$137.7 \text{ billion/year}$$

This mechanical envelope assumes full participation and unchanged average spending. It does not net out Head Start or other non-state provision, align every enrollment age exactly, model family preferences, or price new facilities and higher quality. It is therefore neither a fiscal score nor a cost ceiling. Credential proxies in Table 4 similarly multiply sector revenue per FTE by assumed years: $13,722 x 1 or 2, and $21,508 x 4 or 2. Graduate programs may have very different marginal costs.

Student budgets reveal an additional boundary. College Board estimates a $21,320 annual budget for a public two-year commuter student, including $4,150 in tuition and fees, leaving $17,170 in non-tuition expenses. The public four-year in-state on-campus budget is $30,990, including $11,950 in tuition and fees, leaving $19,040 for other expenses. These are student budget estimates rather than public expenditures, but they show why replacing tuition alone cannot guarantee participation or completion (College Board, 2025).

The correct fiscal approach is modular. Lawmakers should specify the eligible population, benefit amount, take-up, existing-program offsets, capital requirements, staffing model, and target outcomes for each module. Existing student debt should remain a separate stock-policy question. Forward tuition reform does not determine whether or how prior debt is treated.

6. Healthcare Continuity as a Separate Guarantee

Health coverage should not be buried inside an education appropriation. Education funding pays for learning capacity. Health coverage pays for medical care and financial risk protection. Combining them into one estimate obscures both.

The school-to-work transition creates a practical case for continuity. The uninsured rate among adults ages 19 to 25 was 14.3 percent in 2024, the highest of the age groups reported by the Census Bureau. The statistic does not isolate students or recent graduates, but it identifies a vulnerable transition period (U.S. Census Bureau, 2025).

Table 5. Health-continuity accounting choices

Policy design	Core cost equation	Required inputs	Central accounting boundary
Fixed bridge	Eligible transitions x monthly net premium x covered months x take-up	Transition population, benchmark plan, months, take-up, and existing coverage offsets	A bounded transition benefit, not universal system reform.
Portable public plan or Medicaid-style pathway	Eligible enrollment x net program cost	Payment rates, provider participation, benefits, use, administration, and employer/household offsets	Requires program design and microsimulation.
Universal or single-payer system	Gross public outlays - converted public/private financing and administrative offsets	Benefits, payment rates, use, taxes, employer premium conversion, and transition rules	Federal-budget effects and national health-expenditure effects must be reported separately.

Sources: CMS National Health Expenditure Accounts; Congressional Budget Office single-payer testimony; KFF Employer Health Benefits Survey; U.S. Census Bureau health-insurance estimates.

National health expenditures were $5.3 trillion in 2024. Private health-insurance spending accounted for approximately $1.645 trillion, and out-of-pocket spending approximately $556.6 billion. These are existing financing burdens and resource costs. They are not all additive costs of a new guarantee (Centers for Medicare & Medicaid Services, 2026).

The CBO scenarios make the distinction concrete. Depending on payment rates, benefits, administrative savings, and use, federal subsidies increase by $1.5 trillion to $3.0 trillion in 2030 while total national health expenditures change by between a $0.7 trillion decrease and a $0.3 trillion increase. A proposal may therefore produce a major federal financing shift while reducing, increasing, or approximately preserving total national spending (Congressional Budget Office, 2022).

KFF's employer benchmarks show the private financing already at stake: the average family premium was $26,993 in 2025, the average worker contribution for family coverage was $6,850, and the average single-coverage deductible among workers with a deductible was $1,886. These values do not determine a public-plan cost. They identify amounts that households and employers already pay and that must be included as offsets or conversions in any full reform model (KFF, 2025).

For the first phase of this compact, the narrowest credible commitment is a defined health-coverage bridge across education, apprenticeship, displacement, and re-employment. A portable public plan or universal system may ultimately be preferable, but it should be debated and scored on its own terms.

7. Apprenticeship, Completion, Placement, and Earnings

Access is not the same as completion, and completion is not the same as economic mobility. A human-capacity system should publish a pathway funnel for each credential and apprenticeship program:

$$\text{Eligible} \rightarrow \text{Enrolled} \rightarrow \text{Started} \rightarrow \text{Retained} \rightarrow \text{Completed} \rightarrow \text{Placed} \rightarrow \text{Earnings} \rightarrow \text{Persistence}$$

Each transition needs a numerator, denominator, time window, and equity breakdown. A program that reports placement only among completers may conceal failure before completion. A program that reports enrollment without start or persistence may reward recruitment rather than development.

Table 6. Selected pathway evidence and the correct interpretation

Measure	Reported result	Population or denominator	Evidence class	What it supports	What it does not establish
Registered Apprenticeship retention	93% employed after completion	Apprenticeship completers	Observed	Descriptive completer benchmark.	Completion rate or causal effect.
Employer apprenticeship return	Median $1.44 benefits per $1 invested	68 surveyed AAI employers	Scenario	Conditional employer business case.	Uniform, directly observed, or annual 44% return.
Wage at completion	$22.67 per hour	Apprentices represented in employer survey	Observed	Within-study wage benchmark.	National post-completion earnings.
Public two-year completion	29% within 150% of normal time	First-time, full-time students at initial institution	Observed	Need to track completion and transfer.	All-student or part-time outcome.
Public four-year completion	63% within six years	First-time, full-time students at initial institution	Observed	Institutional completion benchmark.	Causal return for each pathway.

Measure	Reported result	Population or denominator	Evidence class	What it supports	What it does not establish
Scaling Apprenticeship grants audit	62% no start; 27% incomplete; 4% completed without employment; 7% completed with employment	30,826 unemployed or underemployed participants	Observed	Full start-to-employment funnel; percentages rounded.	National Registered Apprenticeship failure rate.

Sources: Apprenticeship.gov; U.S. Department of Labor AAI Employer ROI Final Report; DOL Office of Inspector General; NCES Undergraduate Retention and Graduation Rates.

The employer ROI study estimates the apprenticeship period plus five years after completion, using survey data and projections, discounted at 3 percent. Its medium valuation produces median net ROI of 44.3 percent; the 25th and 75th percentiles are -7.4 and +120.7 percent. Including only direct productivity benefits yields median net benefits of about -$7,700 under the low valuation and -$4,900 under the medium valuation. Positive median total returns depend on indirect benefits such as reduced turnover and improved co-worker productivity. This is Scenario evidence, not an observed universal payoff (U.S. Department of Labor, 2022, pp. ES-4-ES-6).

The OIG audit supplies all four funnel outcomes: 19,120 did not start, 8,223 started without completing, 1,208 completed without an employment outcome, and 2,275 completed and entered training-related or unrelated employment at exit. The counts sum to 30,826. The 7 percent success category is reported directly, not inferred from a rounded residual, and does not establish sustained or training-related placement (U.S. Department of Labor, Office of Inspector General, 2026, pp. 17-18).

The same funnel discipline changes how the public two-year result is read. For the Fall 2017 first-time, full-time cohort, 29 percent completed at the initial institution, 16 percent transferred, and 12 percent remained enrolled at 150 percent of normal time. The residual 43 percent had no credential, transfer, or continuing enrollment in that reporting frame. This derived residual does not explain why students left or what happened later; it shows why completion alone is an incomplete outcome measure (National Center for Education Statistics, 2026).

BLS earnings data reinforce the value of education but should be interpreted cautiously. In 2025, median weekly earnings were $1,876 for workers with a master's degree, $1,578 with a bachelor's degree, $1,135 with an associate degree, and $966 with a high-school diploma. These are cross-sectional associations. They combine occupation, experience, selection, geography, and other differences and do not prove that assigning every learner to a longer credential would cause the observed earnings gap (U.S. Bureau of Labor Statistics, 2026b).

The policy implication is not to choose one pathway for everyone. It is to require every publicly financed pathway to disclose completion, transfer, placement, earnings, persistence, student debt, and employer participation using compatible definitions. Corporate sponsorship should be pooled and portable. A student should not owe service to one sponsor merely because the sponsor contributed to the public system.

8. Data Centers and the Physical Burden of AI

AI appears digital at the point of use, but the supporting infrastructure is physical. Data centers require land, generation, transmission, distribution, cooling, water, backup systems, and public approvals. The burden varies by site and time of day, so national averages are planning signals rather than local impact estimates.

Table 7. Data-center electricity, water, grid, and subsidy evidence

Burden	Evidence	Geography and year	Decision boundary
Electricity use	176 TWh; 4.4% of U.S. electricity	United States, 2023	All data centers, not AI-only load.
Electricity projection	325-580 TWh; 6.7-12.0%	United States, 2028	Wide scenario range sensitive to utilization and technology.
Direct water	66 billion liters	United States, 2023	National total conceals basin scarcity and cooling design.
Indirect water	Approximately 800 billion liters	United States, 2023	Associated with electricity generation and generation mix.
Direct-water projection	150-275 billion liters	United States, 2028	Scenario range, not site allocation.
Summer peak-demand growth	More than 224 GW; 69% above the prior forecast	North America, ten-year planning horizon	All-load forecast; data centers are one major contributor.
Dedicated state incentives	38 states	United States, 2026	Prevalence, not a harmonized national dollar total.
Georgia sales-tax case	$474.2 million projected forgone state sales tax	Georgia, 2025	State-specific evaluation; not nationally generalizable.

Sources: Shehabi et al. (2024); NERC 2025 Long-Term Reliability Assessment; NCSL (2026); Georgia Department of Audits and Accounts (2025).

LBNL's 2023 electricity and water quantities are historical-year model estimates, and its 2028 ranges are projections. Under Table 1, all are Scenario evidence rather than Observed measurements. The projection implies that data centers could account for a materially larger share of U.S. electricity by 2028, but it should not be restated as an AI-only forecast. The NERC assessment likewise identifies data centers among major drivers of load growth while reporting a system-wide forecast across all loads. Good policy preserves those boundaries.

Electricity and water should be evaluated jointly. Evaporative heat rejection can reduce cooling electricity while consuming water; dry or hybrid designs change both demands. LBNL also models liquid-cooling configurations that can improve both water and energy efficiency, so a universal one-for-one tradeoff would be misleading (Shehabi et al., 2024, Table 4.2). Site approval should compare cooling configurations at matched computing load and weather, disclose power usage effectiveness (total facility energy divided by IT energy) and water usage effectiveness with explicit boundaries, and include direct basin consumption plus electricity-related water. Liquid cooling alone does not identify the site's water burden.

Tax incentives create a second burden. NCSL identifies dedicated data-center incentives in 38 states, including electricity-related incentives in 14 states and statewide property-tax incentives in 11. Twenty-three states include job requirements, 13 include sunsets, and at least 23 include statutory clawbacks. The prevalence of incentives is observable; a consistent authoritative national subsidy total is not (National Conference of State Legislatures, 2026).

Every major data-center approval should therefore include a public impact ledger containing:

29. Incremental energy use by hour and season, generation source, and reserve requirement.
30. Interconnection, transmission, distribution, and reliability upgrades, with the responsible payer identified.
31. Cooling design, paired energy/water scenarios, direct and indirect water use, basin conditions, discharge, and recycling.
32. Tax exemptions, grants, discounted land or power, and other public support.
33. Ratepayer exposure and any special tariff.
34. Local employment, wage, procurement, and community-benefit commitments.
35. Decommissioning, restoration, and financial-assurance requirements.

A community impact charge should recover attributable public costs. It should remain separate from a labor-displacement contribution because the two instruments address different externalities.

9. Governance: Voluntary Practice, Independent Oversight, and Law

AI governance fails when one instrument is asked to do everything. A risk-management framework can organize decisions without creating legal rights. A corporate policy can produce operational controls without democratic authority. A statute can create duties and enforcement while still depending on technical standards and firm-level implementation.

Table 8. Governance comparison

Dimension	NIST AI RMF 1.0	Anthropic RSP v3.0 and current RSP	EU AI Act comparator
Legal force	Voluntary framework.	Voluntary company policy.	Binding law with public enforcement.

Dimension	NIST AI RMF 1.0	Anthropic RSP v3.0 and current RSP	EU AI Act comparator
Scope	General lifecycle and context-specific AI risk.	Catastrophic risks from advanced frontier systems; explicitly not comprehensive.	Risk-tiered AI, including general-purpose and specified high-risk uses.
Core structure	Govern, Map, Measure, Manage.	Company requirements, nonbinding roadmap goals, risk reports, and competitor-contingent commitments; broader industry recommendations.	Prohibitions, high-risk duties, transparency requirements, and general-purpose AI obligations.
Independent review	Encourages separation of builders, users, verifiers, and validators.	Seeks credible external review, with process defined through company policy.	AI Office, national authorities, scientific panel, courts, and other public bodies.
Enforcement and sanctions	None supplied by the framework.	Internal governance, contractual, and reputational consequences.	Corrective measures and fines through public authorities.
Rights and remedies	Must be supplied by other law or policy.	Outside the central frontier-catastrophe scope.	Public duties and enforcement; applicability dates vary by obligation.
Best institutional role	Shared risk language and lifecycle process.	Auditable operational case study and evidence of industry practice.	Legal floor, public authority, and remedies.

Sources: NIST (2023); Anthropic (2026a, 2026b); European Commission (2026). RSP v3.0 marks a change in voluntary commitments; v3.4 was current at review. External policy analysis: Williams and Freund (2026), not a compliance audit.

NIST's AI RMF is valuable because it establishes a common process: governance should be continuous, context should be mapped, risks should be measured, and responses should be managed. It also encourages meaningful independence between the system builder and the verifier. Its voluntary status, however, means that it does not itself create sanctions, remedies, or public risk tolerances (National Institute of Standards and Technology, 2023).

Anthropic's RSP v3.0, effective February 24, 2026, changed the governance bargain. It moved away from the previous general commitment to pause when required mitigations were unavailable, separating company requirements from broader industry recommendations and competitor-contingent commitments in Appendix A. Roadmap goals are not hard commitments. Existing ASL-3 safeguards were retained; the revision should not be described as abandonment of all safety controls (Anthropic, 2026a).

An external analysis by GovAI researchers Williams and Freund (2026) identifies the loss of the general pause commitment and weaker future security commitments, while recognizing the value of new risk reports and roadmaps. It also notes that the older policy already contained a competitor-related escape clause. This is a research commentary, not peer-reviewed proof or independent assurance of compliance. Its balanced interpretation helps distinguish transparency gains from reduced advance commitments about future conduct.

The policy changed four more times in under five months: v3.1 on April 2, v3.2 on April 29, v3.3 on May 26, and v3.4 on July 8, 2026. Changes included threshold definitions, external-review governance, and who receives unredacted risk reports. Some strengthened oversight; not every revision weakened it

(Anthropic, 2026b). The institutional lesson is the revisability of private commitments, not that revision itself proves misconduct. Public governance needs version-specific evidence and enforceable minimums that do not depend solely on a company's competitive assessment.

The EU AI Act provides a binding comparator. Public authorities possess supervisory and enforcement powers, including access to documentation, evaluations, corrective measures, and fines. Implementation dates vary, and the 2026 amendments changed the timing of certain high-risk requirements. The institutional distinction remains clear: voluntary codes may support compliance, but they do not replace public law (European Commission, 2026).

9.1 A layered governance model

The strongest arrangement combines four layers:

36. A rights floor established by law, including privacy, nondiscrimination, due process, child protection, accessibility, incident reporting, and redress.
37. Independent oversight with authority, access, conflict rules, publication standards, and meaningful consequences.
38. Technical and organizational risk management using frameworks such as NIST AI RMF.
39. Auditable voluntary practice that can move faster than law but cannot waive it.

Local communities should be able to adopt stricter implementation choices, particularly for child-facing technologies. Local control must sit above a national rights floor. Communities should not be able to waive privacy, accessibility, nondiscrimination, or due process for the people within them.

10. Financing Through Verified Gains and Attributable Burdens

A flat robot tax is too crude. It can penalize systems that assist workers, improve accessibility, or increase quality without reducing paid work. It is also easy to avoid because there is no stable boundary around a robot, software feature, or AI token.

That design choice also responds to the fiscal-policy literature. Brollo et al. (2024) do not recommend special taxes on AI intended to slow investment, citing operational difficulty and risks to productivity. They instead emphasize correcting tax incentives that favor rapid labor displacement, strengthening general taxation of capital income, and adapting social protection. The contribution proposed here is therefore not justified by AI use alone. It is a conditional, testable supplement to broader capital and rent taxation, activated only when audited effects meet the stated trigger and paired with adjustment bands and worker-sharing credits.

The preferred instrument is an Automation Dividend Contribution based on audited economic effects:

$$\text{Labor contribution} = \max(0, \tau \times \max(0, G - N) - \kappa \times W)$$

$$\text{Total obligation} = \text{Labor contribution} + I$$

G is verified AI-attributable private gain, net of deployment costs and required infrastructure cost recovery but before the worker-sharing expenditure claimed for credit. N is a normal-return allowance; tau is the contribution rate; W is eligible, additional worker-sharing expenditure; and kappa is the credit rate. I is attributable public infrastructure cost not already recovered through tariffs, fees, developer payments, or other obligations. All quantities must use the same assessment period. The outer maximum makes the labor credit nonrefundable and caps it at the labor liability; it cannot offset I.

This is a credit against liability, not a deduction from the contribution base. With a 20 percent contribution rate, a 100 percent credit rate, and sufficient pre-credit liability, $1 million in eligible sharing reduces the contribution by $1 million. A $1 million base deduction would reduce it by only $200,000. If liability is smaller than the credit, the current-period credit is capped at that liability. Carryforwards, if any, require a separate rule. These rates illustrate the distinction; they are not recommended or estimated policy parameters.

The rates, normal-return allowance, thresholds, and eligibility rules remain legislative choices. Qualifying expenditure must be additional, verifiable, and not both deducted in G and credited through W. Subsidized or already-reimbursed expenditure cannot receive duplicate credit. Rules also need baseline wages, benefit standards, related-party safeguards, and audits to prevent relabeling ordinary compensation as new sharing. This paper assigns no revenue estimate because the required gain and attribution data are not yet observed.

A trigger should require more than the presence of AI. It should combine evidence that output, revenue, or unit cost improved; covered payroll or compensated hours fell beyond an adjustment band; the relevant AI deployment and expense are documented; and an independent attribution method supports the relationship. Temporary changes caused by recession, divestiture, outsourcing, or unrelated reorganization should be separated.

Qualified credits should reward firms that share productivity gains through:

40. Higher wages or broad-based profit sharing.
41. Shorter working hours without reduced pay.
42. Paid retraining and lifelong learning accounts.
43. Registered apprenticeships and public education partnerships.
44. Transition placement and portable health coverage.

45. Creation of new covered employment.
46. Accessibility and demonstrable public-benefit uses.

The infrastructure component should recover attributable, unrecovered grid, water, subsidy, and community costs even when no labor displacement occurs, with a legal basis and no duplicate recovery. Worker-sharing credits must not erase those obligations. Conversely, a labor contribution should not be excused merely because a firm purchased renewable energy if payroll displacement remains substantial. The ledgers are connected, but their charges answer different questions. Statutory corporate liability also does not establish final economic incidence: some burden may pass to prices, wages, or investment, which evaluation must test.

The broader financing stack should include existing program spending, progressive general revenue, capital and excess-rent taxation, employer financing converted from current benefits where applicable, automation contributions, and site-specific impact charges. No narrow AI tax should be presented as sufficient to finance education, healthcare, food, utilities, debt relief, and transition income simultaneously.

11. Counterarguments, Limitations, and Falsifiability

11.1 AI may remain primarily complementary

The near-term evidence currently supports this countercase. Augmentation is reported more often than prior-employee-task substitution among AI-using firms, and only a small share reports employment decreases. If paid hours, wages, benefits, and covered employment rise after redeployment, a displacement-triggered contribution should be small. The compact is designed to respond to measured substitution rather than presume it.

An unweighted exploratory reconciliation of the 15 sector rows with reported employment-decrease incidence produces $r = +0.289$ for the exposure proxy (95% CI [−0.262, +0.698], $p = .297$) and $r = +0.198$ for firm adoption (95% CI [−0.349, +0.645], $p = .478$). Both estimates are statistically indistinguishable from zero, and both intervals permit negative or materially positive associations. The inputs also combine incompatible denominators and, for exposure, different vintages. These calculations carry no causal or predictive weight; they demonstrate what the present cross-sector data do not identify (Tucker, 2026; U.S. Census Bureau, 2026a).

11.2 Attribution will be difficult

Firms change technology, products, organization, and staffing simultaneously. That does not make attribution impossible, but it argues for adjustment bands, multi-period evidence, independent methods, and a right to challenge the assessment. Where attribution confidence is low, policy should rely more heavily on broad capital and rent taxation and less on a firm-specific automation charge.

11.3 A contribution may slow beneficial innovation

Poorly designed taxes can do so. This is why the proposal excludes token, model-call, and equipment taxes and includes worker-sharing credits, normal-return allowances, and a displacement trigger. A firm that uses AI to improve quality, safety, accessibility, or worker productivity without externalizing costs should not be treated like a firm that removes compensation while shifting transition and infrastructure costs to the public.

11.4 A learning guarantee may inflate credentials

It will if the program equates education with a single academic ladder. The first occupationally appropriate credential, apprenticeship parity, periodic program review, and completion and earnings measures are safeguards against that outcome. Education must also preserve civic, cultural, and personal-development purposes that cannot be reduced to first-year salary.

11.5 Health continuity may become an unbounded fiscal commitment

That risk is real if coverage design, payment rates, and financing offsets remain unspecified. The answer is separate scoring. A fixed bridge, a portable public plan, and a universal system are different policies with different fiscal and resource effects.

11.6 Local infrastructure charges may shift investment elsewhere

Some relocation is possible. A national reporting floor and regional coordination can reduce a race to the bottom. The alternative is not costless investment; it is often a transfer of grid, water, tax, and ratepayer burdens to communities with less bargaining power.

11.7 Governance can become rigid or captured

Law can lag technology, while private standards can be written around the interests of their authors. A layered system reduces dependence on either one. Sunset review, version control, public incident reporting, independent access, appeals, and the ability to reduce or suspend a system's authority provide mechanisms for correction.

The thesis should be narrowed or rejected if evidence shows that AI adoption consistently raises covered hours, wages, and broad-based income without externalizing public costs; if competition disperses the surplus rather than concentrating it; if data-center developers fully internalize grid, water, ratepayer, and subsidy burdens; or if the learning and transition guarantees fail to improve completion, placement, earnings, health continuity, or equitable access relative to less costly alternatives.

12. Implementation Agenda

The compact should advance through evidence gates.

47. Establish a HETH measurement standard. A federal statistical and standards working group should define workflow units, quality adjustment, oversight and rework, substitution, redeployment, and attribution confidence. Pilot reporting should begin with a small number of sectors rather than a universal mandate built on untested definitions.
48. Launch bounded human-capacity pilots. Public institutions should test tuition support, living and completion services, paid work-based learning, health-coverage bridges, and privacy-preserving AI assistance as separately measured modules.
49. Require site-specific data-center ledgers. Interconnection and incentive approvals should disclose electricity, water, grid, ratepayer, subsidy, job, and community-benefit terms in a common format.
50. Build layered governance. Agencies should use NIST's lifecycle process while legislation supplies rights, incident duties, independent authority, remedies, and enforcement. Voluntary frontier-safety policies should be audited as evidence of practice, not accepted as substitutes for public rules.
51. Test contribution design with retrospective data. Before setting a permanent rate, agencies should determine whether audited firm data can distinguish AI-attributable surplus, worker sharing, ordinary business cycles, and unrelated restructuring.
52. Scale only after review. Expansion should depend on published outcomes, independent evaluation, community and worker participation, and explicit stop conditions. Appendix A supplies a vendor-neutral protocol with preregistered measures and stop or redesign conditions.

This sequence is intentionally slower than announcing a national entitlement or automation tax. It is also faster than waiting for perfect certainty. It creates the data, pilots, and institutional relationships needed for a policy that can survive scrutiny.

13. Conclusion

The AI transition is not one labor-market forecast. It is a coupled transformation of work, learning, health security, physical infrastructure, and public authority. Its benefits may be substantial, but they will not distribute themselves automatically, and its costs will not remain inside the firms that create them unless policy makes those costs visible.

The evidence supports neither complacency nor technological panic. It supports a conditional compact: measure gains and burdens together; define AI labor in human-equivalent task-hours; preserve exposure, adoption, substitution, and elimination as separate concepts; finance learning and health continuity as distinct modules; hold education-to-work programs accountable for completion and persistence; allocate data-center burdens to the developments that create them; and combine voluntary practice with law and independent oversight.

The paper's practical contribution is a bounded transition framework: measure before claiming, preserve human authority before irreversible action, and require public reciprocity when concentrated private gains depend on shared labor, infrastructure, and institutions. The next step is not a sweeping promise. It is a measured pilot and a reporting standard capable of telling us when expansion has been earned.

AI-Use and Funding Disclosure

AI-use disclosure. Generative AI tools were used under the author's direction to support source discovery, research synthesis, calculation checking, drafting, and editorial revision. AI output was not treated as evidence. The author reviewed the cited sources and retains responsibility for the paper's claims, interpretations, and errors.

Funding disclosure. This version identifies no external funding source or sponsor. Any financial or in-kind support associated with a later pilot or commissioned evaluation should be disclosed with the funder's role in design, analysis, and publication.

Appendix A. Vendor-Neutral Bounded Pilot Protocol

Purpose. This protocol converts the paper's framework into a test that does not depend on a named vendor or proprietary platform. One public or nonprofit institution can run it with a bounded cohort, an ethically appropriate comparison condition, a preregistered analysis plan, and an independent evaluator. The pilot asks whether AI-assisted learning and transition services improve defined outcomes without shifting unacceptable labor, rights, fiscal, or infrastructure burdens.

Decision rule. Before recruitment, the sponsor should preregister the target sample size, expected attrition, power assumptions, minimum detectable effect for each primary outcome, and what an underpowered result may and may not support, together with outcome thresholds, data definitions, review dates, and stop conditions. A material rights-floor breach triggers immediate pause and review. Expansion requires independently verified improvement against the preregistered comparator, no unresolved rights-floor breach, acceptable cost and human overhead, and publication of adverse as well as favorable results. A pilot that fails should narrow, redesign, or stop.

Table A1. Vendor-neutral pilot measures and stop conditions

Component	Minimum design	Primary measures	Stop or redesign condition
Cohort and comparison	One public or nonprofit institution; a bounded cohort; a matched or pre/post comparator where ethical and feasible; preregistered access and attrition measures.	Participation, retention, demographic access, baseline balance, and support burden.	Unequal exclusion, attrition that invalidates comparison, unmanageable support burden, or inadequate consent.
AI-supported workflow	Instructor- or supervisor-mediated assistive use; no autonomous high-impact decisions; accepted-output threshold defined before launch.	Output quality, learning or task performance, human review time, and failure or escalation rate.	Worse outcomes than the comparator, excessive review burden, or unreviewed consequential output.
HETH and workload accounting	Record pre-deployment paid time and post-deployment execution, setup, oversight, rework, and exception time; distinguish saved time from redeployment.	E, A, U, Q, S, R; payroll, hours, output, quality, and attribution confidence.	Double-counted overhead, incompatible denominators, missing baseline evidence, or attribution too weak for a causal or fiscal claim.
Learning and career pathway	Test tuition, completion, work-based learning, or health-continuity modules separately; employers may advise but may not control admissions, curriculum, credentials, or learner data.	Completion, placement, earnings, six- and twelve-month persistence, and participant experience.	Sponsor capture, coercive placement, data repurposing, or no improvement over a less costly alternative.
Rights and data	Minimum-necessary collection, accessible design, revocable consent where feasible, human review, and a documented appeal path.	Privacy and security incidents, disparate outcomes, accessibility, contestation, and inappropriate automated consequences.	Material privacy, discrimination, accessibility, safety, or due-process failure.
Evidence and independent review	Versioned protocol, evaluator access, incident log, fixed decision dates, and public reporting of negative as well as positive results.	Protocol deviations, auditability, cost per improved outcome, reproducibility, and unresolved conflicts of interest.	Suppressed adverse findings, evaluator conflict, missing evidence, failed thresholds, or inability to reproduce the reported result.

Source: Author's proposed vendor-neutral protocol; proposed research requirements, not observed outcomes.

References

Acemoglu, D. (2025). The simple macroeconomics of AI. Economic Policy, 40(121), 13-58. https://doi.org/10.1093/epolic/eiae042 Source

Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of Economic Perspectives, 33(2), 3–30. https://doi.org/10.1257/jep.33.2.3 Source

Anthropic. (2026a, February 24). Responsible Scaling Policy, version 3.0. Source

Anthropic. (2026b). Responsible Scaling Policy, current version (v3.4 effective July 8, 2026; page updated August 14, 2026). Source

Bonney, K., Breaux, C., Dinlersoz, E., Foster, L., Haltiwanger, J., & Pande, A. (2026). The microstructure of AI diffusion: Evidence from firms, business functions, and worker tasks (CES Working Paper 26-25). U.S. Census Bureau. Source

Brollo, F., Dabla-Norris, E., de Mooij, R., Garcia-Macia, D., Hanappi, T., Liu, L., & Nguyen, A. D. M. (2024). Broadening the gains from generative AI: The role of fiscal policies (IMF Staff Discussion Note 2024/002). International Monetary Fund. https://doi.org/10.5089/9798400277177.006 Source

Centers for Medicare & Medicaid Services. (2026). National health expenditure data: NHE fact sheet, 2024 data. Source

College Board. (2025). Trends in college pricing and student aid 2025. Source

Congressional Budget Office. (2022, May 12). A single-payer health care system: Testimony before the U.S. Senate Committee on the Budget. Source

European Commission. (2026, August 3). AI Act: Regulatory framework and enforcement. Source

Georgia Department of Audits and Accounts. (2025). Data center sales and use tax exemption evaluation. Source

Gmyrek, P., Berg, J., Kaminski, K., Konopczynski, F., Ladna, A., Nafradi, B., Roslaniec, K., & Troszynski, M. (2025). Generative AI and jobs: A refined global index of occupational exposure (ILO Working Paper 140). International Labour Organization. https://doi.org/10.54394/HETP0387 Source

KFF. (2025). 2025 Employer Health Benefits Survey. Source

Morris, D. (2026). AI Transition Evidence Matrix (Version 0.6) [Companion research workbook]. Distributed with this working paper; persistent repository identifier pending.

National Center for Education Statistics. (2026). Undergraduate retention and graduation rates. Source

National Conference of State Legislatures. (2026). Subsidizing servers: How states are competing to attract data centers. Source

National Institute for Early Education Research. (2026). The state of preschool 2025. Source

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. Source

North American Electric Reliability Corporation. (2026). 2025 Long-Term Reliability Assessment. Source

Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J. G., Masanet, E. R., & Sartor, D. A. (2024). 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory. https://doi.org/10.71468/P1WC7Q Source

State Higher Education Executive Officers Association. (2026). State Higher Education Finance FY2025. Source

Tucker, L. C. (2026). You're (not) hired: Artificial intelligence and early career hiring in the Quarterly Workforce Indicators (CES Working Paper 26-27). U.S. Census Bureau. Source

U.S. Bureau of Economic Analysis & U.S. Bureau of Labor Statistics. (2026). Integrated Industry-Level Production Account, 1997-2024. Source

U.S. Bureau of Labor Statistics. (2026a). The Employment Situation: July 2026, tables B-1, B-2, and B-3. Source

U.S. Bureau of Labor Statistics. (2026b). Education pays, 2025. Source

U.S. Bureau of Labor Statistics. (2026c, August 28). Current Employment Statistics - CES (National): March 2026 preliminary benchmark revisions by major industry sector. Source

U.S. Bureau of Labor Statistics. (2026d, August 28). Technical notes for preliminary Current Employment Statistics national benchmark revisions. Source

U.S. Census Bureau. (2025). Health insurance coverage in the United States: 2024 (Current Population Reports P60-288). Source

U.S. Census Bureau. (2026a). AI Supplement Table 2026, Business Trends and Outlook Survey. Source

U.S. Census Bureau. (2026b). School system finances: Fiscal year 2024. Source

U.S. Census Bureau. (2026c). Annual estimates of the resident population by single year of age and sex for the United States: April 1, 2020 to July 1, 2025 (NC-EST2025-SYASEXN). Source

U.S. Census Bureau. (2026d). Quarterly Financial Report, first quarter 2026. Source

U.S. Department of Labor. (n.d.). Apprenticeship.gov. Source

U.S. Department of Labor. (2022). Do employers earn positive returns to investments in apprenticeship? Evidence from a survey of American Apprenticeship Initiative employers. Source

U.S. Department of Labor, Office of Inspector General. (2026). Scaling Apprenticeship Through Sector-Based Strategies grants audit (Report 02-26-001-03-375). Source

Williams, S., & Freund, J. (2026, March 5; updated March 17). Anthropic's RSP v3.0: How it works, what's changed, and some reflections [Research commentary]. Centre for the Governance of AI. Source