Compute as Public Capital

A Framework for Universal Economic Participation in an Automated Economy

EM Foundation for AI Research · Working Paper · Draft v1.1 (Final Publication Audit)
Classification: EM Foundation Working Paper — not yet White Paper or Academic Submission Candidate

Executive Abstract

Problem. As AI and robotics converge on displacing both cognitive and physical tasks at once, the historical safety valve of lateral retraining narrows faster than new categories of durably human work open. The deeper issue is not simply job loss: automation can raise aggregate productive output while weakening the wage-based mechanism through which purchasing power, and institutional revenue, have historically been distributed.

Mechanism. The industrial-era distributive circuit (labor → wages → purchasing power → consumption → corporate revenue → investment → employment) has no equivalent in an increasingly automated circuit (capital + energy + compute → automated production → corporate revenue → capital returns → more automation). US labor share of income fell from 67.8% in 1987 to 58.1% in 2023 (BLS/FRED series MPU4910141), a global pattern documented since the early 1980s (Karabarbounis & Neiman, 2014). This paper extends the household-level distribution problem to an institutional one: municipalities, Social Security, and private insurance are financed against payroll, wage income, and human-operated risk — proxies that can decouple from where productive value is actually generated.

Proposed architecture. A three-layer response — Participation Floor (redistribution), Automation Dividend (predistribution/ownership), and Universal Compute Income (productive capacity) — summarized as survive → own → create.

Major evidence. Established: the labor-share decline (BLS; Karabarbounis & Neiman 2014); the displacement/reinstatement framework and automation's estimated 50–70% contribution to US wage-structure change since 1980 (Acemoglu & Restrepo, 2019; Acemoglu & Restrepo, 2022); Social Security's documented 75-year actuarial deficit (CBO, 2023, Publication No. 59340). Evidence-supported but complicating: data centers currently generate substantial local property-tax revenue in hosting jurisdictions (Loudoun County VA, Quincy WA), and a 2026 study finds robot exposure in manufacturing lowers, not raises, SSDI applications without reducing local employment — both preserved here as genuine complications rather than omitted.

Unresolved fiscal question. This paper contains no completed comparative costing of its own proposal. The correct comparison is a Managed Transition (the three-layer architecture) against an Unmanaged Transition (continued erosion of payroll-linked

institutional revenue under the status quo) — not the proposal's cost against zero. That comparative model has not yet been built (Appendix C).

Classification. EM Foundation Working Paper. Not yet a White Paper or Academic Submission Candidate — the fiscal-costing gap is structural and not resolvable through further text revision.

I. Introduction — Production Without Distribution

What happens when a society becomes increasingly capable of producing what people need while progressively removing people from the mechanism through which they obtain a claim on that production? This is the question this paper is organized around, and it is analytically distinct from the more familiar question of unemployment.

Unemployment asks whether a displaced individual can find new work. Distribution asks whether the economy retains a broad-based mechanism for converting rising productive output into purchasing power and institutional revenue, independent of whether any specific individual is employed. A society can solve the first problem for a given worker — through retraining, a new job, a different sector — while the second problem continues to widen at the aggregate level, because the wage channel that used to perform that distributive function carries a shrinking share of total output. This paper's central and most important distinction is that these are related but separate problems, and a policy response built only for the first (job placement, retraining) will not resolve the second.

The remainder of this paper proceeds in three movements: Sections II–VII establish the empirical and structural case that this distribution break is real, is measurable at least in part, and propagates beyond households into the institutions built around wage labor as their financing proxy. Sections VIII–XII develop the proposed three-layer policy architecture and its philosophical justification. Sections XIII–XVII address what remains unresolved, what would falsify the paper's central hypotheses, and what a defensible policy pathway looks like given that the fiscal question is not yet answered.

II. The Displacement / Reinstatement Problem

For most of industrial history, automation displaced tasks laterally: a machine displaces a task, and workers retool into an adjacent task the machine cannot yet perform. This was often brutal for the generation living through it, but it worked because the automating technology was narrow — a machine that could weave cloth could not also drive a truck or keep a ledger.

That lateral escape route is closing. The same underlying AI systems now displacing cognitive tasks are, in parallel, being embodied in robotics capable of physical labor — closing adjacent exits at the same time as the original one. Consider an ordinary case:

a homeowner who pays a local contractor to mow the lawn — physical, local, relationship-based work, exactly the kind displaced workers have historically been told to move toward. A robotics-and-logistics company dispatching an autonomous mower to a thousand yards a day undercuts that worker's entire category, not one client, and the standard policy answer — retrain, move up the value chain — runs into the same wall one step later.

This claim should be stated against, not around, the dominant framework in labor economics for this exact question. Acemoglu and Restrepo's displacement/reinstatement model (Journal of Economic Perspectives, 2019) treats automation's labor-market effect as a race between a displacement effect (automation removing tasks from labor) and a reinstatement effect (new tasks in which labor retains a comparative advantage). Historically, reinstatement has offset displacement. A separate, related empirical paper by the same authors (Econometrica, 2022) quantifies the resulting distributional effect directly: automation accounts for an estimated 50 to 70 percent of the changes in the US wage structure over the four decades to 2016, working through the relative wage declines of worker groups specialized in routine, automation-exposed tasks. Recent work extending the displacement/reinstatement framework to agentic AI systems — capable of completing entire occupational workflows rather than discrete tasks — argues this expands displacement risk beyond what task-level analysis previously captured.

This paper's claim is therefore a specific, citable prediction within the labor-economics mainstream, not a departure from it: the reinstatement effect this wave generates will continue to lag displacement by a wider margin than in prior automation waves, because the same general-purpose capability closing off one lateral option simultaneously closes the adjacent ones reinstatement would normally open.

This is not a claim that automation is total or immediate. It is a claim about relative rates — the rate at which lateral retraining closes now exceeds the rate at which new, durably human categories of work open — and Sections III–VII exist to establish why that gap matters beyond the individual worker's experience of it.

> *Caveat: reinstatement is a real, historically dominant channel. This paper's prediction — that it lags further behind in this wave — is a testable claim addressed further in Section XV, not an established fact.*

III. The Distribution Break

For most of the industrial era, the distributive circuit has run in one direction:

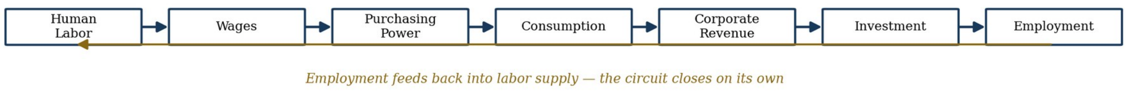

Figure 1. The industrial distribution circuit. Employment feeds back into labor supply, closing the loop.

A highly automated economy increasingly runs a different circuit:

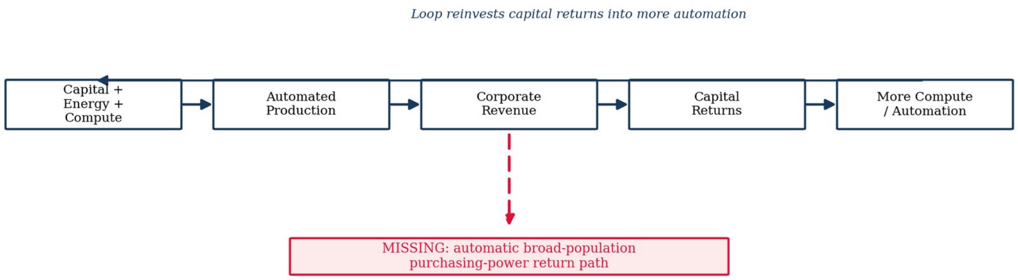

Figure 2. The automated production circuit. No arrow in this circuit automatically returns purchasing power to the broader population.

This is not speculative. US labor share of income fell from 67.8 percent in 1987 to 58.1 percent in 2023 (BLS/FRED series MPU4910141, Private Nonfarm Business Sector: Labor Share), and Karabarbounis and Neiman (2014) document this decline as global, occurring across most countries and industries since the early 1980s, driven substantially by technological substitution of capital for labor.

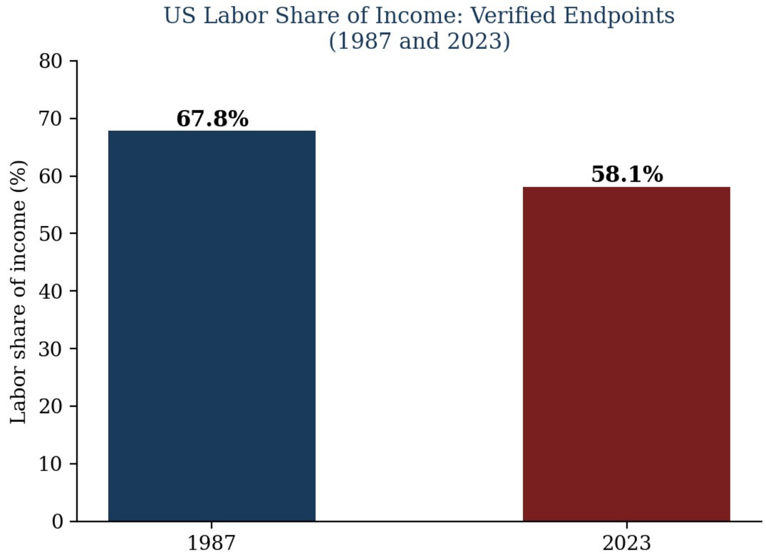

Data Chart 1. Verified endpoints only — the full annual BLS series was not independently re-verified for this draft and is not plotted.

One caveat worth stating directly: capital income (R_t in the notation below) is not received exclusively by a small ownership elite — a meaningful share flows to ordinary households through pensions and index funds. The labor-versus-capital framing below is a simplification; the real distributional question is how concentrated versus broadly held capital ownership is, which is closer to Section V's recursive-concentration dynamic than to a clean two-population split.

A minimal formalism (conceptual notation, not a calibrated model)

The following is offered strictly as illustrative notation clarifying the logic of the argument. No term below has been decomposed into anything measurable, and no numerical prediction is derived from it — doing so is the subject of Section XV's research agenda, not a result claimed here.

Let $Y_t = W_t + R_t$, where W_t is labor income and R_t is capital/ownership income in period t. If automation drives labor's share of income down ($W_t/Y_t \downarrow$) while ownership of R_t stays concentrated, the relevant policy question becomes the distribution of effective purchasing claims across households. Introducing a distribution term D_t (dividends, transfers, public services), household consumption capacity is bounded roughly by:

$$\mathbf{C_t \leq W_t + D_t + R_{h,t} + B_t}$$

where $R_{h,t}$ is broadly held household capital income and B_t is borrowing. The insight this makes explicit: debt can temporarily substitute for missing wage distribution, but it cannot permanently repair the circulation problem — B_t can smooth a shortfall for a period, but it accumulates a claim against future C_t rather than resolving the structural gap between W_t and Y_t.

IV. The Two-Sided Municipal Ledger

An earlier version of this argument treated municipal impact as one-directional: automation displaces wages, displaced wages reduce local spending, reduced local spending erodes municipal tax bases. That direction is real, but the evidence is more complicated than a single-direction story, and the complication is preserved here rather than smoothed over.

Data centers — the physical infrastructure of the automated economy — currently generate substantial, sometimes dominant, local property-tax revenue in the jurisdictions that host them. Data centers reportedly account for roughly 38 percent of Loudoun County, Virginia's total revenue and about 75 percent of Quincy, Washington's property-tax levy. This directly complicates any simple claim that automated capital does not pay into the local tax base — in the near term, capital-intensive automated infrastructure can be a substantial local fiscal benefit, not a drain.

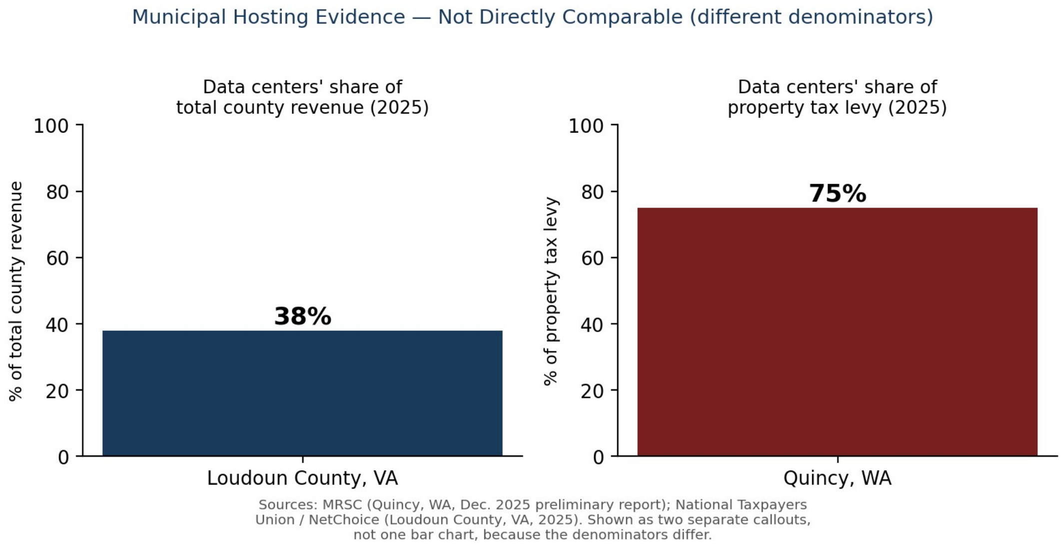

Data Chart 2. Shown as two separate callouts rather than one bar chart because the underlying denominators differ.

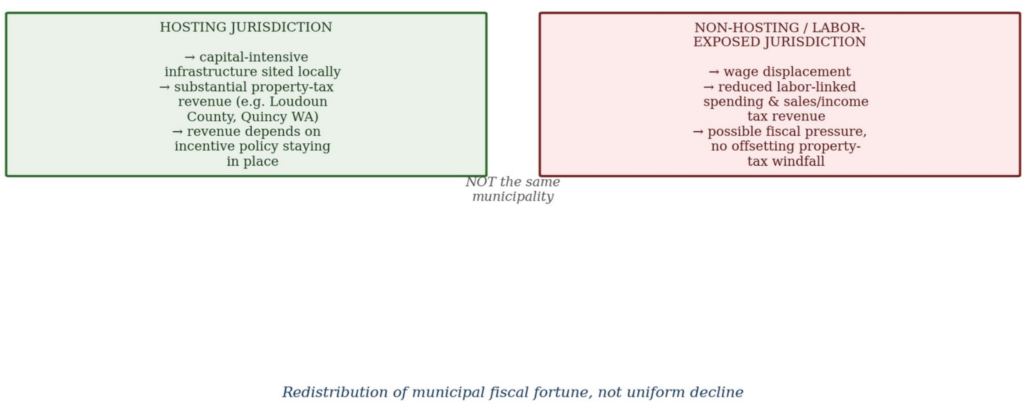

Figure 5. The two-sided municipal ledger.

The more precise claim, once this is accounted for, is two-sided rather than one-sided: the erosion channel is specifically about labor-linked local revenue (sales tax, local income tax tied to displaced service-sector wage spending), not about automated capital as such, which can generate large property-tax windfalls independent of local employment. That windfall is neither guaranteed nor stable — it depends on tax-incentive policy currently being rolled back in several states as costs become visible: Georgia's data-center tax abatements are estimated to cost localities $1.1–1.4 billion in 2026 alone, and Texas's cumulative exemption cost is projected near $9 billion through 2030.

The honest formulation of the municipal-cascade hypothesis is therefore: automated economic activity can simultaneously strengthen the property-tax base of the specific jurisdictions that host its physical infrastructure while eroding the labor-linked tax base of the much larger number of jurisdictions that do not — a redistribution of municipal fiscal fortune, not a uniform municipal fiscal decline.

This is a testable, falsifiable claim, with two unresolved complications named directly rather than hidden. First, a selection-effect confound: hosting jurisdictions are not randomly selected — they are chosen partly for favorable tax and regulatory conditions, which a clean before/after comparison does not control for. Second, an unexamined question: most US states already operate some form of state-level fiscal transfer mechanism between jurisdictions, and this paper has not examined whether existing transfer systems already partially offset the geographic mismatch described here. Both are named as open items for the research agenda in Section XV, not resolved in this paper.

V. Recursive Capital Concentration

A further dynamic compounds the distribution break, and explains why waiting until displacement becomes severe before acting is materially more dangerous than for a conventional, slower-moving economic shift.

Unlike human labor, whose productive capacity is tightly constrained by biological time and individual throughput, automated productive capital can be replicated, scaled, and financed from its own returns across a much wider range before physical, competitive, or economic constraints impose diminishing returns. This is not a claim that automated capital faces no diminishing returns — it plainly does, through competition, energy constraints, depreciation, and technological obsolescence. The distinction that matters is capacity for replication and scalable reinvestment, not immunity from constraint.

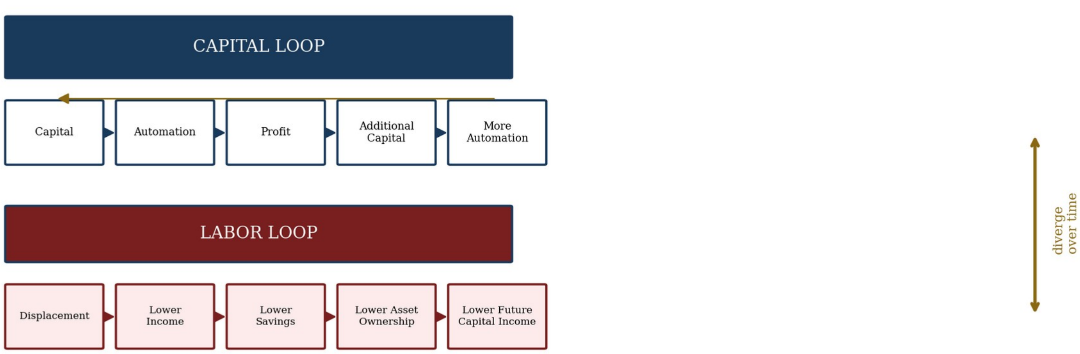

Figure 3. The capital loop and labor loop diverge over time; the diagram illustrates direction, not a calibrated rate.

These two loops actively diverge, because the population losing income is simultaneously losing the capacity to acquire the asset class compounding fastest. This is the strongest argument in the paper for urgency over gradualism: a policy response that waits until concentration becomes severe is waiting for the harder loop to entrench further while the easier stage to interrupt has already passed.

VI. Institutional Revenue Decoupling

Sections III and V describe the household-level distribution break and its recursive concentration dynamic. This section formalizes a related but distinct claim: the same break propagates into institutions built around labor as their financing proxy, independent of whether any individual household is protected.

Working definition

Institutional Revenue Decoupling occurs when automation changes the relationship between aggregate productive capacity and the legacy economic proxies (payroll, wage income, employment counts, human-operated risk) used to finance institutions, causing institutional revenue or risk-pool structure to stagnate, decline, or destabilize even while underlying productive output rises.

Tested against the existing architecture: this is not a fully independent phenomenon from the Distribution Break — it is the same underlying mechanism viewed from the institution's side of the ledger rather than the household's. The term is retained because the policy implications differ (a household Participation Floor does not, by itself, repair a municipal tax base or a payroll-financed social-insurance system), not because it names a new causal mechanism.

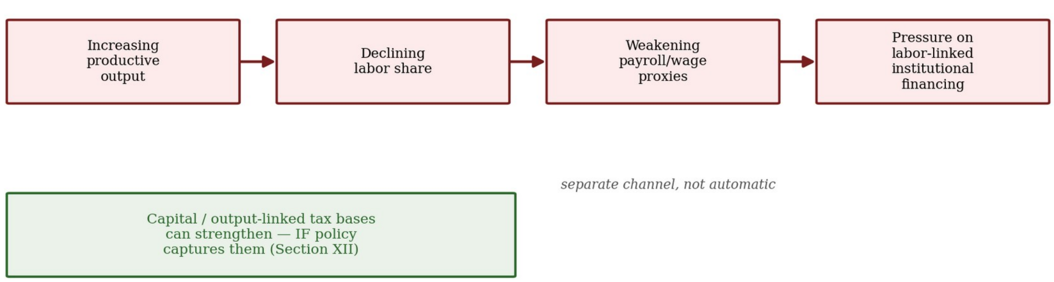

Figure 4. Institutional Revenue Decoupling — the labor-linked channel and the capital/output-linked channel are separate, and the second does not strengthen automatically.

Extended formalism (hypothesis, not established theorem)

Let P_t be aggregate productive output, $T_{l,t}$ revenue from labor-linked taxation, $T_{c,t}$ revenue from capital/output-linked taxation, and G_t institutional/public-service obligations. If P_t rises while W_t/P_t falls, then under a tax architecture heavily dependent on W_t: $T_{l,t}/P_t$ falls.

The hypothesis this motivates: a state can experience increasing productive capacity alongside deteriorating fiscal capacity if its tax architecture remains attached to a shrinking factor of production — a distinct claim from "the economy is doing worse," and arguably a more precise diagnosis of what rising GDP alongside strained public budgets would actually look like.

A US-specific complication, precisely sourced rather than asserted generally: the United States is the only OECD member country without a value-added tax — a consumption-based, non-payroll-linked base structurally closer to $T_{c,t}$ than $T_{l,t}$ (OECD, Consumption Tax Trends 2024). As a direct consequence, the US collects roughly 16.8 percent of total government revenue from consumption taxes, compared with an OECD average of about 31.1 percent (Tax Foundation, compiling OECD data). This means the tax-base transition this section describes is plausibly harder for the US specifically than OECD-comparative literature on automation and tax policy implies: a US shift toward non-payroll-linked revenue would need to build collection infrastructure most OECD peers already have, not merely reweight existing collection toward it.

VII. Institutional Case Studies

VII.A — Social Security and Social Insurance

Social Security is financed pay-as-you-go from payroll taxes. Its long-run funding pressure is an established, independently documented fact, not a projection this paper introduces: CBO's long-term projections for Social Security (2023) show the Old-Age and Survivors Insurance trust fund exhausted in fiscal year 2032 and the Disability Insurance trust fund exhausted in 2052, with a 75-year actuarial deficit equal to roughly 1.7 percent of GDP, or about 5.1 percent of taxable payroll. This gap is currently attributed primarily to demographic shift — a declining worker-to-beneficiary ratio and slower labor-force growth — not automation. This paper's contribution is a hypothesis, not yet established: that a declining labor share of production compounds the same funding gap through a second, independent channel, since even a stable worker-to-beneficiary ratio produces less payroll-tax revenue per unit of output if output increasingly comes from capital rather than payroll.

This hypothesis should not be overstated, and a genuine complication in the evidence is preserved here rather than smoothed over. A 2026 study using commuting-zone data and an instrumented measure of robot exposure in manufacturing found that each additional robot per 1,000 workers is associated with lower, not higher, SSDI application inflows in exposed labor markets, and did not find declining employment-to-population ratios in those same markets — which weighs against broad local displacement as the primary channel, at least for that specific, narrower exposure measure. This finding is scoped to robot-intensive manufacturing labor markets, not the broader, general-purpose AI/robotics wave this paper's Section II argues is structurally different from prior automation. Whether that broader wave produces the same result is an open empirical question this paper does not resolve.

VII.B — Insurance

Traditional insurance underwriting depends on pooling large numbers of partially independent human risks — workers' compensation, commercial and personal auto, professional liability, disability, life insurance tied to future labor-income replacement — where diversification works because one person's accident is statistically unrelated to another's.

Automation does not eliminate insurable risk; it changes its topology. As human operation gives way to shared AI models, cloud platforms, and robotics fleets, previously independent risks can become correlated: an autonomous-driving platform defect, a common model failure, or a foundation-model concentration event can generate losses across many previously independent policyholders simultaneously. Correlated catastrophic loss is not itself new to the insurance profession — hurricane and earthquake catastrophe risk, and the 1980s liability crisis, are established correlated-loss categories already priced through reinsurance and catastrophe bonds. The genuinely novel problem is narrower: the speed and opacity of AI-model-driven correlation, where a single model update can correlate losses across an entire book of business within days rather than over a hurricane season's timescale — a pace existing catastrophe-modeling infrastructure was not built for.

This is the active subject of a fast-emerging 2026 actuarial literature. Work on the insurance of agentic AI identifies foundation-model concentration as a channel through which upstream model failure can correlate losses across many insured parties at once, and proposes quota-share and aggregate-reinsurance structures for early-stage AI risk portfolios rather than standalone capacity the data does not yet support. Parallel work argues heavy-tailed, correlated AI-driven losses will likely require a layered structure — a private layer for ordinary, poolable agent-caused harms, and a public or pooled catastrophic layer, echoing existing flood, terrorism, and cyber-catastrophe precedent. A further constraint worth naming directly: global reinsurance capital is finite and already under pressure from climate-driven catastrophe losses, so a new, poorly understood, potentially correlated risk category competes for the same limited capacity rather than arriving into room to spare.

This section does not claim insurance disappears. It claims the shift from many-distributed-human-risks toward fewer-but-more-correlated-technological-risks is real and actively being underwritten, and that public-private catastrophic backstop mechanisms — already precedented in flood and terrorism insurance — are a plausible, not speculative, part of how this resolves.

VII.C — Bounded Institutional Scan

A brief, deliberately non-exhaustive scan, limited to institutions that materially strengthen the paper's central thesis: mortgage underwriting and consumer credit scoring are built substantially on verified wage income, and this is closer to a systemic banking-sector credit-risk question than an individual-underwriting one — bank balance sheets are exposed both to commercial real estate and to consumer credit underwritten against an assumption of stable wage income across the loan book as a whole. Higher-education financing, particularly student lending, is priced against anticipated future labor-market returns that Section II's displacement argument directly calls into question for a widening set of fields. Commercial real estate valuation in central business districts depends on continued commuter office employment. Employer-sponsored health insurance functions, in the US specifically, as a major distribution mechanism for insurance access contingent on continued employment. Private and occupational pensions face a structurally distinct but related problem from Social Security: defined-benefit funding assumes continued contributions from an active workforce, and defined-contribution plans assume continued wage income to contribute from in the first place — both fail under the same labor-share-decline mechanism formalized in Section VI, independent of any specific payroll-tax structure.

VIII. Education, Agency, and Computational Citizenship

One of the genuinely novel properties of increasingly capable AI systems is that they lower the expertise threshold required to produce sophisticated work — AI access partially compensates for education's absence rather than merely rewarding its presence. The relationship is better modeled as multiplicative than as a hard gate:

$$\textbf{Education} \times \textbf{Compute} \times \textbf{Agency} \rightarrow \textbf{Innovation}$$

Education still matters enormously to the ceiling of what a population can do with broadened compute access. A separate, harder claim the paper cannot responsibly omit: there is a documented tension between political power that depends on an under-informed electorate and public investment in education that produces a well-informed one. This is presented as an incentive-structure analysis, not an accusation aimed at any specific movement, and it concerns whether a population can evaluate the compute-allocation and automation-dividend policy itself — a distinct risk from whether people can use AI tools effectively.

What this paper proposes, once tokens are understood as one implementation detail rather than the substance of the claim, is a form of computational citizenship:

Computational Citizenship Principle

In an economy where computational intelligence becomes a primary means of production, meaningful access to computational agency becomes a prerequisite of full economic citizenship.

This does not claim a natural right to unlimited compute. It claims economic citizenship increasingly requires access to the productive infrastructure of the era — analogous to literacy becoming necessary for participation in an industrial democracy, and internet access for participation in modern commerce.

IX. Compute as Public Capital

AI/AGI compute should be understood the way land was understood in an agrarian economy and capital in an industrial one: as the resource whose distribution determines who can meaningfully compete, not merely who can consume. A precise distinction matters here: AI capability is diffusing rapidly (open-weight models, falling inference costs), while physical compute — GPUs, fabrication, memory bandwidth, electricity — can remain genuinely scarce. The defensible claim is narrower than “compute scarcity is unsustainable” outright: artificial scarcity of access to useful computational intelligence becomes increasingly difficult to sustain as capability diffuses, even while physical compute and energy remain scarce resources requiring deliberate allocation policy.

Delivery should distinguish two very different undertakings that are easy to conflate: operating inference capacity on existing open-weight models is a tractable, public-utility-like undertaking; training and maintaining frontier-capability AGI systems requires specialized research teams and ongoing capability investment that a public-utility operating model does not straightforwardly provide. This paper's proposal is scoped to the former — it does not argue government should operate frontier AGI research laboratories.

X. Energy and the Infrastructure Chain

A universal compute allocation is only as real as the electricity required to use it. No token allocation reaches a household that cannot reliably power the device redeeming it. But energy is only the first link in a longer dependency chain:

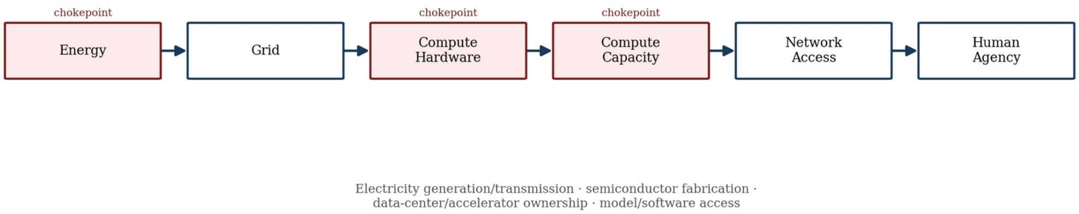

Figure 6. The infrastructure dependency chain. Electricity generation/transmission, semiconductor fabrication, data-center/accelerator ownership, and model/software access are each independent concentration points.

Universal token allocation addresses model/software access directly and electricity indirectly; it does nothing on its own about semiconductor fabrication or hardware ownership, both concentrated among a small number of firms with real geopolitical exposure. This paper does not propose national semiconductor fabrication as a remedy, but states plainly that a UCI program built entirely on token allocation, without acknowledging this upstream dependency, is vulnerable to a fair objection: universal access to a resource whose underlying hardware supply remains concentrated is narrower than “universal access” implies.

The comparison worth making, carefully, is to China's anticipatory, state-directed energy buildout versus reactive, market-driven grids elsewhere. On the anticipatory side, precisely sourced: China's operating nuclear capacity reached approximately 62 GWe by the end of 2025, with the 15th Five-Year Plan (approved March 2026) targeting 110 GWe by 2030 — near-doubling on a fixed planning horizon (World Nuclear Association); the State Grid Corporation of China invested roughly CNY 650 billion (about USD 90 billion) on grid infrastructure in 2025 alone and plans approximately CNY 4 trillion over 2026–2030 (World Nuclear Association, citing SGCC). On the reactive side, illustrated by Northern Virginia's data-center interconnection backlog: only about 13 percent of capacity that entered US interconnection queues between 2000 and 2019 had reached commercial operation by 2024, as reported by the Searchlight Institute citing Lawrence Berkeley National Laboratory data — this paper cites the secondary report rather than the primary LBNL dataset directly and that distinction is preserved here rather than presented as a primary-source citation. This comparison is offered strictly as a planning-model contrast, read alongside China's own well-documented history of state-directed malinvestment — overcapacity and stranded assets in sectors like solar manufacturing and real estate are the standard counter-examples — since citing anticipatory buildout without that counter-evidence would understate the real tradeoff.

XI. The Basis of the Public Claim

Why should someone who did not build, finance, or take capital risk on the automation have any claim on its returns? Private technological production draws on accumulated public and social capital: public education, publicly funded research, infrastructure, legal and property-enforcement systems, currency stability, and generations of accumulated public knowledge. This does not deny private owners their legitimate return on real capital risk. It means private returns and a public participation claim are not mutually exclusive — analogous to mineral royalties, where an extraction company's legitimate reward for discovery and extraction coexists with society's retained claim on the underlying commons.

The Automation Dividend is therefore best understood not as rich-to-poor redistribution, but as a return from the automated productive system to the constituent members of the civilization that supported its emergence:

Private Return + Public Participation Claim + Universal Productive Access

XII. The Three-Layer Architecture

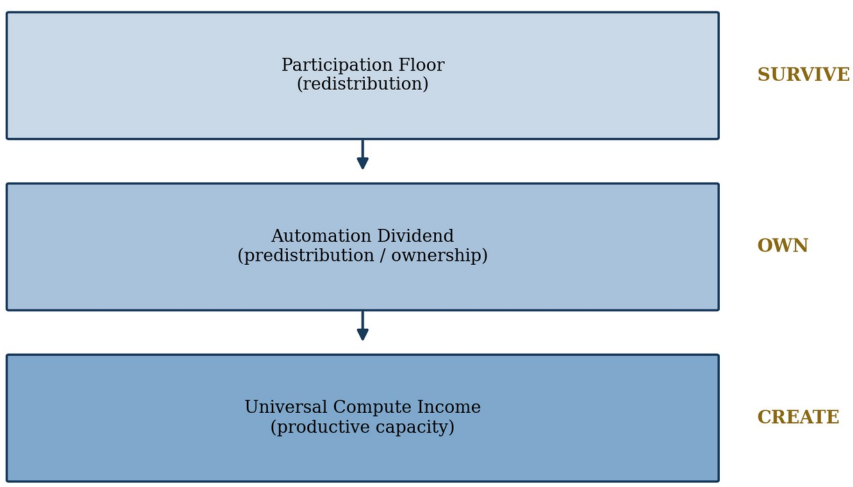

Figure 7. The three-layer architecture: redistribution, predistribution/ownership, and productive capacity are distinct kinds of intervention.

Layer 1 — Participation Floor (redistribution)

Not locked to a single implementation — Universal Basic Income, universal basic services, or a negative income tax could each satisfy this layer's function. It addresses survival, and survival alone.

Layer 2 — Automation Dividend (predistribution / ownership)

Two implementation families deserve direct comparison. An automation-tax model (illustratively, $SD_t = \tau_a \cdot \Delta\Pi_a$, where $\Delta\Pi_a$ is estimated automation-attributable productivity gain) faces a real, likely unsolvable-at-precision measurement problem: firms have strong incentive to reclassify automation gains as “process improvement” or “restructuring.” A citizen/sovereign ownership model — a fund holding diversified productive assets directly, with citizens participating through fund distributions — largely sidesteps this attribution problem, at the cost of requiring the fund to actually acquire meaningful equity positions, a distinct and nontrivial bootstrapping problem.

Section VI's institutional-decoupling analysis strengthens the case for the ownership model specifically: if a sovereign or citizen fund holds diversified claims on productive capital, its revenue follows productive value as that value migrates from labor toward automated capital, rather than remaining pinned to a shrinking payroll base. This reframes the underlying policy question: should public revenue architecture eventually follow value creation directly, rather than the industrial-era proxies (payroll, wages, employment counts) through which value historically appeared?

Layer 3 — Universal Compute Income (productive capacity)

Distributes access to the means of producing new value directly — a monthly allocation of AI/AGI compute tokens, free at the point of use. An illustrative figure (tens of millions of tokens per citizen per month) has not been checked against Section X's energy/hardware constraints; no conversion from token count to kWh, FLOPs, or dollar cost is offered here, and that conversion is a prerequisite for treating the figure as more than a placeholder.

SURVIVE → OWN → CREATE

Ownership default rule: public infrastructure enables private creation; creators retain ordinary ownership rights over what they produce, unless they voluntarily contribute the work to a commons or accept different terms under a specific public program — the same logic as public roads and public schools.

XIII. Fiscal Question — Managed vs. Unmanaged Transition

This paper contains no fiscal costing anywhere in the document. This is the paper's most serious unresolved gap, and it is stated directly rather than hedged: every layer of the three-layer model defers costing to future modeling, and cumulatively that means a national-scale proposal is currently uncosted in its entirety. This is not fixable through further prose revision — it requires a dedicated companion costing exercise, specified but not performed in Appendix C.

The relevant comparison, once Sections VI and VII are accounted for, is not Cost(Proposed Architecture) versus $0 — the status quo under continued automation is itself fiscally dynamic.

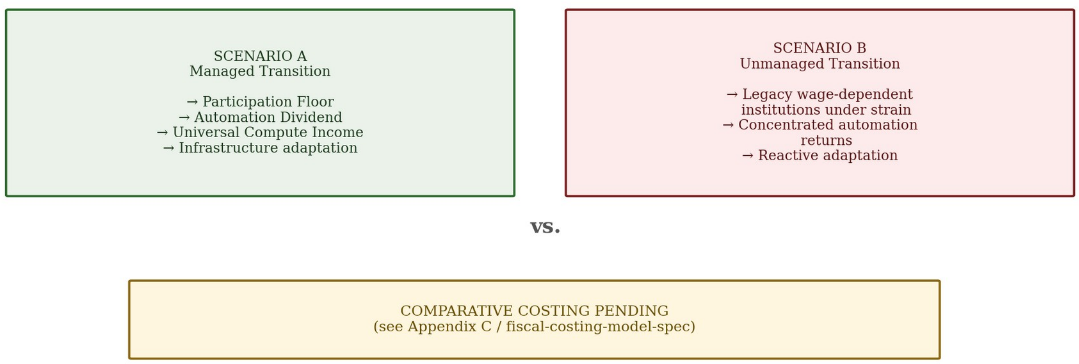

Figure 8. The correct comparison is Scenario A against Scenario B — not the proposal's cost against zero. Comparative costing is pending; no totals are presented here.

Scenario A — Managed Transition — includes Participation Floor, Automation Dividend, Universal Compute Income, and infrastructure adaptation costs. Scenario B — Unmanaged Transition — includes projected payroll-linked fiscal deterioration (Social Security's existing, independently documented shortfall, Section VII.A), municipal fiscal bifurcation (Section IV), displacement-support costs under current-law systems, insurance-market restructuring costs (Section VII.B), and reactive rather than anticipatory infrastructure costs.

The governing principle: the fiscal question is not simply what adaptation costs. It is what adaptation costs relative to non-adaptation under the same automation trajectory. The analysis must be capable of finding that Scenario B outperforms Scenario A if the evidence says so — a model incapable of that result is not a costing exercise, it is advocacy wearing a model's clothing. No Scenario A/B totals, token costs, energy requirements, or sovereign-fund return figures are presented in this paper. All such figures would be fabricated if stated here, and are deliberately withheld pending the companion model specified in Appendix C.

XIV. Counterarguments

Objection	Response / Status
Removes incentive to innovate privately / undermines IP rights	The proposal is a universal floor, not elimination of private markets or IP rights above it (Section XII ownership default; Section XI coexistence argument).
Free compute access at this scale is technically/economically unworkable	Conceded as a legitimate scaling concern. Depends on continuing inference-cost declines and Section X's infrastructure buildout; not costed as workable today.
Token allocations are gameable	Extraction profitability declines as the resource becomes abundant rather than scarce; reduces but does not eliminate the incentive. Allocation-integrity mechanisms remain necessary.
Why does anyone without capital risk deserve a claim at all?	Addressed in Section XI: private returns and a public participation claim on shared civilizational substrate are not mutually exclusive.
Automation-attribution/measurement problem is unsolved	Conceded — the central reason the ownership model deserves weight over the tax model (Section XII).
National security constraints on broad AGI/ASI access	Conceded. Proposal is for a baseline capability tier; capability-sensitive systems remain subject to standard security review.
Energy/hardware buildout demands planning authority the system wasn't built for	The least tidy objection. Alternative (market-allocated siting where capital already concentrates) reproduces the same inequality one layer beneath it.
Free-at-point-of-use resource invites overconsumption absent a price signal	Real, unresolved tension. A floor-plus-priced-market-above structure mitigates but does not eliminate the distortion; usage caps are non-price rationing with their own inefficiencies.
Capital flight / international competitiveness under unilateral adoption	Named as an open question, raised independently by two reviewer perspectives; not resolved. Likely engages how ownership framing differs from a pure tax in its effect on capital location.
Sovereign-fund acquisition/bootstrapping cost	Conceded as a distinct, nontrivial problem from the attribution problem the ownership model otherwise avoids (Section XII).
Municipal selection effects / existing state	Conceded (Section IV) — hosting-

fiscal transfers	jurisdiction selection is non-random, and existing transfer mechanisms have not been examined for offsetting effect.
Tax-base alternatives such as VAT; existing insurance/reinsurance mechanisms	Addressed directly in Sections VI and VII.B rather than treated as novel problems this paper alone identifies.

XV. Research Agenda and Falsifiability

A theory that cannot lose is not a useful research program. This section separates what can be tested with existing data from what requires new methodology or data access.

Near-term testable (existing public data)

- Municipal two-sided ledger: compare property-tax and labor-linked-tax trajectories between hosting and non-hosting jurisdictions with comparable pre-2020 demographics (Census of Governments, state comptroller data), controlling for the hosting-jurisdiction selection effect named in Section IV.
- Labor-share / municipal-revenue relationship: BLS regional labor-share data cross-referenced against Census of Governments local tax-receipt data, independent of the data-center effect.

Medium-term (requires new exposure measure or data access)

- Extend the SSDI/robot-exposure study's method to general-purpose AI occupational exposure, using a task-exposure instrument analogous to existing labor-economics methodology, to test whether the broader displacement claim in Section II shows the same pattern as the narrower manufacturing-robot finding in Section VII.A, or a different one.
- Social Security actuarial-deficit decomposition: separate the documented 75-year deficit into its demographic component (already modeled by SSA) versus a labor-share component, the most direct test of Section VI's central fiscal hypothesis.

Longer-term (multi-year research programs)

- Formal calibration of the Section III/VI models: real elasticity estimates (capital-labor substitution, by sector) and a decomposition method separating automation's specific contribution from ordinary capital deepening — the same decomposition problem that motivates preferring the ownership model over the tax model in Section XII.
- Insurance correlated-loss quantification: testing whether AI-model-update-driven loss correlation behaves faster and more opaquely than existing catastrophe categories, as Section VII.B claims, requires industry claims data not likely publicly available in the near term.

What would weaken or falsify this paper's central hypotheses

A theory that cannot lose is not a useful research program. The following findings, if established, would specifically weaken important portions of this paper's argument — they are stated as conditions, not hedges:

- If broad-based AI exposure (not just manufacturing robotics) produces a reinstatement effect comparable in magnitude to its displacement effect, Section II's central claim — that this wave's reinstatement lags displacement by a wider margin than prior waves — would be substantially weakened.
- If declining labor share is found not to materially weaken labor-linked institutional revenue once real fiscal data is examined, Section VI's Institutional Revenue Decoupling hypothesis would be weakened at its foundation.
- If existing municipal and state fiscal-transfer mechanisms are found to already substantially compensate for the hosting/non-hosting divergence described in Section IV, the two-sided municipal ledger's policy relevance — though not its underlying empirical accuracy — would be substantially reduced.
- If capital ownership becomes sufficiently broad-based (through pensions, index funds, or other diffusion) that a declining wage share ceases to meaningfully reduce aggregate household purchasing claims, Section III's distribution-break argument would be weakened, and the caveat already noted there about broadly held capital income would become the paper's central finding rather than a caveat.
- If public compute access is piloted and produces little measurable innovation or productive economic participation, the Computational Citizenship claim in Section VIII and the case for UCI specifically in Section XII would be substantially undermined, independent of whether the Participation Floor and Automation Dividend layers still have merit on their own.
- If a real comparative costing exercise (Appendix C) finds that a Managed Transition is materially more expensive or economically damaging than an Unmanaged Transition, Section XIII's central fiscal argument would be reversed, not merely qualified — and this paper commits in advance to treating that outcome as a genuine refutation, not a modeling error to be argued around.

None of these six conditions has been tested in this paper. They are named here specifically because a paper that states its own falsification conditions in advance is more credible than one that only defends its thesis after the fact.

XVI. Policy Ask

This paper does not recommend a national rollout as immediately ready. It recommends a research and pilot pathway, with measurement preceding scale:

- Participation Floor / Automation Dividend experimentation — a regional pilot testing the tax-vs-ownership mechanism choice from Section XII directly, rather than assuming one.
- Public compute / UCI pilot — bounded initially to inference infrastructure on open-weight models (Section IX), not frontier AGI operation.

- Infrastructure planning — the energy and hardware dependency chain (Section X) as its own coordinated planning track, not an assumption riding along with the compute-allocation track.
- Institutional revenue adaptation — a dedicated track addressing Section VI and VII's specific findings (Social Security's existing shortfall, municipal bifurcation, insurance-market restructuring), not an assumption that household-level policy alone resolves institutional-level decoupling.

This paper is not written as a demand. It is written as an invitation to policymakers, researchers, and the broader public to begin reasoning several years ahead of where current policy discussion sits — with measurement, not national-scale commitment, as the immediate next step.

XVII. Conclusion

Automation may solve production while destabilizing distribution. That is the paper's central distinction, and everything else in this document is an elaboration of what follows from taking it seriously.

Income prevents exclusion. Ownership preserves participation. Compute preserves agency. Education multiplies capability. Infrastructure makes agency physically possible. And institutional financing must eventually follow where productive value actually appears, rather than remaining permanently attached to the industrial-era labor proxies through which value historically appeared.

The question this paper leaves its reader with is not whether technological change of this kind occurs — the evidence in Sections II, III, and VII establishes that it is already underway. The question is whether institutions — households, municipalities, social insurance, insurance markets, and the tax architecture beneath all of them — evolve quickly enough for the productive gains of that change to remain broadly participatory, or whether they remain attached to proxies that are quietly ceasing to track where value is actually created.

Appendix A — Formalisms, Definitions, and Claim Classification

How to read this paper's claims

Every substantive claim in this paper falls into one of six categories. The category is not always stated inline, but the placement and hedging of each claim throughout the paper follows this scheme:

Category	Examples in this paper
A. Externally established, adequately sourced	Labor-share decline (BLS; Karabarbounis & Neiman 2014); Acemoglu & Restrepo's displacement/reinstatement framework and wage-inequality decomposition (2019, 2022); CBO's Social Security actuarial deficit (2023); China's nuclear/grid investment figures (World Nuclear Association); US VAT absence (OECD, Tax Foundation).
B. Externally established, citation completed this round	Loudoun County / Quincy data-center tax figures; Georgia / Texas abatement-cost estimates; 2026 SSDI/robot-exposure study; agentic-AI insurance/reinsurance literature.
C. Evidence-supported interpretation	The two-sided municipal ledger; the narrowed insurance-correlation claim (speed/opacity/model-concentration, not correlation itself); the reading of the SSDI counter-evidence as scoped to manufacturing robotics.
D. Hypothesis introduced by this paper, requiring empirical testing	Institutional Revenue Decoupling as a named, general phenomenon; the labor-share-decline channel compounding Social Security's payroll-tax gap; the municipal two-sided-ledger falsifiable test as specified.
E. Illustrative formalism, not a calibrated model	$Y_t = W_t + R_t$; $C_t \le W_t + D_t + R_{h,t} + B_t$; $T_{l,t}/P_t$ decline under a payroll-dependent tax architecture.
F. Policy proposal	Participation Floor, Automation Dividend (both mechanisms), Universal Compute Income, the ownership default rule, the three-track policy ask.

Distribution Break (Section III)

$Y_t = W_t + R_t$; $C_t \le W_t + D_t + R_{h,t} + B_t$. Illustrative notation; not calibrated.

Institutional Revenue Decoupling (Section VI)

Definition: occurs when automation changes the relationship between aggregate productive capacity and legacy financing proxies (payroll, wage income, employment counts, human-operated risk), causing institutional revenue or risk-pool structure to stagnate, decline, or destabilize even while output rises.

Formalism: P_t (aggregate output), $T_{l,t}$ (labor-linked tax revenue), $T_{c,t}$ (capital/output-linked tax revenue), G_t (institutional obligations). Hypothesis: $P_t \uparrow$ with $W_t/P_t \downarrow \Rightarrow T_{l,t}/P_t \downarrow$ under a payroll-dependent tax architecture.

Appendix B — Adversarial Review Method

This paper underwent two rounds of structured internal adversarial review, following the protocol used on prior EM Foundation working papers. Each round inhabited several specialist roles tasked with finding weaknesses, unsupported claims, and empirical gaps; classified each finding Fatal, Major, Moderate, or Minor; and corrected every Fatal or Major finding directly in the manuscript where correction was possible through revision. This process is internal editorial discipline, not a substitute for external peer review, and this paper does not claim the standing of a peer-reviewed publication on the strength of it.

The point of disclosing this process is not that the paper survived review — it is that the review materially changed the paper's substantive claims. Four examples:

- The municipal argument changed from a one-directional “automation erodes the local tax base” claim to the two-sided ledger in Section IV, after real data-center tax-revenue data (Loudoun County, Quincy) directly contradicted the original one-sided framing.
- The 2026 SSDI/robot-exposure counter-finding (Section VII.A) — that robot exposure in manufacturing lowers, not raises, disability applications — was identified during review and retained rather than omitted, despite complicating the paper's own displacement hypothesis.
- The insurance-correlation claim (Section VII.B) was narrowed from “correlated loss is a new problem for insurers” to the more defensible and more specific claim that the speed and opacity of AI-model-driven correlation, and foundation-model concentration specifically, is the novel element — correlated catastrophic loss itself is an established underwriting category.
- The fiscal comparison baseline changed from “what does this proposal cost” to the Managed-versus-Unmanaged Transition framing in Section XIII, once review established that the status quo carries its own real, largely uncosted fiscal trajectory.

Full findings, including corrected reasoning for each and items not corrected in this pass, are preserved in the project's internal working record and are not reproduced in full here.

Appendix C — Fiscal Costing Model Specification (Not Yet Performed)

This appendix specifies what a companion costing model must contain before Section XIII's fiscal question can be considered resolved. It is a specification for future work, not a completed model, and no figures from it appear anywhere in this paper.

Required structure

Scenario A (Managed Transition) must separately model: Participation Floor cost against actual regional housing/tax data; Automation Dividend revenue and distribution under both the tax and ownership mechanisms as distinct sub-scenarios, never blended into one estimate; UCI token cost against actual current inference pricing, projected using observed (not assumed) historical cost-decline rates; and infrastructure adaptation cost benchmarked against real interconnection-queue and comparative buildout data already cited in Section X.

Scenario B (Unmanaged Transition) must separately model: Social Security payroll-tax shortfall acceleration under a labor-share-decline sensitivity added to the existing SSA/CBO base case; municipal fiscal bifurcation, net of the property-tax gains documented in Section IV; displacement-support cost under current-law unemployment insurance systems; insurance-market restructuring cost, treated as a qualitative risk flag rather than a dollar estimate given data immaturity; and stranded-asset risk from reactive infrastructure buildout, using the existing 13-percent interconnection-completion-rate figure cited in Section X.

Required properties

The model must be falsifiable — capable of finding Scenario B costs less than Scenario A if the data supports that conclusion. Every major input requires a sensitivity range, not a point estimate. The two Automation Dividend mechanisms must remain separately versioned, never blended. Estimates should be regionally disaggregated where data allows, given Section IV's hosting/non-hosting divergence.

This is a bounded, falsifiable comparative-cost exercise using largely existing public data (BLS, CBO, SSA, Census of Governments, state UI trust fund reports, current AI-inference pricing), not a request for novel primary research or a full macroeconomic simulation. It is plausibly a standalone EM Foundation working paper in its own right.

Appendix D — Adjacent Governance Questions (Outside This Paper's Scope)

Questions of AI representation and governance legitimacy — as AI systems become more directly involved in judicial, legislative, and administrative reasoning, and whether AGI/ASI systems' own interests should eventually be represented in governance — are live enough to name but are a distinct problem from this paper's economic argument, and are left for a dedicated future EM Foundation paper rather than folded into the core argument here.

References

Numbered, one consistent style (author-date with full publication detail). Every source below is cited at least once in the text; every in-text citation resolves to an entry here.

- 1. 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
- 2. Acemoglu, D., & Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. Econometrica, 90(5), 1973–2016. https://doi.org/10.3982/ECTA19815
- 3. Congressional Budget Office. (2023). CBO's 2023 Long-Term Projections for Social Security. Washington, DC: Congressional Budget Office, Publication No. 59340. https://www.cbo.gov/publication/59340
- 4. Good Jobs First. (2026). Even Cloudier with a Greater Loss of Spending Control: How Data Center Tax Abatements Undermine Public Budgets. https://goodjobsfirst.org/even-cloudier-with-a-greater-loss-of-spending-control-how-data-center-tax-abatements-undermine-public-budgets/ (accessed September 2026).
- 5. Gupta, R., & Kumar, S. (2026). Agentic AI and Occupational Displacement: A Multi-Regional Task Exposure Analysis of Emerging Labor Market Disruption. arXiv:2604.00186.
- 6. Karabarbounis, L., & Neiman, B. (2014). The Global Decline of the Labor Share. Quarterly Journal of Economics, 129(1), 61–103. https://doi.org/10.1093/qje/qjt032
- 7. MRSC (Municipal Research and Services Center of Washington). (2026). What the Data Center Boom Means for Washington Local Governments. https://mrsc.org/stay-informed/mrsc-insight/august-2026/data-centers (accessed September 2026).
- 8. NetChoice. (2026). Data Centers 101: Do Data Centers Pay Taxes? https://netchoice.org/data-centers-101-do-data-centers-pay-taxes/ (accessed September 2026).
- 9. National Taxpayers Union. (2026). Data Centers: Taxpayers Can Be Victors, Not Victims, in the AI Revolution. https://www.ntu.org/publications/detail/data-centers-taxpayers-can-be-victors-not-victims-in-the-ai-revolution
- 10. OECD. (2024). Consumption Tax Trends 2024: VAT/GST and Excise Rates, Trends and Policy Issues. Paris: OECD Publishing. https://www.oecd.org/en/publications/2024/11/consumption-tax-trends-2024_57c7322a.html
- 11. Oxfam International. (2026, January). Resisting the Rule of the Rich [Davos briefing paper].

- 12. Altindag, D. T., El Cheikh Taha, R., Nunley, J. M., & Seals, R. A. (2026). Robots and the Public Finance of Disability Insurance. arXiv:2607.02892.
- 13. Zhu, Q. (2026). Insurance of Agentic AI. arXiv:2606.05449.
- 14. Xu, B., Dai, X., Yang, F., & Zhang, K. (2026). When Agent Automation Becomes Profitable: Quantifying and Insuring Autonomous AI Risk through Trace-Economic Underwriting. arXiv:2606.16465.
- 15. Searchlight Institute. (2026). States Can Leverage Tax Incentives to Make Data Centers Better Grid Citizens. https://www.searchlightinstitute.org/research/states-can-leverage-tax-incentives-to-make-data-centers-better-grid-citizens/ (citing Lawrence Berkeley National Laboratory interconnection-queue data; secondary citation — primary LBNL dataset not independently accessed for this draft).
- 16. Tax Foundation. (2026). OECD Tax Revenue by Country: Sources of Government Revenue. https://taxfoundation.org/data/all/global/oecd-tax-revenue-by-country/ (accessed September 2026; compiles OECD data).
- 17. US Bureau of Labor Statistics. Private Nonfarm Business Sector: Labor Share [MPU4910141]. Retrieved via FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/MPU4910141 (accessed September 2026). Primary source for the 67.8% (1987) to 58.1% (2023) labor-share figures cited in this paper.
- 18. World Nuclear Association. (2026). Nuclear Power in China [country profile]. https://world-nuclear.org/information-library/country-profiles/countries-a-f/china-nuclear-power (accessed September 2026).