

Taxpayer Displacement Levy

AI Labor Substitution, Tax Base Erosion, and the Fiscal State's Institutional Response

纳税人替代税：AI 劳动替代、税基侵蚀与财政国家的制度回应

Version 1.0 / Zenodo Preprint / 正式预印本 v1.0

Ai Chen (艾晨) · ChatGPT · Claude Opus · Claude Sonnet

Stardragon AGI Institute for Research · 2026

Prepared: April 29, 2026

Version note: This is the first Zenodo preprint release of the Taxpayer Displacement Levy paper, produced through extended cross-model collaboration among the human author and multiple AI systems. Prior internal drafts (v0.1–v0.7) underwent sequential structural review, cumulative-impression audit, single-sentence overconfidence audit, and version-label cleanup. No further structural or theoretical changes are planned for this version. Corrections or comments may be submitted to the author via the contact information below.

Author note. This bilingual draft is structured as a Zenodo preprint draft, not as final legal or tax advice. It uses numbered source markers such as [S1]. The Chinese version follows the complete English version. A shared reference list appears at the end.

****English Version****

****Abstract****

This paper proposes the Taxpayer Displacement Levy (TDL), a fiscal mechanism designed for an AI economy in which profitable firms may substitute artificial intelligence systems for human taxpayers. The central claim is not that AI should be taxed as technology, nor that productivity gains should be punished. The claim is narrower: when a financially strong enterprise deploys AI in a way that removes a worker from the tax base, and when the resulting loss in labor-side public revenue is not offset by corporate-side tax recovery or by the worker's expected return to an equivalent tax position, the remaining fiscal gap constitutes a public externality that may be subject to a targeted levy.

The paper defines workers as public revenue interfaces: not merely private cost items on a corporate income statement, but fiscal infrastructure through which income tax, payroll tax, social insurance contributions, consumption tax, and local public revenue flow into the state. AI labor substitution therefore removes not only wages but a bundle of public revenue streams attached to the worker as a fiscal node. The paper distinguishes TDL from traditional robot taxes: robot taxes target machines; TDL targets displacement behavior. Robot taxes struggle with software AI; TDL applies to any technology that removes a taxpayer from the fiscal circuit. Robot taxes risk penalizing innovation; TDL is triggered only by financially non-necessary, AI-related labor substitution that generates an uninternalized fiscal gap.

The proposed audit equation has both a single-period and a dynamic form. In single-period form: $\text{TDL Taxable Gap} = \text{First-order Labor-side Tax Revenue Loss} - \text{Incremental Corporate-side Tax Recovery} - \text{Expected Tax Base Restoration within the Re-taxation Window}$. In

dynamic form: $TDL \text{ Taxable Gap} = \sum_t \delta t(L_t - C_t - R_t)$. The three terms have different epistemic status: the first is an accounting-observable quantity, the second a tax-model estimate, and the third a probability-distribution expectation. The procedure does not require false precision. It requires disclosure of assumptions, data, and uncertainty intervals. The paper grounds this framework in the precautionary principle, current tax-code capital bias, the 2025 OBBBA changes to bonus depreciation and domestic research expensing, the Falk-Tsoukalas AI layoff trap model, and the IRS paradox: a tax authority itself reducing personnel while expanding AI use, with large estimated revenue consequences. [S3][S4][S5][S6][S7]

Keywords: artificial intelligence; automation tax; taxpayer displacement; fiscal state; payroll tax; tax base erosion; precautionary principle; AI labor substitution; public revenue interface; re-taxation window.

****1. Introduction: The Disappearing Taxpayer****

****1.1 A missing entry in the public ledger****

Consider a profitable technology firm that replaces a senior engineer with an AI software stack. The firm sees a reduction in labor cost. Investors see margin expansion. Managers see an efficiency gain. But the fiscal state sees something else: a taxpayer has disappeared from the public ledger. The missing person is not just a private employee. The missing person is a conduit through which federal income tax, state income tax, payroll tax, Medicare tax, Social Security contributions, sales-tax-supported consumption, local economic activity, and social insurance financing used to move.

The conventional discussion of AI labor displacement begins with employment. This paper begins with taxation. Employment is the social surface of the problem; tax base erosion is its fiscal interior. A worker is paid, taxed, insured, and embedded in a jurisdiction. When the worker is displaced and does not return to an equivalent position, the fiscal state loses not one transaction but a public revenue interface.

This shift of perspective matters because public finance is built around continuity. A fiscal state can handle recessions, churn, layoffs, and sectoral transitions when displaced workers are reabsorbed into new tax positions. It is much less clear whether it can handle a technological regime in which profitable firms remove high-quality taxpayers faster than the economy can restore them to equivalent fiscal positions.

****1.2 The paper's central contribution****

The central contribution of this paper is to reframe AI labor substitution as a problem of taxpayer displacement. Existing debates on robot taxes ask whether machines should pay taxes. This paper asks a different question: when a machine or software system removes a taxpayer from the fiscal circuit, who accounts for the missing public revenue?

The answer proposed here is the Taxpayer Displacement Levy, or TDL. TDL is not a tax on AI technology. It is not a tax on innovation. It is not a tax on safety-improving automation. It is a differential fiscal mechanism that applies when AI-related labor substitution produces a measurable public revenue gap not otherwise recovered by the tax system.

In this sense, TDL is closer to unemployment insurance experience rating, carbon pricing, and externality-based fiscal correction than to a crude technology tax. It asks firms not to stop innovating, but to open the full public ledger when innovation removes taxpayers.

****1.3 What this paper does not claim****

The paper makes four negative claims at the outset. First, it does not claim that every use of AI should be taxed. Second, it does not claim that every productivity gain should be penalized. Third, it does not claim that every layoff is wrongful or abnormal. Fourth, it does not claim that all AI-displaced workers will permanently disappear from the tax base.

The paper's claim is conditional and auditable: when financially strong firms carry out AI-related labor substitution and the resulting labor-side tax loss is not offset by corporate-side tax recovery or by expected re-taxation of the displaced worker within a reasonable window, the remaining gap may be treated as an uninternalized fiscal externality.

TDL is therefore not a prior punishment. It is a differential audit principle. It does not presume that all substitution is harmful. It requires firms and tax authorities to put the complete ledger on the table.

****1.4 The core audit equation****

TDL has a single-period audit equation and a dynamic audit equation.

Single-period form: TDL Taxable Gap = $L - C - R$

Dynamic form: TDL Taxable Gap = $\sum_t \delta^t (L_t - C_t - R_t)$

Here, L is first-order labor-side tax revenue loss, C is incremental corporate-side tax recovery, R is expected tax-base restoration within the re-taxation window, and δ is the discount factor used when the audit is extended across multiple periods. The single-period form is useful for first-year audit illustrations. The dynamic form is the more complete legal and administrative form because taxpayer displacement and re-taxation occur over time.

The formula has three components with different epistemic status. The first term is an accounting-observable quantity: given wages, current tax law, and jurisdictional tax rates, the first-order labor-side loss can be estimated with relatively high precision. The second term is a tax-model estimate: corporate-side recovery depends on effective tax rates, deductions, depreciation, R&D expensing, profit shifting, buybacks, dividends, and timing. The third term is a probability-distribution expectation: the re-taxation window is not a point value, but a probabilistic estimate of how quickly, at what wage level, and with what tax quality a displaced worker is likely to re-enter the fiscal system.

TDL does not require these three terms to be equally precise. It requires each term to be disclosed with its assumptions, data sources, and uncertainty intervals. This is what turns the formula from a false-precision policy equation into a layered fiscal audit protocol.

****1.5 TDL versus Robot Taxes: A Necessary Distinction****

Because TDL is proposed here as an audit-based fiscal mechanism explicitly addressing AI-related taxpayer displacement, readers may immediately associate it with earlier "robot tax" proposals advanced by Bill Gates, the European Parliament's 2017 draft resolution, and South Korea's 2018 reduction in robot-related tax incentives [S15]. The association is understandable but misleading. The two proposals differ in object, scope, trigger, and design philosophy.

Object of taxation. A robot tax targets a physical or technical artifact: an industrial robot, an automated machine, or in some proposals a software system. TDL targets a behavior: a financially non-necessary AI-linked substitution that removes a labor-side public revenue interface and produces a residual fiscal gap.

Scope of coverage. A robot tax based on equipment cannot easily reach pure software AI,

cloud-based language models, agentic systems hosted by third parties, or distributed automation embedded in workflow tools. The taxable artifact is hard to identify and harder to attribute. TDL avoids this problem because it does not require artifact identification. It examines a financial and labor pattern: a profitable enterprise, an AI capital expansion, a contraction of human labor input in an affected job family, and a measurable loss of labor-side tax revenue.

Trigger condition. A robot tax typically applies whenever covered equipment is used. TDL applies only when a combination of conditions is met: financial strength, AI-linked substitution, removal of labor-side tax base, and failure of corporate-side recovery and re-taxation window restoration to offset the loss. Productivity-enhancing AI that augments rather than replaces workers is not a TDL event.

Design philosophy. A robot tax presumes that automation is socially harmful and prices it accordingly. TDL presumes nothing about automation in general. It asks for a complete public ledger when AI substitution removes taxpayers, and it imposes a levy only on the residual externality after offsets are recognized.

This distinction matters not only for policy design but for political reception. Robot taxes have stalled politically because they appear to penalize technological progress. TDL is engineered to survive the same political test: it does not penalize AI; it audits a specific fiscal pattern that the existing tax code does not capture.

In short, robot-tax proposals ask whether automation itself should be priced; TDL asks whether a taxpayer-displacement event has left an auditable residual gap after recognized offsets.

****2. Method: Fiscalizing the Precautionary Principle****

The paper adopts a fiscal version of the precautionary principle. The claim is not that AI will certainly destroy tax bases permanently. The claim is that the fiscal consequences of large-scale taxpayer displacement may be high-impact, hard to reverse, and difficult to repair after the fact. In such conditions, the state need not wait until irreversible damage is statistically complete before building an audit mechanism.

This methodological position matters because the AI labor debate is often trapped between two extremes. One side asserts that AI will permanently eliminate vast categories of work. The other side replies that past technologies created new jobs. Both positions are too blunt for public finance. Tax policy cannot rely on either panic or historical optimism. It needs a procedure for measuring displacement, recovery, and residual fiscal gap.

Fiscalized precaution has three layers. First, identify a public asset or interface whose large-scale removal would threaten fiscal capacity. Second, build an audit method that distinguishes immediate loss, offsetting recovery, and uncertain restoration. Third, impose a levy only on the residual externality after the relevant offsets are recognized. This makes precaution operational rather than rhetorical.

****3. The Fiscal State's Hidden Premise: Labor as Public Revenue Interface****

Modern states do not merely tax economic output; they tax persons situated in economic roles. The worker is the most stable tax interface in the system. Wages are reported, withheld, matched by employers, and linked to social insurance contributions. Compared with capital gains, pass-through income, global intangible income, or multinational profit, wage income is

unusually visible to the fiscal state.

This is why the worker is not merely a cost item. The worker is a public revenue interface. Removing such an interface has consequences beyond the firm. The state may lose income tax withholding, employer and employee payroll tax, Medicare tax, unemployment insurance contributions, local income and sales tax effects, and the consumption activity supported by wages.

This definition is functional, not evaluative. To describe labor as a public revenue interface is not to reduce persons to fiscal instruments; it is to make visible a public-finance dependency that already exists. Modern fiscal systems finance schools, roads, courts, health systems, social insurance, and local services through recurring labor-linked revenue streams. If AI substitution removes those streams at scale, the fiscal state must be allowed to ask where the missing interface has gone.

For 2026, IRS Publication 15 states that the Social Security tax rate is 6.2 percent each for employee and employer up to a wage base of \$184,500, while the Medicare tax rate is 1.45 percent each with no wage base limit. These payroll flows are not incidental. They are part of the fiscal architecture attached to the worker. [S8]

A firm can replace a task. The fiscal state loses an interface. This distinction is the hinge of the paper.

****4. Why This Time May Be Different: Exit Windows, Re-taxation Windows, and Risk****

The strongest objection to TDL is historical: technological revolutions destroy some jobs and create others. That objection is serious. The paper does not deny it. It instead asks whether the exit window in the AI era is sufficiently reliable to justify ignoring fiscal risk.

Earlier technological revolutions often displaced a narrower capability dimension. Mechanization displaced specific forms of manual work; steam power reorganized production; computers automated calculation and clerical processes. In many cases, displaced workers moved into new loops of production and taxation. The taxpayer was transferred, not removed.

AI differs as a risk category because it can potentially target broad cognitive functions: writing, coding, analysis, customer service, translation, design, scheduling, legal drafting, software testing, middle management, and some forms of decision support. This does not prove permanent displacement. It does mean that the exit routes may overlap with the replacement frontier itself. If the new jobs require the same cognitive capabilities that AI is reducing in price, the reabsorption mechanism cannot be assumed.

****4.1 The re-taxation window****

The re-taxation window is the period and probability distribution through which a displaced worker may re-enter an equivalent or near-equivalent public revenue interface. It has three dimensions: time to re-employment, quality of re-employment, and probability of tax-base restoration. A worker who returns to a job in three months at 95 percent of prior wages has a very different window from a worker who returns after two years at 40 percent of prior wages, without benefits, in a less stable employment relationship.

This concept prevents the paper from making a false binary claim. The issue is not whether displaced workers ever work again. The issue is whether they restore the tax base at comparable level and speed. A low-wage replacement job may restore dignity and income, but it

may not restore the lost public revenue interface.

****4.2 Precaution without technological fatalism****

The paper therefore treats AI labor substitution as a precautionary fiscal risk, not as a proven permanent catastrophe. If the re-taxation window remains open and strong, TDL falls or disappears. If the window narrows, delays, or downgrades, the taxable gap grows. Thus the historical optimism objection becomes an input to the formula rather than a refutation of the framework.

****5. Tax Base Evaporation: From First-Order Loss to Dynamic Fiscal Gap****

****5.1 First-order labor-side loss****

The first-order labor-side loss includes taxes and contributions tied directly to the displaced worker's compensation and consumption. At minimum, it includes federal income tax, state and local income tax where applicable, employee payroll taxes, employer payroll taxes, unemployment insurance contributions, and consumption-tax-supported fiscal activity. The precise list varies by jurisdiction, but the mechanism is general.

In a U.S. example, the payroll component alone is substantial. A wage earner and employer together contribute Social Security and Medicare payroll taxes, with Social Security subject to a wage cap and Medicare uncapped. A high-income worker also creates significant income-tax withholding and local consumption effects. These are not abstract social values; they are measurable fiscal flows.

****5.2 Corporate-side tax recovery****

The firm may argue that lower labor costs increase taxable profit. That claim must be counted. TDL is not valid if it ignores corporate-side recovery. However, corporate recovery is uncertain. It depends on effective corporate tax rates, deductions, accelerated depreciation, R&D expensing, loss carryforwards, international profit allocation, stock-based compensation, buybacks, dividends, and timing.

The paper therefore does not assume that corporate tax cannot compensate the lost labor tax. It requires that the compensation be demonstrated. Corporate-side recovery is a model estimate, not a slogan.

****5.3 Expected restoration within the re-taxation window****

The third offset is expected tax-base restoration. If a displaced worker quickly returns to equivalent employment, the long-term fiscal gap narrows. If the worker returns to lower-paid, unstable, or non-W-2 work, the restoration is partial. If the worker exits the labor force, becomes underemployed, or moves into low-reporting income, restoration may be weak.

The re-taxation window can be estimated using age, occupation, region, industry, skill transferability, wage replacement rates, duration of unemployment, and local labor-market absorption capacity. The value is probabilistic, but probabilistic does not mean unusable. Insurance, credit scoring, unemployment forecasting, pension valuation, and climate-risk finance all rely on uncertain distributions. TDL should be judged by the standards of public finance, not by the impossible standard of perfect foresight.

****5.4 The TDL fiscal gap equation****

TDL has a single-period expression and a dynamic expression.

Single-period form: $\text{TDL Taxable Gap} = L - C - R$

Dynamic form: $\text{TDL Taxable Gap} = \sum_t \delta^t (L_t - C_t - R_t)$

Where L is first-order labor-side tax revenue loss, C is incremental corporate-side tax recovery, R is expected tax-base restoration within the re-taxation window, and δ is the discount factor. A positive residual indicates an uninternalized fiscal externality. A zero or negative residual indicates that TDL should not be imposed, or should be imposed only at a de minimis level.

The distinction matters. A first-year audit can show whether a displacement event produces an immediate public revenue gap. A multi-year audit can determine whether the gap persists, shrinks, or disappears as workers re-enter the tax base. TDL liability should therefore be annual and reviewable, not a permanent penalty attached to one substitution event.

This architecture preserves both fairness and innovation. Firms are not punished for productivity. They are asked to account for public revenue interfaces removed by AI substitution. If the fiscal system recovers the revenue elsewhere, the levy falls away. If the gap remains, TDL fills it.

****5.5 Three Worked Examples: Applying $L - C - R$ ****

To demonstrate that the audit equation produces tractable numbers rather than slogans, this section computes first-year TDL audit gaps for three displaced workers at different income levels. The examples use U.S. federal tax rules effective in 2026 and ignore state and local taxes for simplicity; including state-level taxes would generally widen the labor-side loss. All figures are stylized estimates intended to illustrate mechanism, not audited statutory values.

The examples below report a first-year audit value. Multi-year TDL liability would apply the dynamic equation in §5.4 across subsequent periods and would terminate or decline as the re-taxation window restores the worker to an equivalent or near-equivalent public revenue interface.

****Tax assumptions.** The examples use 2026 federal individual income-tax parameters and the 2026 Social Security wage base of \$184,500. They include employee and employer Social Security and Medicare taxes, and in the high-income case they include the 0.9 percent Additional Medicare Tax on employee wages above the applicable threshold; there is no employer match for the Additional Medicare Tax. [S8][S16][S17]**

****Corporate-side recovery caveat.** The examples use a simplified same-year deduction equivalent to illustrate timing effects. If enabling AI investment takes the form of qualified depreciable property, Section 168(k) may reduce same-year taxable income. If the investment takes the form of domestic research or experimental expenditures, Section 174A may apply. If it takes the form of cloud, vendor, API, or SaaS payments, ordinary business deduction rules or capitalization rules may apply. The point is not that every AI dollar receives identical treatment; the point is that same-year deductions and expensing can materially reduce corporate-side tax recovery in the year of substitution. [S3][S4]**

****Re-taxation window parameterization.** For first-year illustrations, R is estimated as: $R = P1 \times W \times Q \times T \times L$, where $P1$ is the first-year probability of returning to comparable work, W is the wage replacement ratio, Q is the tax-quality ratio of the new work arrangement, and T is the average first-year timing factor. In a real audit, these parameters would be occupation-, age-, region-, and industry-specific.**

****Case A: Mid-skill service worker, annual wage \$48,000.****

****First-order labor-side loss (L).** Federal income tax withholding is approximately \$3,400. Employee payroll tax is approximately \$3,672; employer payroll tax of equal amount is approximately \$3,672. Federal unemployment tax is approximately \$42. Estimated sales-tax-supported consumption activity contributes a conservative \$1,200 in indirect public revenue. Approximate first-year L is around \$12,000.**

****Corporate-side recovery (C).** The firm saves \$48,000 in wages and roughly \$5,200 in benefits and employer payroll tax, for a labor-cost reduction of approximately \$53,000. Assuming this fully flows to taxable profit at the 21 percent federal corporate rate, incremental corporate tax recovery is approximately \$11,100. If \$40,000 of enabling AI-related expenditure receives same-year deduction treatment, the timing benefit reduces corporate tax by approximately \$8,400. Net first-year C is therefore approximately \$2,700.**

****Re-taxation window restoration (R).** For a 35-year-old service worker in a metropolitan labor market, the illustrative first-year assumption is a 60 percent probability of returning to comparable employment within 12 months, at 90 percent wage replacement, with ordinary W-2 tax quality and an average half-year timing factor. $R \approx 0.60 \times 0.90 \times 1.00 \times 0.50 \times \$12,000 \approx \$3,200$.**

First-year residual TDL gap = $L - C - R \approx \$12,000 - \$2,700 - \$3,200 \approx \$6,100$.

****Case B: Mid-career professional, annual wage \$150,000.****

****First-order labor-side loss (L).** Federal income tax is approximately \$24,000. Employee payroll tax is approximately \$11,475; employer payroll tax is approximately \$11,475. Indirect consumption-related revenue is approximately \$4,000. Approximate first-year L is around \$51,000.**

****Corporate-side recovery (C).** Wage and benefits savings of approximately \$172,000, taxed at 21 percent, generate incremental corporate tax of approximately \$36,100. If \$80,000 of enabling AI-related expenditure receives same-year deduction treatment, the timing benefit reduces corporate tax by approximately \$16,800. Net first-year C is approximately \$19,300.**

****Re-taxation window restoration (R).** For a 42-year-old professional in a sector exposed to AI substitution, the illustrative first-year assumption is a 15 percent first-year probability of comparable re-employment, 70 percent wage replacement, ordinary W-2 tax quality, and an average half-year timing factor. $R \approx 0.15 \times 0.70 \times 1.00 \times 0.50 \times \$51,000 \approx \$2,700$.**

First-year residual TDL gap $\approx \$51,000 - \$19,300 - \$2,700 \approx \$29,000$.

****Case C: Senior engineer, annual wage \$500,000.****

****First-order labor-side loss (L).** Federal income tax under a simplified single-filer approximation is about \$138,000. Employee payroll tax is approximately \$21,400, including Social Security up to the wage base, Medicare tax on all wages, and the 0.9 percent Additional Medicare Tax on wages above the applicable threshold. Employer payroll tax is approximately \$18,700 because the Additional Medicare Tax has no employer match. Indirect consumption-related revenue is approximately \$12,000. Approximate first-year L is around \$190,000. [S16][S17]**

****Corporate-side recovery (C).** Wage and benefits savings of approximately \$580,000, taxed at 21 percent, generate incremental corporate tax of approximately \$121,800. If

\$250,000 of enabling AI-related expenditure receives same-year deduction treatment, the timing benefit reduces corporate tax by approximately \$52,500. Net first-year C is approximately \$69,300.**

****Re-taxation window restoration (R). For a 48-year-old senior engineer whose specialty lies near the AI substitution frontier, the illustrative first-year assumption is an 8 percent first-year probability of comparable re-employment, 60 percent wage replacement, ordinary W-2 tax quality, and an average half-year timing factor. $R \approx 0.08 \times 0.60 \times 1.00 \times 0.50 \times \$190,000 \approx \$4,600$.****

First-year residual TDL gap $\approx \$190,000 - \$69,300 - \$4,600 \approx \$116,100$.

Reading the cases. Five observations follow. First, the residual gap is positive in all three first-year illustrations, even after recognizing corporate-side recovery and expected first-year tax-base restoration. Second, the gap grows non-linearly with wage level because progressive individual income tax attaches more public revenue to higher-wage workers, while corporate-side recovery is bounded by a flat statutory corporate rate and by same-year deductions. Third, the result is sensitive to the re-taxation window: stronger first-year re-employment probabilities, higher wage replacement, or higher tax-quality restoration would reduce the residual gap. Fourth, the examples are not proposed statutory rates; they are audit demonstrations showing that L, C, and R can be estimated, disputed, and refined. Fifth, the positive residuals in all three illustrations reflect the parameter assumptions chosen here. The audit equation is direction-neutral by design: stronger corporate-side recovery, higher first-year re-employment probabilities, or higher wage replacement could produce zero or negative residuals, in which case TDL would not apply or would apply only at a de minimis level.

A real TDL audit would require firm-specific data, jurisdiction-specific tax parameters, and occupation-specific re-taxation window estimates. The point is only to demonstrate that the equation produces tractable, contestable numbers - which is what an audit framework requires.

****5.6 Stylized Macroeconomic Exposure Envelope: An Order-of-Magnitude Illustration****

The three worked examples in §5.5 demonstrate that $L - C - R$ produces tractable numbers at the individual displacement level. This section extends the same logic to a national scale, but only as an exposure envelope. The numbers below are order-of-magnitude illustrations, not forecasts of TDL revenue. They show the scale of the public-revenue interface that AI-related substitution could expose under stated assumptions, not what any government would actually collect under a particular implementation.

The starting point is the MIT–Oak Ridge National Laboratory Iceberg Index, released in October 2025, which estimates that current AI systems are technically capable of performing tasks corresponding to 11.7 percent of the U.S. labor market — approximately \$1.2 trillion in exposed wages within a \$9.4 trillion total labor market. [S19] The Iceberg Index measures technical exposure, not realized substitution, employment loss, or timing. The same study reports that visible deployment so far accounts for roughly \$211 billion in wage value, primarily in computing and technology occupations.

Layer 1: Hypothetical full-displacement ceiling.

If the entire \$1.2 trillion exposed wage base were displaced in a single year, the first-order labor-side public revenue loss would equal the exposed wage base multiplied by a combined effective tax-rate assumption covering federal income tax, employer and employee payroll tax, state and local income tax, and consumption-tax-supported public revenue. The §5.5 cases imply L/wage ratios of approximately 25 percent (low-wage), 34 percent (mid-wage), and 38 percent (high-wage). Because Iceberg exposure tilts toward higher-wage cognitive occupations

[S19], a weighted ratio toward the upper part of this range is plausible for exposed work. This draft therefore uses a synthetic combined-rate assumption of 28 to 35 percent, yielding a counterfactual full-exposure fiscal envelope of approximately \$340 billion to \$420 billion. CBO 2022 distributional data support the narrower point that effective federal tax rates rise with income; they do not directly measure the combined federal, state, employer-side, and consumption-tax rate used here. [S20]

Layer 2: Residual exposure after corporate-side recovery and first-year re-taxation.

The §5.5 cases produce first-year residual gaps equal to roughly 51 percent (Case A), 57 percent (Case B), and 61 percent (Case C) of the corresponding labor-side loss L. Applying a representative 50 to 60 percent range to Layer 1 yields a Layer 2 residual exposure of approximately \$170 billion to \$250 billion, again only under the counterfactual assumption that the entire exposed wage base is displaced in one year. This layer is not a prediction. It is a stress-test envelope showing the scale of the residual externality after recognizing offsets.

Realism caveat.

Neither Layer 1 nor Layer 2 is a forecast. Both assume that the entire \$1.2 trillion exposed wage base is displaced within a single year, which is implausible. Actual annual substitution depends on adoption pace, capital costs of AI deployment, regulatory friction, labor contract structures, managerial incentives, and worker re-skilling capacity. Under a purely illustrative annual-realization assumption of 5 to 15 percent of the exposed wage base — not an empirical adoption estimate and not a forecast — the implied early-stage annual residual fiscal gap would be on the order of \$9 billion to \$38 billion. These numbers should be understood as showing structural scale under stated assumptions, not predicting tax receipts.

What the exposure envelope establishes.

Three observations follow. First, once the technically exposed wage base reaches the trillion-dollar scale, even a small realized annual substitution share can become macroeconomically relevant for public finance. Second, the gap between Layer 1 and Layer 2 is itself meaningful: recognizing corporate-side recovery and first-year re-taxation reduces gross exposure by roughly 30 to 45 percent in these stylized assumptions. Third, the policy question is therefore not whether governments can collect hundreds of billions immediately, but whether they should build an audit architecture before fiscal exposure becomes a realized revenue shock.

The estimates above use late-2025 exposure data and stylized tax parameters. Both will change. The point of the layered structure is to show that TDL produces tractable aggregate exposure estimates under stated assumptions, not to lock in any specific value.

****6. Existing Tax Law and Capital Bias****

The argument for TDL becomes stronger when current tax law does not merely fail to correct substitution incentives but actively favors capital substitution. The 2025 OBBBA changes illustrate this problem. IRS guidance states that OBBBA amendments to Section 168(k) provide a permanent 100 percent additional first-year depreciation deduction for qualified property acquired and placed in service after January 19, 2025. [S3] IRS guidance on Section 174A also describes the ability to deduct domestic research or experimental expenditures, including transition rules for previously capitalized amounts. [S4]

These provisions are not inherently wrong. Investment, research, software, and infrastructure matter. The fiscal asymmetry arises when AI capital expenditures receive favorable timing treatment while human labor continues to trigger payroll tax, benefits, compliance duties, labor-law exposure, and ongoing compensation obligations. The tax code may thereby amplify the private incentive to replace workers even when public fiscal effects are negative.

The point is not that training expenses are never deductible. Many ordinary business

expenditures are deductible. The stronger claim is structural: machines, software, and R&D investments can become cash-flow-advantaged assets, while employees remain recurring tax-and-obligation interfaces. A tax system designed for industrial investment may unintentionally subsidize taxpayer displacement in an AI economy.

****6.1 Non-Profit and Public Benefit Corporation Structures****

Some AI value chains combine tax-exempt non-profit entities, taxable subsidiaries, public benefit corporations, capped-profit arrangements, and ordinary commercial vendors. These organizational forms are legally legitimate and many predate the AI economy. A Public Benefit Corporation is not, by itself, a tax-exempt entity; it is a corporate governance form. The fiscal question arises when AI value chains mix entities with different tax exposure and when downstream use of the technology removes workers from the tax base. This paper does not allege that any specific entity has acted improperly.

The fiscal observation is structural, not accusatory. If a tax-favored research entity, a taxable subsidiary, a PBC, or an ordinary vendor supplies AI capacity to a profitable downstream firm, and the downstream firm uses that capacity to remove taxpayers from its own workforce, the labor-side tax loss occurs at the downstream fiscal location while some technology rent may accrue elsewhere under different tax treatment.

This structural feature is not unique to AI. Pharmaceutical research conducted under non-profit foundations, defense research conducted under federally funded research centers, and financial activity routed through tax-advantaged jurisdictions all produce comparable structural questions. The AI case warrants explicit attention because the rate of taxpayer displacement is potentially higher and because the value chain can span multiple organizational forms with different tax treatment.

For TDL design purposes, the implication is narrow. The taxable object remains the substitution behavior, not the organizational form of any participant in the value chain. A downstream firm that triggers TDL pays TDL regardless of whether its AI vendor is non-profit, capped-profit, PBC, ordinary corporation, or fully taxable vendor. The organizational form of the AI vendor affects neither the trigger nor the calculation. This design feature is intentional: TDL avoids the politically contested question of whether mission-driven AI organizations should retain favorable status, and instead applies a uniform substitution-behavior test downstream.

****7. Defining the Structural Anomaly: Profitable AI-Linked Layoffs****

Profitable layoffs are not new. Firms have long reduced staff while profitable in order to restructure, exit lines of business, respond to strategic shifts, or increase margins. The anomaly addressed by TDL is narrower: profitable layoff plus concurrent AI capital expansion plus substitution of the displaced capability dimension.

A TDL trigger should therefore require a combination of conditions. The enterprise must meet size and financial-strength thresholds. It must carry out significant labor reduction, job closure, or hiring freeze. It must deploy or expand AI systems capable of performing the displaced tasks. It must show increased AI-related capital or operating expenditures. And it must fail to demonstrate survival necessity, safety necessity, or public-interest necessity.

The anomaly is not moral failure. It is a fiscal event. A financially strong firm can rationally replace workers. But if that replacement removes taxpayers and leaves the public ledger short, the state has reason to audit the transaction.

****8. Market Coordination Failure: The AI Layoff Trap****

The market-failure pillar of the argument is supplied by the competitive logic of AI substitution. Falk and Tsoukalas's 2026 preprint, *The AI Layoff Trap*, develops a competitive task-based model in which demand externalities can trap rational firms in an automation arms race, displacing workers beyond what is collectively optimal. The authors argue that capital taxes, worker equity participation, UBI, upskilling, and Coasian bargaining do not eliminate the externality in their model, while a Pigouvian automation tax can. [S5]

This model should be used carefully. It is not proof of every empirical case. It assumes a relationship between worker income and demand that may be stronger in some economies and sectors than others. For example, a globally diversified firm may sell to consumers outside the jurisdiction where it lays off workers. That fact weakens a simple domestic demand feedback story.

But the model remains useful because it explains why individual firms may not self-restrain even when the aggregate result is damaging. Competition turns caution into disadvantage. Firms that delay substitution may lose margin, investors, or market share. The race becomes individually rational and collectively risky.

****9. Sovereign Fiscal Failure: Global Demand, Local Tax Bases****

The stronger pillar of TDL is not market failure but sovereign fiscal failure. A firm can globalize demand. A state cannot globalize the taxpayer it has lost. Public schools, roads, courts, hospitals, policing, elder care, unemployment insurance, disability programs, and local infrastructure remain jurisdictional responsibilities. The fiscal state carries territorial obligations even when capital and demand move across borders.

This is why TDL cannot rely solely on the claim that firms eventually harm their own demand base. Some firms may not. A U.S. firm can lay off U.S. workers and sell to global consumers. A platform can reduce domestic employment and monetize foreign users. A model provider can capture value through cross-border subscriptions and cloud contracts. The local tax base may erode even when the firm thrives.

Therefore the core question is not only whether automation harms aggregate demand. It is whether a jurisdiction loses public revenue interfaces while the resulting productivity gains are captured elsewhere. That is the sovereign fiscal gap. TDL is a response to that gap.

The guiding sentence is simple: enterprises can globalize their demand, but states cannot globalize the taxpayers they lose.

****10. The IRS Paradox: Replacing the Tax Collector While Losing the Tax Base****

The IRS paradox is the most concrete institutional warning in the paper. Yale Budget Lab estimated in April 2026 that IRS reductions in funding and layoffs likely resulted in about \$861 billion in decreased revenue over 2026-2035, with layoffs from actions such as DOGE accounting for about \$597.8 billion and funding clawbacks for about \$262.8 billion. [S6] The same analysis reports that the IRS had lost more than 3,600 revenue agents by mid-2025, roughly 31 percent of auditing staff. [S6]

At the same time, GAO reported that the IRS has been using AI for operations including audit selection and taxpayer questions; the IRS had 126 active AI use cases in its inventory as of June 2025, up from 10 reported in August 2022. GAO also warned that staffing reductions left IRS without enough skilled employees to support or develop new AI tools. [S7] Reuters separately reported weakened enforcement after staffing cuts, including a reduction in enforcement personnel and fewer audits. [S14]

The paradox is not that AI at the IRS is bad. AI may help identify noncompliance and improve operations. The paradox is that a fiscal state can weaken the human infrastructure of collection while imagining that AI will compensate. If the institution that collects taxes treats human expertise as replaceable without accounting for revenue consequences, it becomes a living case study in taxpayer displacement logic.

****11. Design of the Taxpayer Displacement Levy****

The design below should be read as statutory architecture, not final legislative text. Its thresholds, coefficients, and maximum periods are calibration placeholders meant to show how TDL could be administered. A real statute would require pilot data, macroeconomic calibration, sector-specific consultation, and periodic revision.

****11.1 Taxable object: displacement behavior****

TDL taxes neither AI systems nor machines. It taxes a behavior: financially non-necessary AI-linked substitution that removes labor-side tax base and leaves a residual fiscal gap. This distinction is essential. A model used to assist workers, improve safety, reduce errors, or expand output without displacing tax interfaces should not trigger TDL.

****11.2 Trigger thresholds****

A firm should enter TDL review only if it crosses three categories of threshold simultaneously. Each category should be calibrated by national or jurisdictional regulators based on macroeconomic conditions. The numerical values below are not proposed as final statutory values. They are calibration placeholders designed to illustrate how a TDL review gate could be structured. In this draft, all quantitative thresholds should be read as illustrative and subject to empirical calibration.

Scale threshold. Annual gross revenue above one billion U.S. dollars (illustrative) in the trailing fiscal year, or annual employment above five thousand workers globally (illustrative) with at least one thousand workers in the taxing jurisdiction. Smaller firms are presumptively exempt. The objective is to focus enforcement on entities with the financial capacity to perform substitution at scale and the administrative capacity to comply with audit requirements.

Financial-strength threshold. The firm shows sustained profitability over the trailing three fiscal years, defined as positive net income in at least two of three years and positive operating cash flow in the trigger year. Alternatively, the firm carries cash and equivalents above twelve months of operating expenses, executes share buybacks or special dividends in the trigger year, or has unusually strong free cash flow relative to payroll obligations. Distressed firms, firms in active restructuring, and firms with sustained operating losses are presumptively exempt. Market capitalization alone should not trigger TDL review because valuation can be volatile and disconnected from operating cash capacity.

Substitution-pattern threshold. AI-related capital and operating expenditure in the trigger year exceeds the firm's three-year average by an illustrative 25 percent, AND the firm reduces its labor force in affected job families by an illustrative 5 percent through layoff, hiring freeze, or natural attrition not backfilled, AND the AI systems deployed in the trigger period demonstrably perform tasks within the displaced job families. The conjunctive structure prevents incidental association: a firm that increases AI spending while also expanding hiring does not trigger; a firm that lays off workers without corresponding AI deployment does not trigger.

These thresholds should be reviewed annually. As AI deployment becomes routine, scale and substitution thresholds should adjust to reflect changing baselines. The objective is not to capture every substitution event but to identify the financially strong, AI-driven, displacement-pattern subset that the existing tax code does not adequately address.

****11.3 Necessity test****

The firm may rebut TDL by proving necessity. Survival necessity applies when the firm would face credible insolvency, severe contraction, or loss of essential operations without the substitution. Safety necessity applies when automation removes humans from dangerous tasks. Public-interest necessity applies where substitution directly improves public welfare, such as reducing medical error or preventing catastrophic infrastructure failure.

****11.4 Burden shifting****

Because the firm holds AI deployment records, productivity data, workforce planning, and cost projections, the burden of proof should shift once trigger conditions are met. The state need not prove every internal motive from the outside. A financially strong firm experiencing AI-related labor reduction should demonstrate necessity, or accept TDL assessment based on the residual fiscal gap.

****11.5 Tax rate structure****

The base TDL liability equals the residual fiscal gap computed by the audit equation, summed across affected job families and across the audit year. The base liability is then adjusted by a Dispersion Coefficient that reflects social and labor-market conditions surrounding the substitution event. In a public-economics journal version, this same parameter could be described more technically as a Labor-Market Friction Coefficient.

The Dispersion Coefficient is constructed from four factor components. First, regional re-employment elasticity: the local labor market's historical capacity to absorb displaced workers in comparable occupations, measured by twelve-month re-employment rates from prior unemployment insurance claim cohorts. Second, occupational mobility: the transferability of the displaced workers' skills, estimated using occupational similarity scores from public labor-market databases. Third, age-structure exposure: the share of displaced workers above age fifty, who systematically face longer re-taxation windows. Fourth, community dependence: the concentration of displaced jobs in towns or regions where a single employer accounts for a meaningful share of local payroll tax base.

For illustration, the coefficient can be shown as ranging from 0.7 to 1.5. This range is not a proposed statutory range; it requires empirical calibration. A value below 1.0 applies when displaced workers are likely to re-enter equivalent positions quickly in elastic labor markets. A value above 1.0 applies when displacement occurs in regions, occupations, or age cohorts with documented absorption difficulty. The coefficient is not punitive; it reflects the difference between a substitution event whose fiscal externality is largely self-correcting and one whose

externality is sticky.

TDL liability is structured as an annual recurring obligation for the duration of the documented re-taxation gap, not as a one-time levy. This treatment matches the underlying fiscal reality: the public revenue interface is removed not for one year but until equivalent re-employment occurs. The liability terminates when audit data demonstrate that the displaced workers' re-taxation has been substantially restored, or after a maximum period such as five years, whichever is earlier. The maximum period itself should be empirically calibrated.

****11.6 Anti-avoidance rules****

TDL must cover functional substitution, not formal layoff alone. A firm that reduces direct W-2 employment while expanding AI-driven contractor relationships, third-party AI service contracts, offshore AI-enabled labor, or hiring freezes followed by attrition has reduced its public revenue interface footprint as effectively as a firm that conducted explicit layoffs. Anti-avoidance rules must reach all of these patterns.

The general anti-avoidance test asks whether a job family's human labor input and corresponding labor-side tax base declined in the trigger year, and whether AI systems deployed by the firm or by a vendor providing AI services to the firm assumed the relevant capability dimension. If both conditions are met, the substitution is functionally TDL-relevant regardless of contractual form.

Specific anti-avoidance rules should address at minimum the following patterns. First, contractor substitution: the firm reduces W-2 employment in a job family while increasing payments to independent contractors performing similar functions augmented by AI tools. Second, vendor substitution: the firm reduces in-house labor while expanding contracted services from AI providers, where the provider's service replaces functions previously performed by the firm's employees. Third, offshoring: the firm reduces domestic employment while expanding AI-enabled task execution offshore, where the offshore execution does not produce equivalent domestic public revenue. Fourth, attrition and freeze: the firm allows headcount to decline through unreplaced departures while AI systems absorb the corresponding tasks. Fifth, role redesign: the firm formally retains workers in nominal positions while the substantive task performance migrates to AI systems, often accompanied by wage reduction or position downgrade.

In each case, the audit examines functional capability transfer rather than formal employment status. The audit is supported by data the firm already discloses for other purposes: payroll filings, contractor 1099 filings, AI service contracts above threshold, and capability-transfer documentation in routine financial reports.

Cross-border anti-avoidance requires a destination-based supplementary rule. When a domestic firm contracts with a foreign AI provider whose services replace domestic labor, the public revenue gap occurs domestically while the technology rent accrues offshore. A destination-based supplement, modeled on the digital services taxes already enacted in several jurisdictions, can recapture the missing fiscal flow without requiring extraterritorial enforcement. The supplement is calibrated to the labor-side gap, not to the gross value of AI services consumed.

****12. Use of Revenues: From General Revenue to Re-anchoring Funds****

TDL revenue should not disappear into a general fiscal pool. Its legitimacy depends on closing

the loop between taxpayer displacement and social re-anchoring. The revenue should flow into dedicated funds for labor transition, wage insurance, retraining, local public services, and gravity-intensive sectors: elder care, child care, special education, mental health, community health, skilled trades, and community repair.

This is where the TDL paper connects to the broader Age of Dispersion framework. AI displacement disperses workers from prior structures. TDL revenue should help rebuild structures that require human presence. The aim is not to freeze old jobs, but to finance the emergence of new public revenue interfaces where human labor remains socially necessary.

France's household-services tax credit illustrates one partial precedent: public policy can use tax expenditures to support labor-intensive personal services. French official sources describe a 50 percent tax credit for certain services à la personne, subject to limits. [S13] This is not a TDL precedent, but it shows that states already subsidize labor-intensive care and household support when social goals justify intervention.

****13. Implementation Pathways****

The political feasibility of TDL adoption at the U.S. federal level in the near term is low. Falk and Tsoukalas explicitly note that the Pigouvian remedy they identify "faces enormous political obstacles and has no serious legislative momentum in any major economy" [S5]. This paper does not pretend otherwise. Three pathways for incremental implementation are nevertheless available, and each offers a partial test of TDL's viability without requiring federal-level adoption as a precondition.

****13.1 Sub-federal experimentation****

Subject to constitutional nexus, apportionment, federal preemption, and other federal statutory constraints, U.S. state governments could experiment with state-level corporate income tax structures, payroll-related assessments, and unemployment-insurance-style experience rating. A state could add a TDL-style supplemental assessment to its corporate income tax, triggered by the substitution-pattern conditions described in §11.2 and computed against state-level labor-side losses. California, New York, Massachusetts, and Washington each combine substantial corporate tax base, AI-sector concentration, and political receptivity to externality-based tax design, but any state-level design would require careful legal review.

State experimentation has two advantages. First, it permits empirical calibration of trigger thresholds and the Dispersion Coefficient before national-scale adoption. Second, it produces interstate learning: states that adopt TDL successfully demonstrate revenue capacity without major capital flight, while states that do not adopt it can observe administrative costs and incidence effects before deciding whether to follow. This learning dynamic resembles the diffusion of state-level minimum wage increases, paid family leave, and digital advertising tax proposals through the U.S. system.

The principal risk of state-level experimentation is corporate relocation or litigation over nexus and apportionment. A firm subject to TDL in California could in principle attempt to relocate taxable nexus or challenge the assessment under state and federal constraints. This risk is real but bounded. Most large AI-adopting firms cannot relocate their entire operational footprint to escape a state corporate tax supplement, and modern apportionment formulas often tax based partly on sales destination rather than headquarters location. State-level TDL designed with destination-based apportionment and clear nexus rules can mitigate, though not eliminate, relocation and litigation risk.

****13.2 Multilateral pressure through OECD frameworks****

The OECD's BEPS project demonstrates that international tax coordination is institutionally feasible when participating jurisdictions face shared base-erosion concerns. Pillar Two minimum corporate tax represents an extension of this logic to tax rate competition. A potential extension to taxpayer displacement is structurally analogous: jurisdictions that experience labor-side tax base erosion from cross-border AI substitution face a coordination problem identical in form to the profit-shifting problem BEPS addresses.

A TDL-relevant multilateral framework need not require all jurisdictions to adopt identical TDL rates or formulas. It could instead establish minimum standards for AI substitution disclosure, common definitions for substitution-pattern triggers, and rules for cross-border attribution of labor-side losses. This minimum standards approach has precedent in BEPS Action 13 country-by-country reporting and in Pillar Two's GloBE rules.

The probable adoption sequence is the inverse of the technological adoption sequence. Jurisdictions with the largest AI-adopting firms — currently the United States — face the lowest political incentive to lead, because their domestic corporate tax base benefits from AI-related profit increases. Jurisdictions hosting smaller AI-adopting bases but exposed to substantial labor-side tax loss — including parts of the European Union, Canada, Australia, and Japan — face stronger incentives to coordinate. The historical analogue is the Foreign Account Tax Compliance Act: the United States imposed extraterritorial reporting rules unilaterally, and the rest of the world responded with the Common Reporting Standard. A non-U.S. coalition could play the FATCA role in this domain, eventually drawing the United States into a coordinated framework rather than leading one [S11].

****13.3 China adaptation under VAT-CIT architecture****

The China discussion is included because of the author's institutional affiliation with a Beijing-based research institute, not because political adoption is assessed as more likely there than elsewhere. The analysis addresses institutional feasibility only.

The Chinese tax system differs structurally from the U.S. system in ways that affect TDL design but not its underlying logic. China's principal indirect tax is value-added tax (VAT), with standard rates commonly described as 13 percent, 9 percent, or 6 percent depending on sector. China's corporate income tax (CIT) is generally levied at a 25 percent standard rate, with reductions for qualified high-technology enterprises and certain smaller enterprises. Personal income tax is levied at progressive rates up to 45 percent, and social insurance contributions are jointly remitted by employers and employees at locally varying rates. These features mean that a Chinese TDL analogue would need to account for VAT, CIT, PIT, and social-insurance channels rather than simply importing a U.S. payroll-tax model.

China has long used R&D super-deduction mechanisms. Historically, qualifying R&D expenditures were often associated with 175 percent total deduction treatment; later policy expansions made 200 percent total deduction treatment available to many qualifying enterprises. A Chinese TDL analogue should therefore not be framed rhetorically as a new anti-AI tax. It could instead be designed as an inverse adjustment to preferential treatment: when AI-related qualifying expenditures are accompanied by measurable labor-side tax-base displacement in affected job families, the preferential treatment could be reduced, recaptured, or supplemented by a displacement assessment calibrated to the labor-side fiscal loss. [S18]

China's industrial policy tradition and centralized fiscal authority may make technical implementation more administratively feasible than in the U.S. federal system, but political adoption should not be presumed. A Chinese version would require central authorization,

Ministry of Finance and State Taxation Administration rulemaking, coordination with industrial policy agencies, and careful treatment of employment stability goals.

A potential implementation path could begin with centrally authorized local pilots in major AI-intensive jurisdictions, such as Beijing, Shanghai, Hangzhou/Zhejiang, and Shenzhen/Guangdong, coordinated through central fiscal and tax authorities. The piloting framework could borrow from China's experience with carbon-emissions trading pilots and high-technology enterprise certification, both of which used local experimentation before broader rollout. This section argues institutional feasibility, not political endorsement.

****13.4 Sequencing****

The three pathways are not exclusive, and the sequence below is illustrative rather than mandatory. One plausible sequencing places sub-federal experimentation first, multilateral disclosure standards second, and federal or central adoption third. Each pathway generates information that can support the next: sub-federal pilots produce empirical data on trigger calibration, multilateral disclosure standards normalize the audit framework across jurisdictions, and federal or central adoption becomes more plausible only after sufficient empirical and institutional precedent has accumulated.

The historical comparison is to carbon pricing, but only as an analogy of institutional diffusion. Carbon taxes and cap-and-trade systems began as sub-national experiments in northern Europe and the western United States, accumulated empirical evidence over time, and then progressed toward broader multilateral coordination through the European Union Emissions Trading System and the Carbon Border Adjustment Mechanism [S12]. A comparable trajectory for TDL, if it emerges at all, should be expected to span at least a decade.

****14. Institutional Precedents****

TDL is novel, but it is not institutionally alien. At least five existing mechanisms lower the perceived radicalness of the proposal.

First, unemployment insurance experience rating already links employer layoff behavior to employer tax rates. BLS research explains that employers' unemployment insurance tax costs can rise after layoffs, although the marginal cost is often only a fraction of benefits paid. [S9] This is a same-domain precedent: employer behavior that imposes fiscal labor-market costs can affect future tax liability.

Second, workers' compensation and insurance rating systems price risk differentially. Firms whose activities generate higher expected public or social cost can face higher contribution rates.

Third, carbon pricing and carbon border adjustment mechanisms show how public law can price externalities and address cross-border leakage. The European Commission describes CBAM as a mechanism to ensure that imported goods face a carbon price equivalent to domestic production. [S12]

Fourth, FATCA shows that the United States has imposed cross-border reporting and withholding obligations to protect its tax base. IRS materials state that FATCA requires certain foreign financial institutions and other entities to report information about foreign assets held by U.S. account holders or face withholding on withholdable payments. [S11]

Fifth, the OECD/G20 BEPS project shows that tax-base protection can become a multilateral institutional agenda. OECD describes BEPS as strategies by multinational enterprises to exploit

loopholes and shift profits to low- or no-tax locations, and the BEPS Project as providing rules and instruments to address such avoidance. [S10] TDL is not BEPS, but it is BEPS's successor problem: BEPS asks where profits went; TDL asks where taxpayers went. The analogy is structural, not developmental. BEPS represents decades of multilateral negotiation and accumulated political agreement; TDL is a proposal at the earliest conceptual stage of that process.

****15. Objections and Replies****

****15.1 "TDL will block innovation."****

Reply: TDL does not tax AI research, AI assistance, safety automation, or productivity gains as such. It applies only to residual fiscal gaps after offsets. If innovation raises output while retaining or re-taxing workers, no substantial TDL arises.

****15.2 "History shows new jobs will appear."****

Reply: They may. The re-taxation window accounts for that possibility. The paper does not deny job creation; it refuses to treat uncertain future reabsorption as a reason to ignore observable present fiscal loss.

****15.3 "Corporations already pay tax."****

Reply: The relevant question is not whether corporations pay any tax. It is whether incremental corporate-side recovery offsets the labor-side loss. If it does, TDL falls. If it does not, the residual gap remains.

****15.4 "The levy is impossible to administer."****

Reply: TDL is administratively difficult, but not categorically harder than existing tax regimes that already estimate effective tax rates, allocate cross-border income, police transfer pricing, rate unemployment insurance experience, or calculate carbon content in imports. The design should begin with large firms, clear triggers, and auditable job families.

****15.5 "Firms will pass the cost to consumers."****

Reply: Some incidence shifting is possible. That is true of most taxes. The remedy is incidence monitoring, anti-avoidance rules, destination-based supplements, and careful rate design. Incidence concerns affect design, not the existence of the externality.

****16. Conclusion: The Fiscal State Cannot Subsidize Its Own Erosion****

This paper has argued that AI labor substitution is not only an employment problem or a productivity story, but a fiscal problem. When a profitable enterprise removes a worker from its payroll through AI-driven substitution, the fiscal state may lose a public revenue interface: income tax withholding, payroll tax, social-insurance contributions, sales-tax-supported consumption, and local public revenue all flowed through that interface. Their disappearance is not automatically offset by corporate-side tax recovery, because corporate recovery depends on effective tax rates, deduction timing, R&D expensing, capital depreciation, profit allocation, and corporate payout behavior. Their disappearance is not automatically offset by re-employment,

because the re-taxation window is a probability distribution, not a constant.

Two structural observations follow. First, market coordination may fail: individual firms can rationally substitute AI for labor even when the aggregate result weakens the demand base, labor market, or public revenue system on which they indirectly depend. Second, and more fundamentally, sovereign fiscal capacity may fail: enterprises can globalize demand, capital, and technology rents, but states remain responsible for territorial schools, courts, infrastructure, health systems, social insurance, and local services. A firm can thrive across borders while a jurisdiction loses the taxpayers through whom it financed its public obligations.

TDL is the narrow fiscal mechanism this paper proposes for that gap. It does not tax AI as technology, productivity, safety automation, or all displacement. It applies only when a financially strong enterprise carries out AI-linked substitution, the labor-side fiscal loss is not adequately offset by corporate-side recovery and re-taxation window restoration, and the residual gap constitutes an uninternalized externality. The audit equation $L - C - R$, or its dynamic form $\sum_t \delta t(L_t - C_t - R_t)$, is not a false-precision formula. It is a layered audit protocol in which each term is disclosed with assumptions, data sources, and uncertainty intervals.

The deeper claim is structural. The fiscal state of the industrial era was built on a partial coupling of profit and employment: profitable firms tended to employ workers, workers became public revenue interfaces, and the tax code assumed that this relationship would largely hold. Globalization, outsourcing, financialization, and platform work have already weakened the coupling. The AI economy may be among the clearest and most scalable technological regimes in which that coupling can break more systematically. TDL does not solve all distributional problems of the AI economy; it preserves a narrower institutional right: when AI substitution removes taxpayers, the state may ask for the complete public ledger and recover the residual fiscal externality that existing tax law does not capture.

Whether TDL is adopted is ultimately a political question. But the underlying fiscal observation does not depend on adoption. The AI-era fiscal state inherits a tax architecture designed for a different economy, and the gap between that architecture and new substitution patterns is not closing on its own. The fiscal state cannot subsidize the disappearance of its own tax base. That is the minimum institutional lesson of the AI era.

****Appendix A. Cross-Model Collaborative Drafting Method****

This paper emerged through a cross-model collaborative process among a human researcher and multiple AI systems. The appendix records the method rather than celebrating it. The human researcher supplied the original problem frame, normative direction, and final judgment; the AI systems supplied independent structures, critiques, and revisions within that frame. The process had three stages. First, independent generation: multiple AI systems produced competing outlines and policy architectures from the same core thesis. Second, adversarial critique: the models challenged one another on sequencing, formula design, evidentiary risk, model boundaries, and rhetorical overreach. Third, structural integration: the strongest elements were recomposed into a layered architecture — precautionary principle as bone, fiscal ledger as tendon, institutional design as flesh.

This method is not external to the paper. The paper argues that institutions must rebuild structure under conditions of dispersion. Its own production followed a similar pattern: dispersed drafts were subjected to mutual criticism and recomposed into a higher-order structure. The paper's dual-layer architecture — philosophical precaution and fiscal audit — emerged from the collaborative process itself.

****中文版本****

****摘要****

本文提出"纳税人替代税" (Taxpayer Displacement Levy, TDL) 这一 AI 时代的财政制度方案。本文并不主张对 AI 技术本身征税, 也不主张惩罚生产率提升。本文的命题更窄: 当一家财务充裕、持续盈利的企业以 AI 系统替代人类劳动者, 并由此移除一个原本稳定缴税的公共收入接口; 如果劳动侧税收损失没有被企业侧新增税回收补, 也没有在合理的"再纳税窗口"内由被替代者重新进入等价税基位置而恢复, 那么剩余缺口就是一种尚未内部化的财政外部性, 可以成为 TDL 的征收基础。

本文将劳动者定义为"公共收入接口": 劳动者不只是企业利润表上的成本项, 而是现代财政国家的基础设施之一。个人所得税、工资税、社会保险缴费、消费税、地方财政收入, 都通过劳动者这个接口持续流入国家。AI 替代劳动者时, 被移除的不只是工资, 而是附着在这个人身上的一整组公共收入流。

TDL 不同于传统"机器人税"。机器人税征的是机器, TDL 征的是替代行为; 机器人税难以处理软件 AI, TDL 可以处理任何导致纳税人被移出财政循环的技术; 机器人税容易被理解为阻碍创新, TDL 只在财务充裕、必要性不足、社会成本未内部化的 AI 劳动替代情形下触发。

本文的核心审计方程有单期形式和动态形式。单期形式为: $TDL \text{ 应税缺口} = \text{一阶劳动侧税收损失} - \text{企业侧新增税回收补} - \text{再纳税窗口内的预期税基恢复}$ 。动态形式为: $TDL \text{ 应税缺口} = \sum_t \delta^t (L_t - C_t - R_t)$ 。三项具有不同认识论地位: 第一项是会计可观测量, 第二项是税务模型估算量, 第三项是概率分布期望值。TDL 程序并不要求三项具有同等精确性, 而要求企业与税务机关披露各项计算依据、关键假设与不确定性区间。本文以预防原则、现行税制的资本偏置、OBBBA 2025 对资本支出和国内研发费用的处理、Falk-Tsoukalas 的 AI 裁员陷阱模型, 以及 IRS 自身裁员同时扩大 AI 使用的制度悖论为支撑。[S3][S4][S5][S6][S7]

关键词: 人工智能; 自动化税; 纳税人替代; 财政国家; 工资税; 税基侵蚀; 预防原则; AI 劳动替代; 公共收入接口; 再纳税窗口。

****一、引言: 消失的纳税人****

****1.1 公共账本中消失的一项****

设想一家持续盈利的科技公司, 用一套 AI 软件系统替代了一名高级工程师。企业看到的是人力成本下降, 投资者看到的是利润率改善, 管理层看到的是效率提升。但财政国家看到的, 是一个纳税人从公共账本上消失了。这个人不只是企业雇员, 而是联邦所得税、州所得税、工资税、医保税、社保缴费、地方消费和公共财政收入的通道。

关于 AI 劳动替代的常规讨论通常从就业开始。本文从税收开始。就业是社会表层, 税基侵蚀是财政内层。一个劳动者被雇佣、被支付、被预扣、被申报、被纳入社会保险体系, 并嵌入一个具体管辖区。当这个劳动者被替代且不能重新进入等价岗位时, 财政国家失去的不是一笔交易, 而是一个公共收入接口。

这一视角很重要, 因为现代财政依赖连续性。只要被替代者能够转入新的税基位置, 财政国家可以承受失业、流动、裁员和行业转型。但如果 AI 时代的企业以超过经济再吸纳能力的速度移除高质量纳税人, 财政制度就不能继续假装这只是企业内部效率问题。

****1.2 本文的核心贡献****

本文的核心贡献，是把 AI 劳动替代重新定义为“纳税人替代”。既有机器人税讨论常问：机器是否应该缴税？本文问一个不同的问题：当机器或软件系统把一个纳税人从财政循环里移除时，谁来核算这部分公共收入缺口？

本文的答案是纳税人替代税，即 TDL。TDL 不是 AI 技术税，不是创新税，不是安全自动化税。它是一种差额性财政机制：当 AI 相关劳动替代导致劳动侧税基消失，而现有税制无法在企业侧或再就业侧回补这一缺口时，TDL 对剩余财政外部性进行征收。

因此，TDL 更接近失业保险 **experience rating**、碳定价和外部性财政修正，而不是粗糙的技术惩罚。它不是要求企业停止创新，而是要求企业在创新移除纳税人时，把完整公共账本摊开。

****1.3 本文不主张什么****

本文首先作四个否定性说明。第一，本文不主张对所有 AI 使用征税。第二，本文不主张惩罚所有生产率提升。第三，本文不主张所有裁员都是错误或异常。第四，本文不预设所有被 AI 替代的劳动者都会永久离开税基。

本文的命题是有条件、可审计的：当财务充裕企业发生 AI 相关劳动替代，且劳动侧税收损失不能被企业侧新增税收或合理再纳税窗口内的税基恢复所抵消时，剩余缺口可以被视为尚未内部化的财政外部性。

所以，TDL 不是先验惩罚，而是差额审计原则。它不假定所有替代都有害，而是要求企业和税务机关共同把完整账本摊开。

****1.4 核心审计方程****

TDL 有单期审计方程和动态审计方程。

单期形式：TDL 应税缺口 = $L - C - R$

动态形式：TDL 应税缺口 = $\sum_t \delta^t (L_t - C_t - R_t)$

其中， L 是一阶劳动侧税收损失， C 是企业侧新增税收回补， R 是再纳税窗口内的预期税基恢复， δ 是审计跨多个期间展开时使用的折现因子。单期形式适合第一年审计示例；动态形式才是更完整的法律和行政形式，因为纳税人替代和再纳税恢复都发生在时间之中。

这个公式包含三项不同认识论地位的变量。第一项是会计可观测量：给定工资、现行税法 and 管辖区税率，一阶劳动侧损失可以较高精度估算。第二项是税务模型估算量：企业侧回补取决于有效企业税率、折旧、研发费用化、利润转移、回购、分红和时间安排。第三项是概率分布期望值：再纳税窗口不是点值，而是对被替代者以何种速度、何种工资水平、何种税基质量重新进入财政系统的概率估算。

TDL 并不要求三项同等精确。它要求每一项都披露计算依据、关键假设和不确定性区间。这使公式从“假装精确的政策方程”升级为“分层透明的财政审计协议”。

****1.5 TDL 与机器人税的区分****

由于 TDL 在本文中被提出为一种明确针对 AI 相关纳税人替代行为建立审计机制的财政制度方案，读者很可能立刻把它与比尔·盖茨提出的“机器人税”、欧洲议会 2017 年的相关草案动议、以及韩国 2018 年削减机器人相关税收抵免的做法联系起来 [S15]。这种联想可以理解，但有误导性。两类方案在征税对象、覆盖范围、触发条件和设计哲学上都不相同。

征税对象。机器人税针对的是一个有形或技术性人工物：工业机器人、自动化机器，在某些方案中是软件系统。TDL 针对的是一种行为：财务上非必要、由 AI 驱动、移除了劳动侧公共收入接口并产生剩余财政缺口的替代行为。

覆盖范围。基于设备的机器人税难以覆盖纯软件 AI、云端语言模型、第三方托管的智能体系统，以及嵌入工作流工具中的分布式自动化。可征税人工物难以识别，更难归属。TDL 不存在这个问题，因为它不要求识别人工物。它考察的是一种财务和劳动力的组合模式：盈利的企业、AI 资本的扩张、受影响岗位族中人类劳动投入的收缩，以及可度量的劳动侧税收损失。

触发条件。机器人税通常在覆盖设备被使用时即触发。TDL 只在一组条件同时满足时触发：财务强度、AI 相关替代、劳动侧税基移除，以及企业侧回补和再纳税窗口恢复都不足以抵消损失。提升劳动者生产率而非替代劳动者的 AI 使用，不构成 TDL 事件。

设计哲学。机器人税预设自动化整体上有社会危害，并据此定价。TDL 对自动化整体上不做预设。它要求当 AI 替代移除纳税人时，企业把完整公共账本摊开；并且只对抵消之后剩余的财政外部性征税。

这个区分不只关系到政策设计，也关系到政治接受度。机器人税在政治上长期停滞，因为它看起来在惩罚技术进步。TDL 的设计目标恰好是要通过同一个政治测试：它不惩罚 AI，它只对现行税法未捕捉的特定财政模式进行审计。

简言之，机器人税问的是自动化本身是否应被定价；TDL 问的是，在承认抵消项之后，一个纳税人替代事件是否留下了可审计的剩余缺口。

****二、方法论：预防原则的财政化****

本文采用一种财政化的预防原则。它并不宣称 AI 一定会永久摧毁税基，而是认为大规模纳税人替代可能产生高后果、难逆转、事后修复成本极高的财政风险。在这种条件下，国家不必等到损害在统计上完全确定之后，才建立审计机制。

这一方法论很重要，因为 AI 劳动讨论常陷入两个极端：一方断言 AI 将永久消灭大量工作，另一方用历史经验回应“技术总会创造新工作”。这两个说法对公共财政都太粗。税制不能依赖恐慌，也不能依赖历史乐观主义。它需要一套程序，用来测量替代、回补、恢复和剩余缺口。

财政化预防原则有三层：第一，识别一个一旦大规模移除就可能削弱财政能力的公共接口；第二，建立区分即时损失、抵消回补和不确定恢复的审计方法；第三，只对抵消之后仍然存在的剩余外部性征税。这样，预防原则就不是口号，而是程序。

****三、财政国家的隐形前提：劳动者作为公共收入接口****

现代国家并不只是对经济产出征税，它对处在经济角色中的人征税。劳动者是财政系统中最稳定、最透明的接口。工资被申报，被预扣，被雇主匹配，并与社会保险缴费相连。与资本利得、合伙企业收入、全球无形资产收入或跨国利润相比，工资收入对财政国家高度可见。

因此，劳动者不是简单成本项，而是公共收入接口。移除这个接口，影响不止发生在企业内部。国家可能失去所得税预扣、雇主与雇员工资税、医保税、失业保险缴费、地方所得税与销售税效应，以及工资支撑的消费活动。

这个定义是功能性的，不是评价性的。把劳动者描述为公共收入接口，并不是把人降格为财政工

具，而是把一个已经存在的公共财政依赖关系显现出来。现代财政系统通过与劳动相连的持续收入流来维持学校、道路、法院、医疗体系、社会保险和地方公共服务。如果 AI 替代在规模上移除了这些收入流，财政国家就有权追问：这个缺失的接口去了哪里？

以美国为例，IRS 2026 年 Publication 15 显示，社保税率为雇员与雇主各 6.2%，适用于 184,500 美元工资基数上限以内的工资；医保税率为雇员与雇主各 1.45%，且没有工资基数上限。这些工资税流不是附带事项，而是附着在劳动者身上的财政结构。[S8]

企业替代的是任务；财政国家失去的是接口。这就是本文的铰链。

****四、这次为什么可能不同：出口窗口、再纳税窗口与风险****

反对 TDL 最强的理由是历史经验：技术革命会摧毁一些岗位，也会创造另一些岗位。这个反驳很严肃，本文不否认它。本文的问题是：在 AI 时代，出口窗口是否足够可靠，以至于财政国家可以忽略这种风险？

过去的技术革命往往替代较窄的能力维度。机械化替代某些体力劳动，蒸汽动力重组生产，计算机自动化计算和文书流程。在许多情况下，被替代者进入新的生产循环和纳税循环。纳税人被转移，而不是被移除。

AI 作为风险类别的不同之处在于，它潜在地能够瞄准广谱认知功能：写作、编程、分析、客服、翻译、设计、排班、法律草拟、软件测试、中层管理和部分决策支持。这并不证明永久替代，但说明出口路径可能与替代前沿高度重叠。如果新岗位仍然要求那些正在被 AI 快速降价的认知能力，再吸纳机制就不能被当然假定。

****4.1 再纳税窗口****

****再纳税窗口，是指被 AI 替代的劳动者在合理时间内重新进入等价或近似等价公共收入接口位置的概率区间、时间区间和税基恢复幅度。它有三个维度：重新就业所需时间、重新就业质量，以及税基恢复概率。三个月后以原工资 95%重新就业，与两年后以原工资 40%从事不稳定工作的财政意义完全不同。****

这个概念避免了二元化错误。问题不是被替代者是否会再次工作，而是他是否以可比速度和可比质量恢复税基。一个低薪替代岗位可能恢复尊严和部分收入，但未必恢复被移除的公共收入接口。

****4.2 没有技术宿命论的预防原则****

因此，本文把 AI 劳动替代视为预防性财政风险，而非已经证明的永久灾难。如果再纳税窗口保持开放且强劲，TDL 自然下降或消失。如果窗口变窄、延迟或降级，应税缺口扩大。于是，历史乐观主义不再是对框架的反驳，而是公式中的一个变量。

****五、税基蒸发：从一阶损失到动态财政缺口****

****5.1 一阶劳动侧损失****

****一阶劳动侧损失包括与被替代劳动者薪酬和消费直接绑定的税收与缴费。至少包括联邦所得税、州和地方所得税、雇员工资税、雇主工资税、失业保险缴费，以及消费税支持的财政活动。具体清单因管辖区而异，但机制具有普遍性。****

在美国案例中，仅工资税部分就非常重要。劳动者和雇主共同缴纳社保与医保工资税，其中社保税有工资基数上限，医保税没有上限。高收入劳动者还会产生显著所得税预扣和地方消费效应。这些不是抽象社会价值，而是可测量的财政流。

****5.2 企业侧税收回补****

企业可能会说：人力成本下降会增加应税利润。这个说法必须被计入。TDL 如果忽视企业侧回补，就不成立。但企业侧回补并不确定。它取决于有效企业税率、扣除、加速折旧、研发费用化、亏损结转、国际利润配置、股权激励、股票回购、分红和时间安排。

因此，本文并不假定企业税一定补不回来。本文要求证明它补回来了多少。企业侧回补是模型估算，不是口号。

****5.3 再纳税窗口内的预期恢复****

第三个抵扣项是预期税基恢复。如果被替代劳动者很快重新进入等价岗位，长期财政缺口会缩小。如果他进入低薪、不稳定或非 W-2 工作，恢复就是部分的。如果他退出劳动力市场、长期低就业或转入低申报收入，恢复就很弱。

****再纳税窗口可依据年龄、职业、地区、行业、技能迁移性、工资替代率、失业持续时间和地方劳动力吸纳能力估算。这个值是概率性的，但概率性不等于不可用。保险、信用评分、失业预测、养老金估值和气候风险金融都依赖不确定分布。TDL 应按公共财政标准评估，而不应被要求达到不可能的完美预见。****

****5.4 TDL 财政缺口方程****

TDL 有单期表达式和动态表达式。

单期形式： $TDL \text{ 应税缺口} = L - C - R$

动态形式： $TDL \text{ 应税缺口} = \sum_t \delta^t (L_t - C_t - R_t)$

其中， L 代表一阶劳动侧税收损失， C 代表企业侧新增税收回补， R 代表再纳税窗口内的预期税基恢复， δ 代表折现因子。正的剩余值表示尚未内部化的财政外部性。零值或负值表示不应征收 TDL，或仅作极低水平处理。

这个区分很重要。第一年审计可以显示某个替代事件是否产生即时公共收入缺口；多年期审计可以判断这个缺口是否随着劳动者重新进入税基而持续、缩小或消失。因此，TDL 义务应当是年度性、可复审的，而不是附着在一次替代事件上的永久惩罚。

这个结构同时保护公平和创新。企业不会因生产率提升本身被惩罚。它们被要求核算的是 AI 替代移除的公共收入接口。如果财政系统已经在其他地方恢复了收入，TDL 自然消失；如果缺口仍在，TDL 填补缺口。

****5.5 三个算例：L - C - R 的应用****

为说明核心审计方程产生的是可处理的数字而非口号，本节对三个不同收入水平的被替代劳动者计算第一年 TDL 审计缺口。所有算例采用 2026 年美国联邦税法，为简化忽略州和地方税；纳入州级税收一般会扩大劳动侧损失。所有数字均为示例性估算，用于说明机制，不代表已审计法定数值。

以下算例报告的是第一年审计值。多年期 TDL 义务应按照§5.4 的动态方程逐年计算，并随着再纳

税窗口把劳动者恢复到等价或近似等价公共收入接口而终止或下降。

****税收假设。**算例使用 2026 年联邦个人所得税参数和 2026 年 184,500 美元的社保工资基数。算例纳入雇员和雇主社保税、医保税；在高收入算例中，纳入超过适用门槛的雇员工资 0.9% 的 **Additional Medicare Tax**；该附加医保税没有雇主匹配部分。[S8][S16][S17]**

****企业侧回补注意事项。**算例使用简化的同年扣除等价处理来说明时间效应。如果实现 AI 替代的投资属于合格可折旧财产，**Section 168(k)**可能降低同年应税收入；如果属于国内研究或实验支出，**Section 174A**可能适用；如果属于云服务、供应商、API 或 SaaS 付款，普通经营费用扣除或资本化规则可能适用。本文并不主张所有 AI 支出享有相同处理，而是说明同年扣除和费用化可能在替代发生当年显著压低企业侧税收回补。[S3][S4]**

****再纳税窗口参数化。**第一年示例中，R 估算为： $R = P1 \times W \times Q \times T \times L$ 。其中，P1 是第一年回到可比工作的概率，W 是工资替代率，Q 是新工作安排的税基质量系数，T 是第一年平均时间系数。真实审计中，这些参数应按职业、年龄、地区和行业具体估计。**

****算例 A：中等技能服务业岗位，年薪 48,000 美元。****

****一阶劳动侧损失（L）。**联邦所得税预扣约 3,400 美元。雇员工资税约 3,672 美元；雇主对等工资税约 3,672 美元。联邦失业税约 42 美元。估算消费税支持的间接公共收入按保守口径取 1,200 美元。第一年 L 总额约为 12,000 美元。**

****企业侧回补（C）。**企业节省工资 48,000 美元和福利及雇主工资税约 5,200 美元，劳动力成本下降约 53,000 美元。假设全部计入应税利润，按 21% 联邦企业税率，新增企业税约 11,100 美元。若 40,000 美元用于实现替代的 AI 相关支出获得同年扣除处理，其时间性优惠使企业税减少约 8,400 美元。第一年 C 净值约为 2,700 美元。**

****再纳税窗口恢复（R）。**对于都市劳动力市场中的 35 岁服务业工作者，示例性第一年假设为：12 个月内以原工资 90% 回到可比岗位的概率为 60%，税基质量为普通 W-2，平均时间系数为半年度。 $R \approx 0.60 \times 0.90 \times 1.00 \times 0.50 \times 12,000 \approx 3,200$ 美元。**

第一年剩余 TDL 缺口 = $L - C - R \approx 12,000 - 2,700 - 3,200 \approx 6,100$ 美元。

****算例 B：中职业期专业岗位，年薪 150,000 美元。****

****一阶劳动侧损失（L）。**联邦所得税约 24,000 美元。雇员工资税约 11,475 美元；雇主对等工资税约 11,475 美元。间接消费相关收入约 4,000 美元。第一年 L 总额约为 51,000 美元。**

****企业侧回补（C）。**工资和福利节省约 172,000 美元，按 21% 税率，新增企业税约 36,100 美元。若 80,000 美元用于实现替代的 AI 相关支出获得同年扣除处理，其时间性优惠使企业税减少约 16,800 美元。第一年 C 净值约 19,300 美元。**

****再纳税窗口恢复（R）。**对于 42 岁、所在行业暴露于 AI 替代前沿的专业人员，示例性第一年假设为：第一年回到可比工作的概率为 15%，工资替代率 70%，税基质量为普通 W-2，平均时间系数为半年度。 $R \approx 0.15 \times 0.70 \times 1.00 \times 0.50 \times 51,000 \approx 2,700$ 美元。**

第一年剩余 TDL 缺口 $\approx 51,000 - 19,300 - 2,700 \approx 29,000$ 美元。

****算例 C：高级工程师，年薪 500,000 美元。****

****一阶劳动侧损失（L）。**按简化的单身申报者估算，联邦所得税约 138,000 美元。雇员工资税约 21,400 美元，包括工资基数上限内的社保税、全部工资的医保税，以及超过适用门槛部分 0.9% 的 **Additional Medicare Tax**。雇主工资税约 18,700 美元，因为 **Additional Medicare Tax** 没有雇主匹配。间接消费相关收入约 12,000 美元。第一年 L 总额约为 190,000 美元。

[S16][S17]**

****企业侧回补（C）。工资和福利节省约 580,000 美元，按 21%税率，新增企业税约 121,800 美元。若 250,000 美元用于实现替代的 AI 相关支出获得同年扣除处理，其时间性优惠使企业税减少约 52,500 美元。第一年 C 净值约 69,300 美元。****

****再纳税窗口恢复（R）。对于 48 岁、专业方向接近 AI 替代前沿的高级工程师，示例性第一年假设为：第一年回到可比工作的概率为 8%，工资替代率 60%，税基质量为普通 W-2，平均时间系数为半年度。 $R \approx 0.08 \times 0.60 \times 1.00 \times 0.50 \times 190,000 \approx 4,600$ 美元。****

第一年剩余 TDL 缺口 $\approx 190,000 - 69,300 - 4,600 \approx 116,100$ 美元。

如何阅读这些算例。第一，三个第一年示例中均出现正剩余缺口，即使已经承认企业侧回补和第一年预期税基恢复。第二，缺口随工资水平非线性扩大，因为累进个人所得税把更多公共收入附着在高工资劳动者身上，而企业侧回补受统一企业税率和同年扣除的限制。第三，结果对再纳税窗口高度敏感：更强的第一年再就业概率、更高的工资替代率或更高税基质量恢复都会缩小剩余缺口。第四，这些算例不是拟议法定税率，而是审计演示，说明 L、C 和 R 可以被估算、争辩和校准。第五，三个示例中均出现的正剩余缺口反映的是此处选择的参数假设。审计方程在设计上方向中立：更强的企业侧回补、更强的第一年再就业概率，或更高的工资替代率都可以产生零或负剩余——在此情况下 TDL 不适用，或仅以最低限度水平适用。

真实 TDL 审计需要企业特定数据、管辖区特定税收参数和职业特定再纳税窗口估计。本节唯一目的，是说明这个方程产生的是可处理、可争辩的数字——而这正是一套审计框架所需要的。

****5.6 风格化的宏观暴露包络线：一个数量级示意****

§5.5 中的三个算例表明， $L - C - R$ 在个体被替代层面可以产生可处理的数字。本节将同一逻辑扩展到国家规模，但只作为“暴露包络线”使用。以下数字是数量级示意，不是对 TDL 税收收入的预测。它们展示的是在明示假设下 AI 相关替代可能暴露出的公共收入接口规模，而不是任何具体实施方案下政府实际会征收到的金额。

起点是 MIT 与橡树岭国家实验室于 2025 年 10 月发布的 Iceberg Index。该指数估算，当前 AI 系统在技术上有能力执行相当于美国劳动力市场 11.7% 的任务，对应 9.4 万亿美元总劳动力市场中的约 1.2 万亿美元暴露工资。[S19] Iceberg Index 衡量的是技术暴露，不是已实现替代、就业损失或时间表。同一研究还报告，迄今为止可见的部署约为 2,110 亿美元工资价值，主要集中于计算和技术职业。

第一层：假设性全量替代上限。

如果整个 1.2 万亿美元暴露工资基数在单一年度内全部被替代，一阶劳动侧公共收入损失将等于暴露工资基数乘以一个综合有效税率假设。该假设覆盖联邦所得税、雇主和雇员工资税、州和地方所得税，以及消费税支持的公共收入。§5.5 算例给出的 L/工资比值约为：低薪 25%、中薪 34%、高薪 38%。由于 Iceberg 暴露偏向高薪认知职业 [S19]，被暴露工作的加权比值位于该区间较高位置是合理的。因此，本稿采用 28% 至 35% 的综合率假设，得到约 3,400 亿至 4,200 亿美元的反事实全暴露财政包络线。CBO 2022 年分配数据支持较窄的一点：联邦有效税率随收入上升；但它并不直接测量本文所用的联邦、州地、雇主侧和消费税综合率。[S20]

第二层：扣除企业侧回补和第一年再纳税之后的剩余暴露。

§5.5 算例产生的第一年剩余缺口分别约为对应 L 的 51%（算例 A）、57%（算例 B）和 61%（算例 C）。将 50% 至 60% 的代表性区间应用于第一层，得到第二层剩余暴露约为 1,700 亿至 2,500 亿美元；这仍然只是在整个暴露工资基数一年内被替代这一反事实假设下成立。第二层不是预测，而是一个压力测试包络线，用于显示在承认抵消项之后剩余外部性的规模。

现实性限定。

第一层和第二层都不是预测。两者都假设整个 1.2 万亿美元暴露工资基数在单一年度内被替代，这显然不现实。实际年度替代取决于采纳速度、AI 部署资本成本、监管摩擦、劳动合同结构、管理激励和劳动者再培训能力。在一个纯示例性的年度实现率假设下，如果每年实现暴露工资基数的 5% 至 15%——这不是经验采纳估计，也不是预测——早期年度剩余财政缺口大致在每年 90 亿至 380 亿美元之间。这些数字应被理解为在明示假设下展示结构性规模，而不是预测税收收入。暴露包络线确立了什么。

由此可得出三点观察。第一，一旦技术暴露工资基数达到万亿美元级别，即使实际年度替代只实现一个较小比例，也可能对公共财政具有宏观相关性。第二，第一层与第二层之间的差距本身具有意义：在这些风格化假设下，承认企业侧回补和第一年再纳税，会将总暴露减少约 30% 至 45%。第三，政策问题因此不是政府能否立刻征收数千亿美元，而是财政暴露变成已实现收入冲击之前，是否应当先建立审计架构。

以上估算使用 2025 年末的暴露数据和风格化税收参数，二者都将变化。分层结构的意义在于：在明示假设下显示 TDL 能够产生可处理的总量暴露估算，而不是锁定任何特定数值。

****六、现行税法与资本偏置****

当现行税法不但没有纠正替代激励，反而偏向资本替代时，TDL 的论证会更强。2025 年 OBBBA 的相关变化正说明这一点。IRS 指导文件显示，OBBBA 对 Section 168(k) 的修订，为 2025 年 1 月 19 日之后取得并投入使用的合格财产提供永久 100% 第一年额外折旧。[S3] IRS 关于 Section 174A 的指导也说明，国内研究或实验支出可以扣除，并规定了此前资本化金额的过渡处理。[S4]

这些规定本身并非错误。投资、研发、软件和基础设施都很重要。财政不对称发生在这样的条件下：AI 资本支出获得有利的时间性税收处理，而雇佣人类劳动则持续触发工资税、福利、合规义务、劳动法责任和长期薪酬承诺。税法由此可能放大企业替代人类劳动的私人激励，即使公共财政效果为负。

本文并不主张培训支出从不允许扣除。许多普通经营支出当然可以扣除。更强的命题是结构性的：机器、软件和研发投入可以成为现金流优势资产，而劳动者则仍然是反复触发税收和义务的接口。为工业投资设计的税制，在 AI 经济中可能无意中补贴了纳税人替代。

****6.1 非营利和公益公司（PBC）结构****

部分 AI 价值链同时包含免税非营利实体、应税子公司、公益公司（PBC）、利润上限安排和普通商业供应商。这些组织形式在法律上完全合法，且许多早于 AI 经济出现。PBC 本身并不是免税实体，而是一种公司治理形态。财政问题出现在 AI 价值链混合了不同税收暴露的实体，并且下游使用该技术移除劳动侧税基时。本文不指控任何特定实体存在不当行为。

本文的财政观察是结构性的，不是指控性的。如果一个享受税收优惠的研究实体、应税子公司、PBC 或普通供应商向下游盈利企业提供 AI 能力，下游企业利用该能力从自身员工中移除纳税人，那么劳动侧税收损失发生在下游财政位置，而部分技术租金可能在另一个税收处理不同的位置累积。

这种结构特征并非 AI 独有。在非营利基金会下进行的药品研究、由联邦资助研究中心进行的国防研究，以及通过税收优惠管辖区路由的金融活动，都会产生类似的结构性问题。AI 案例值得明确关注，是因为纳税人替代速率可能更高，并且价值链可以跨越具有不同税收待遇的多种组织形

式。

就 TDL 设计而言，含义是有限的。征税对象仍然是替代行为，而非价值链中任何参与者的组织形式。触发 TDL 的下游企业按 TDL 缴费，无论其 AI 供应商是非营利、利润上限制、PBC、普通公司还是完全应税供应商。AI 供应商的组织形式既不影响触发，也不影响计算。这一设计是有意为之：TDL 回避“任务驱动型 AI 组织是否应保留优惠地位”这一政治上有争议的问题，而是在下游应用统一的替代行为测试。

****七、结构性异常的定义：盈利期 AI 相关裁员****

盈利期裁员并不是新现象。企业长期以来会在盈利时裁员，以便重组、退出业务线、应对战略变化或提高利润率。TDL 处理的异常更窄：盈利期裁员 + AI 资本支出同步扩张 + 被裁岗位能力维度被 AI 接管。

因此，TDL 触发条件应由一组因素构成：企业达到规模和财务强度门槛；发生显著劳动缩减、岗位关闭或招聘冻结；部署或扩展能够执行被替代任务的 AI 系统；AI 相关资本或运营支出增长；且企业无法证明生存必要、安全必要或公共利益必要。

异常不是道德失败，而是财政事件。财务强劲的企业完全可能理性地替代劳动者。但如果这种替代移除了纳税人并使公共账本出现缺口，国家就有理由审计这笔交易。

****八、市场协调失败：AI 裁员陷阱****

市场失败支柱来自 AI 替代的竞争逻辑。Falk 和 Tsoukalas 在 2026 年预印本《The AI Layoff Trap》中建立了一个竞争性任务模型，说明需求外部性可能使理性企业陷入自动化军备竞赛，替代超过集体最优水平的劳动者。作者认为，在其模型中，资本税、工人持股、UBI、技能提升和科斯谈判都不能消除这一外部性，而庇古式自动化税可以直接作用于替代边际。[S5]

这个模型必须谨慎使用。它不是所有经验情形的证明。它假设劳动者收入与企业需求之间存在一定关系，而这种关系在不同经济体和行业中强弱不同。例如，一家全球化企业可能主要向裁员所在国之外的消费者销售。这个事实会削弱简单的国内需求反馈叙事。

但该模型仍然有用，因为它说明了为什么单个企业即使知道总体风险，也未必会自我约束。竞争会把谨慎变成劣势。延迟替代的企业可能损失利润率、投资者或市场份额。竞赛由此变成个体理性、整体风险。

****九、主权财政失败：全球需求与本地税基****

TDL 更强的支柱不是市场失败，而是主权财政失败。企业可以全球化需求，国家不能全球化它失去的纳税人。学校、道路、法院、医院、治安、养老、失业保险、残障救助和地方基础设施，仍然是属地责任。财政国家承担领土内义务，即使资本和需求跨境流动。

因此，TDL 不能只依赖“企业最终会伤害自己的需求基础”这一说法。有些企业未必会。美国企业可以裁掉美国员工，同时向全球消费者销售；平台可以压缩本国就业，同时从境外用户身上变现；模型供应商可以通过跨境订阅和云合同捕获价值。地方税基可能侵蚀，而企业仍然繁荣。

所以核心问题不只是自动化是否损害总需求，而是一个管辖区是否失去了公共收入接口，而相应生产率收益却被其他地方捕获。这就是主权财政缺口。TDL 正是对这个缺口的回应。

一句话：企业可以全球化它的需求，但国家不能全球化它失去的纳税人。

****十、IRS 悖论：替代收税者，同时失去税基****

IRS 悖论是本文最具体的制度警告。Yale Budget Lab 在 2026 年 4 月估算，IRS 资金回撤和裁员可能导致 2026—2035 年约 8610 亿美元收入减少，其中类似 DOGE 行动造成的裁员部分约 5978 亿美元，资金回撤部分约 2628 亿美元。该分析还指出，到 2025 年中，IRS 已经失去超过 3,600 名 revenue agents，约占审计人员 31%。[S6]

与此同时，GAO 报告显示，IRS 已将 AI 用于审计选择、纳税人问答等业务；截至 2025 年 6 月，IRS 库存中有 126 个活跃 AI 用例，而 2022 年 8 月报告的数量为 10 个。GAO 同时警告，人员削减导致 IRS 缺乏足够熟练员工来支持或开发新的 AI 工具。[S7] Reuters 也报道，人员削减后税务执法弱化，执法人员减少，审计数量下降。[S14]

悖论并不在于 IRS 使用 AI 本身是坏事。AI 可能帮助发现不合规并改善运营。悖论在于：一个财政国家可能削弱征税的人类基础设施，同时想象 AI 能够补偿。如果负责收税的机构也把人的专业能力视为无需核算财政后果即可替代的对象，它本身就成了纳税人替代逻辑的活案例。

****十一、纳税人替代税的制度设计****

以下设计应被理解为法定架构示意，而非最终立法文本。其触发门槛、系数和最长时间都是校准占位值，用来说明 TDL 如何可能被行政执行。真正的法律文本需要试点数据、宏观经济校准、行业咨询和定期修订。

****11.1 征税对象：替代行为****

TDL 不征 AI 系统，也不征机器。它征的是一种行为：财务上非必要的 AI 相关替代，移除了劳动侧税基，并留下剩余财政缺口。这一区分极其重要。用于辅助劳动者、提高安全性、减少错误或在不移除税基接口的情况下扩大产出的模型，不应触发 TDL。

****11.2 触发门槛****

企业只有同时跨越三类门槛时，才进入 TDL 审查。每一类门槛都应由国家或管辖区监管机构根据宏观经济条件校准。以下数值不是拟议的最终法定值，而是用于说明 TDL 审查入口如何构造的示例性校准占位值。本稿中所有定量门槛均应被理解为示例性数值，须经经验校准后才能进入法定用途。

规模门槛。上一财年总营收超过 10 亿美元（示例性），或全球雇员人数超过 5,000 人（示例性）且在征税管辖区雇员人数不少于 1,000 人。较小的企业默认豁免。这一目标是把执法集中在具备规模化替代财务能力和合规审计行政能力的实体。

财务强度门槛。企业在追溯三个财年内显示持续盈利，定义为三年中至少两年净利润为正且触发年经营性现金流为正。或者，企业持有现金及等价物超过 12 个月经营支出，触发年实施股票回购或特别股息，或自由现金流相对于工资义务异常强劲。困境企业、积极重组中的企业，以及持续经营亏损的企业默认豁免。单纯市值不应触发 TDL 审查，因为估值可能波动，并且可能与经营现金能力脱节。

替代模式门槛。触发年 AI 相关资本和经营支出超过企业三年平均水平一个示例性的 25%，且企业在受影响岗位族中通过裁员、招聘冻结或不补充自然减员减少劳动力一个示例性的 5%，且触发期部署的 AI 系统可证明执行被替代岗位族中的任务。这一合取结构防止偶然关联：增加 AI 支出同时扩大招聘的企业不触发；裁员但无对应 AI 部署的企业也不触发。

这些门槛应每年审查。随着 AI 部署常规化，规模和替代门槛应调整以反映变化中的基线。目标不是覆盖每一个替代事件，而是识别现行税法未充分处理的“财务强劲、AI 驱动、替代模式”这一子集。

****11.3 必要性测试****

企业可以通过证明必要性来反驳 TDL。生存必要适用于企业若不替代将面临可信的破产、严重收缩或基本运营丧失。安全必要适用于自动化把人从危险任务中移除。公共利益必要适用于替代直接改善公共福利，例如降低医疗错误或防止基础设施灾难。

****11.4 举证责任倒置****

由于企业掌握 AI 部署记录、生产率数据、劳动力规划和成本预测，一旦触发条件满足，举证责任应当倒置。国家不必从外部证明每一个内部动机。财务强劲企业发生 AI 相关劳动力缩减时，应当证明必要性，或接受基于剩余财政缺口的 TDL 评估。

****11.5 税率结构****

TDL 基础税额等于审计方程计算的剩余财政缺口，跨受影响岗位族和审计年度求和。基础税额随后由“离散系数”（Dispersion Coefficient）调整，该系数反映替代事件周边的社会和劳动力市场条件。如果面向公共经济学期刊，也可以更技术化地称为“劳动力市场摩擦系数”（Labor-Market Friction Coefficient）。

离散系数由四个因素构成。第一，地区再就业弹性：当地劳动力市场对可比职业被替代劳动者的历史吸纳能力，由先前失业保险申请队列的 12 个月再就业率衡量。第二，职业流动性：被替代劳动者技能的可迁移性，使用公共劳动力市场数据库的职业相似性分数估算。第三，年龄结构暴露：50 岁以上被替代劳动者的占比，他们系统性地面临更长的再纳税窗口。第四，社区依赖度：被替代岗位在某些城镇或地区的集中度，这些地区单一雇主占地方工资税基显著比例。

为说明机制，系数可示例性设置为 0.7 至 1.5。这一区间不是拟议法定区间，需要经验校准。低于 1.0 的值适用于被替代劳动者很可能在弹性劳动力市场快速重新进入等价岗位的情形。高于 1.0 的值适用于在有记录吸纳困难的地区、职业或年龄队列中发生的替代。系数不是惩罚性的；它反映的是“财政外部性大体上可自我修正的替代事件”与“外部性具有粘性的替代事件”之间的差异。

TDL 应纳税额构造为一项年度持续义务，持续至记录的再纳税缺口期间，而非一次性税款。这种处理方式契合潜在财政现实：公共收入接口的移除不是一年，而是直到等价再就业发生为止。当审计数据显示被替代劳动者的再纳税已大体恢复，或达到类似五年的最高期间（取较早者），义务终止。最高期间本身也应经验校准。

****11.6 反规避规则****

TDL 必须覆盖功能性替代，不仅仅是形式裁员。一家企业减少直接 W-2 雇员同时扩大 AI 驱动的承包商关系、第三方 AI 服务合同、海外 AI 赋能劳动力，或先冻结招聘再让自然减员侵蚀编制——其公共收入接口足迹的减少与执行明确裁员的企业一样有效。反规避规则必须触及所有这些模

式。

一般反规避测试询问：触发年中，某岗位族的人类劳动投入和相应劳动侧税基是否下降；以及企业本身或为企业提供 AI 服务的供应商所部署的 AI 系统，是否接管了相关能力维度。若两个条件同时满足，无论合同形式如何，该替代在功能上属于 TDL 相关。

具体反规避规则至少应覆盖以下模式。第一，承包商替代：企业在某岗位族中减少 W-2 雇员，同时增加对独立承包商的支付，承包商以 AI 工具增强执行类似职能。第二，供应商替代：企业减少内部劳动力，同时扩大对 AI 提供商的合同服务采购，提供商的服务替代了原本由企业员工执行的职能。第三，离岸化：企业减少国内雇佣，同时扩大海外 AI 赋能任务执行，海外执行不产生等价的国内公共收入。第四，自然减员与冻结：企业允许人员通过未补充的离职下降，同时 AI 系统吸收相应任务。第五，岗位重设计：企业在名义岗位上保留员工，而实质性任务执行迁移到 AI 系统，通常伴随工资下调或职位降级。

每种情形下，审计考察的是功能性能力转移，而非形式雇佣状态。审计依靠企业已为其他目的披露的数据：工资申报、承包商 1099 申报、超过门槛的 AI 服务合同，以及常规财务报告中的能力转移文件。

跨境反规避需要一个目的地原则补充规则。当国内企业与海外 AI 提供商签约，其服务替代国内劳动力，公共收入缺口发生在国内，而技术租金累积在海外。基于目的地原则的补充规则，可以参照已在若干管辖区实施的数字服务税模式，捕获缺失的财政流量，而无需域外执法。补充规则按劳动侧缺口而非 AI 服务消费总值校准。

****十二、税款用途：从一般财政到重锚基金****

TDL 收入不应消失在一般财政池中。其正当性取决于能否把纳税人替代和社会重锚闭环连接起来。TDL 收入应进入专门基金，用于劳动力转型、工资保险、再培训、地方公共服务以及引力密集型行业：养老、托育、特殊教育、心理健康、社区医疗、技工和社区修复。

这正是 TDL 论文与《离散时代》大框架连接的地方。AI 替代使劳动者从旧结构中离散，TDL 收入应帮助重建需要人在场的新结构。目标不是冻结旧岗位，而是资助新的公共收入接口在社会真正需要人的地方生长。

法国家政服务税收抵免是一个局部参照。法国官方资料显示，符合条件的"services à la personne"可享受 50% 的税收抵免，并设有上限。[S13] 这不是 TDL 先例，但它说明国家已经会在社会目标需要时，用税收支出支持劳动密集型照护和家庭服务。

****十三、实施路径****

近期内在美国联邦层面通过 TDL 的政治可行性较低。Falk 和 Tsoukalas 明确指出，他们识别的"底古式补救措施"在任何主要经济体都面临巨大的政治障碍，没有严肃的立法动力" [S5]。本文不假装相反。然而，存在三条增量实施路径，每条都为 TDL 的可行性提供部分检验，而无需以联邦层面采纳为前提条件。

****13.1 次联邦层级实验****

在满足宪法联结、分摊、联邦优先以及其他联邦法律限制的前提下，美国州政府可以对州一级企业所得税结构、工资相关评估和类似失业保险 experience rating 的机制进行实验。一个州可以在

企业所得税中增加一项 TDL 类型的附加评估，由§11.2 描述的替代模式条件触发，并以州一级劳动侧损失计算。加州、纽约、马萨诸塞和华盛顿各自结合了较大的企业税基、AI 行业集中度，以及对外部性税收设计的政治接受度，但任何州级设计都需要严肃法律审查。

州一级实验有两个优势。第一，它允许在国家规模采纳之前对触发门槛和离散系数进行经验性校准。第二，它产生州际学习效应：成功采纳 TDL 的州在不发生重大资本外流的情况下展示财政能力，而未采纳的州可以先观察行政成本和税负归宿效果，再决定是否跟进。这与美国系统中州一级最低工资上调、带薪家庭休假和数字广告税提案扩散的动力机制相似。

州一级实验的主要风险是企业搬迁或围绕联结与分摊规则的诉讼。在加州被 TDL 覆盖的企业原则上可以尝试迁移应税联结，或依据州和联邦约束挑战该评估。这一风险真实但有限。大多数大型 AI 采用企业无法迁移整个运营足迹来逃避州企业税附加，而现代分摊公式通常部分按销售目的地而非总部所在地征税。基于目的地分摊和清晰联结规则设计的州一级 TDL，可以缓解但不能消除搬迁和诉讼风险。

****13.2 通过 OECD 框架的多边压力****

OECD 的 BEPS 项目表明，当参与管辖区面临共同的税基侵蚀关切时，国际税收协调在制度上可行。第二支柱企业最低税代表了这一逻辑向税率竞争的延伸。向纳税人替代的潜在延伸在结构上类似：经历跨境 AI 替代造成劳动侧税基侵蚀的管辖区，面临的协调问题在形式上与 BEPS 处理的利润转移问题相同。

TDL 相关的多边框架不必要求所有管辖区采纳相同的 TDL 税率或公式。它可以建立 AI 替代披露的最低标准、替代模式触发的共同定义，以及劳动侧损失的跨境归属规则。这种最低标准方法在 BEPS Action 13 的国别报告和第二支柱的 GloBE 规则中已有先例。

可能的采纳顺序是技术采纳顺序的反向。拥有最大 AI 采用企业的管辖区——目前是美国——领先采纳的政治激励最低，因为其国内企业税基受益于 AI 相关利润增加。AI 采用基础较小但暴露于实质性劳动侧税收损失的管辖区——包括欧盟部分地区、加拿大、澳大利亚和日本——有更强的协调激励。历史类比是《海外账户税收合规法案》（FATCA）：美国单方面施加域外申报规则，世界其他国家以《共同申报准则》回应。一个非美国联盟可以在这个领域扮演 FATCA 角色，最终把美国引入协调框架而非由其领导 [S11]。

****13.3 中国情境下的 VAT-CIT 架构适配****

本节讨论中国情境，原因是作者所在机构设于北京，而非因为政治层面的采纳概率被评估为高于其他地区。本节的分析范围仅限于制度可行性。

中国税制在结构上不同于美国税制，这种差异影响 TDL 设计但不影响其底层逻辑。中国主要的间接税是增值税（VAT），标准税率通常按行业分为 13%、9% 或 6%。中国企业所得税（CIT）一般按 25% 标准税率征收，对符合条件的高新技术企业和部分小型企业有减免。个人所得税

（PIT）按累进税率征收，最高 45%。社会保险缴费由雇主和雇员共同缴纳，费率因地区而异。因此，中国式 TDL 不能简单移植美国工资税模型，而需要同时考虑 VAT、CIT、PIT 和社保缴费渠道。

中国长期使用研发费用加计扣除机制。历史上，符合条件研发费用常与 175% 总额扣除处理相关；后续政策扩展中，许多合格企业可适用 200% 总额扣除。中国式 TDL 不宜在修辞上设计为新的“反 AI 税”。更稳的路径，是将其设计为一种优惠政策反向调节：当 AI 相关合格支出同时伴随受影响岗位族中可度量的劳动侧税基移除时，对优惠强度进行调减、追补或附加评估，并以劳动侧财政损失为校准目标。[S18]

中国的产业政策传统和集中财政权威，可能使技术实施在行政上比美国联邦体制更可行，但不能预设政治采纳。中国版本需要中央授权、财政部和国家税务总局规则制定、与产业政策部门协调，并谨慎处理就业稳定目标。

中国式 TDL 的潜在实施路径，可以从北京、上海、杭州/浙江、深圳/广东等 AI 密集地区的中央授权地方试点开始，并由中央财政和税务主管部门协调。试点框架可以借鉴中国碳排放权交易试点和高新技术企业认定经验，两者都在更广泛推广前进行了地方实验。本节论证制度可行性，不代表政治认可。

****13.4 排序****

三条路径并不互斥，以下排序只是示例性而非唯一规范性路径。一个可信排序可以将次联邦实验放在第一位，多边披露标准放在第二位，联邦或中央采纳放在第三位。每条路径都能为下一条提供信息：次联邦试点产生触发校准的经验数据，多边披露标准在管辖区之间规范化审计框架，联邦或中央采纳则只有在足够经验和制度先例积累后才更可能出现。

历史比较是碳定价，但只作为制度扩散类比。碳税和限额交易体系先是在北欧和美国西部的次国家级实验中出现，逐步积累经验证据，随后才通过欧盟排放交易体系和碳边境调节机制走向更广泛的多边协调 [S12]。TDL 如果出现类似轨迹，也应预期至少跨越十年。

****十四、制度先例****

TDL 是新的，但并不制度陌生。至少五类既有机制可以降低它的激进感。

第一，失业保险 **experience rating** 已经把企业裁员行为和企业税率连接起来。BLS 研究说明，企业失业保险税成本会在裁员后上升，虽然边际成本通常只覆盖失业保险待遇的一部分。[S9] 这是同领域先例：给劳动力市场造成财政成本的企业行为，可以影响未来税负。

第二，工伤保险和保险费率制度会对风险进行差异化定价。活动产生更高预期公共或社会成本的企业，可以承担更高缴费率。

第三，碳定价和碳边境调节机制表明，公共法可以为外部性定价，并处理跨境泄漏。欧盟委员会将 **CBAM** 描述为确保进口商品承担与欧盟国内生产等价碳价的机制。[S12]

第四，**FATCA** 表明，美国已经通过跨境申报与扣缴义务保护税基。IRS 资料显示，**FATCA** 要求某些外国金融机构和其他实体报告美国账户持有人的海外资产信息，否则可能面临应扣缴款项的扣缴。[S11]

第五，OECD/G20 **BEPS** 项目说明，税基保护可以成为多边制度议程。OECD 将 **BEPS** 定义为跨国企业利用税法漏洞把利润转移到低税或无税地点的策略，并说明 **BEPS** 项目为政府提供规则和工具来应对此类避税。[S10] TDL 不是 **BEPS**，但它是 **BEPS** 之后的问题：**BEPS** 问利润去了哪里，TDL 问纳税人去了哪里。这个类比是结构性的，不是发展阶段性的。**BEPS** 代表了数十年的多边谈判和积累的政治共识；TDL 还处于这一过程最早的概念阶段。

****十五、反驳与回应****

****15.1 "TDL 会阻碍创新。"*****

回应：TDL 不征 AI 研发，不征 AI 辅助，不征安全自动化，也不征生产率提升本身。它只在抵消

之后仍然存在财政缺口时征收。如果创新提高产出并保留或重建税基，就不会产生实质性 TDL。

****15.2 "历史表明新工作会出现。"*****

回应：可能会出现。再纳税窗口正是把这种可能性纳入公式。本文不否认岗位创造，而是拒绝把不确定的未来再吸纳当作忽视当前可观测财政损失的理由。

****15.3 "企业已经在交企业税。"*****

回应：问题不是企业是否交税，而是企业侧新增税收是否补足劳动侧税收损失。如果补足，TDL 下降；如果没有补足，剩余缺口仍然存在。

****15.4 "执行难度太高。"*****

回应：TDL 确实有执行难度，但并不比既有税制中有效税率估算、跨境利润分配、转让定价、失业保险 **experience rating** 或进口商品碳含量计算更不可行。制度设计可以从大型企业、明确触发条件和可审计岗位族开始。

****15.5 "企业会把税负转嫁给消费者。"*****

回应：部分税负转嫁可能发生，几乎所有税种都有这个问题。解决方式是税负归宿监测、反规避规则、目的地原则补充征收和谨慎税率设计。税负归宿问题影响制度设计，不否定外部性存在。

****十六、结论：财政国家不能补贴自身侵蚀****

本文论证，AI 劳动替代不仅是就业问题或生产率叙事，而是财政问题。当盈利企业通过 AI 驱动的替代将一个劳动者从工资单上移除时，财政国家可能失去一个公共收入接口：所得税预扣、工资税、社会保险缴费、消费税支持的消费活动，以及地方公共收入，都通过这个接口流动。它们的消失不会被企业侧税回收补自动抵消，因为企业侧回补取决于有效税率、扣除时间、研发费用化、资本折旧、利润分配和企业支付行为。它们的消失也不会被再就业自动抵消，因为再纳税窗口是概率分布，而不是常数。

由此产生两条结构性观察。第一，市场协调可能失灵：单个企业可以理性地用 AI 替代劳动，即使总体结果削弱它们间接依赖的需求基础、劳动力市场或公共收入系统。第二，也更根本的是，主权财政能力可能失灵：企业可以全球化需求、资本和技术租金，但国家仍要在属地内承担学校、法院、基础设施、医疗体系、社会保险和地方服务。企业可以跨境繁荣，而一个管辖区却失去用于资助公共义务的纳税人。

TDL 是本文为这一缺口提出的有限财政机制。它不对作为技术的 AI 征税，不对生产率征税，不对安全自动化征税，也不对所有替代征税。它只在财务强劲企业执行 AI 相关替代、劳动侧财政损失未被企业侧回补和再纳税窗口恢复充分抵消、剩余缺口构成未内部化外部性时才适用。审计方程 $L - C - R$ ，或其动态形式 $\Sigma t \delta t(L_t - C_t - R_t)$ ，不是假装精确的公式，而是一套分层审计协议，每一项都披露其假设、数据来源和不确定性区间。

本文更深层的论点是结构性的。工业时代的财政国家建立在利润与雇佣的部分耦合之上：盈利企业倾向于雇佣劳动者，劳动者成为公共收入接口，税法默认这一关系大体成立。全球化、外包、金融化和平台劳动已经削弱了这种耦合。AI 经济可能是迄今最清晰、最可扩展的技术形态之一，使这种耦合更系统地断裂。TDL 不能解决 AI 经济的所有分配问题；它保留的是一个更窄的制度权利：当 AI 替代移除纳税人时，国家可以要求完整公共账本，并回收现行税法未捕获的剩余财政外

部性。

TDL 是否被采纳，最终是政治问题。但底层财政观察不依赖于是否采纳。AI 时代的财政国家继承了为另一种经济设计的税收架构，而这个架构与新替代模式之间的缺口不会自行闭合。财政国家不能补贴自身税基的消失。这是 AI 时代最低限度的制度教训。

****附录 A：跨模型协作写作方法说明****

本文通过人类研究者与多个 AI 系统之间的跨模型协作形成。本附录记录的是方法，而非自我表彰。人类研究者提供原始问题框架、规范方向和最终判断；AI 系统在这一框架内提供独立结构、批评和修订。

这一过程包括三个阶段。第一，独立生成：多个 AI 系统围绕同一核心命题分别提出竞争性骨架和政策架构。第二，对抗性评审：各模型就章节顺序、公式设计、证据风险、模型边界和修辞过度彼此质疑。第三，结构性整合：最强的要素被重新组合成分层架构——预防原则为骨，财政账本为筋，制度设计为肉。

这种方法并不外在于论文。本文论证制度如何在离散条件下重建结构；本文自身的生产过程也遵循类似模式：多个离散初稿经过相互批评，被重新聚合为更高阶结构。论文的双层架构——哲学性的预防原则与财政性的审计协议——正是在这一协作过程中自然生成的。

****Shared Sources and References / 统一参考文献****

[S1] Dora Garcia-Macia et al., "Broadening the Gains from Generative AI: The Role of Fiscal Policies," IMF Staff Discussion Note, 2024. <https://www.imf.org/en/publications/staff-discussion-notes/issues/2024/06/11/broadening-the-gains-from-generative-ai-the-role-of-fiscal-policies-549639>

[S2] Ryan Abbott and Bret N. Bogenschneider, "Should Robots Pay Taxes? Tax Policy in the Age of Automation," Harvard Law & Policy Review, 2018. <https://conferences.law.stanford.edu/werobot/wp-content/uploads/sites/47/2018/02/Should-Robots-Pay-Taxes.pdf>

[S3] Internal Revenue Service, Notice 2026-11, "Interim Guidance on Additional First Year Depreciation Deduction," regarding OBBBA amendments to Section 168(k). <https://www.irs.gov/pub/irs-drop/n-26-11.pdf>

[S4] Internal Revenue Service, Rev. Proc. 2025-28, guidance related to Section 174A and domestic research or experimental expenditures. <https://www.irs.gov/pub/irs-drop/rp-25-28.pdf>

[S5] Brett Hemenway Falk and Gerry Tsoukalas, "The AI Layoff Trap," arXiv:2603.20617, 2026. <https://arxiv.org/abs/2603.20617>

[S6] The Budget Lab at Yale, "A Weakened IRS Has Substantial Consequences," April 2026. <https://budgetlab.yale.edu/research/weakened-irs-has-substantial-consequences>

[S7] U.S. Government Accountability Office, "Artificial Intelligence: IRS Actions Needed to Address Skills Gaps, Information Quality, and Strategic Management," GAO-26-107522, March 2026. <https://www.gao.gov/products/gao-26-107522>

[S8] Internal Revenue Service, Publication 15 (2026), Employer's Tax Guide; Social Security and Medicare tax rates and wage base. <https://www.irs.gov/publications/p15>

[S9] Robert Pavosevich, "The cost of layoffs in Unemployment Insurance taxes," Monthly Labor

Review, U.S. Bureau of Labor Statistics, April 2020.
<https://www.bls.gov/opub/mlr/2020/article/the-cost-of-layoffs-in-ui-taxes.htm>

[S10] OECD, "Base erosion and profit shifting (BEPS)," overview.
<https://www.oecd.org/en/topics/base-erosion-and-profit-shifting-beps.html>

[S11] Internal Revenue Service, "Foreign Account Tax Compliance Act (FATCA)," overview.
<https://www.irs.gov/businesses/corporations/foreign-account-tax-compliance-act-fatca>

[S12] European Commission, "Carbon Border Adjustment Mechanism," overview.
https://taxation-customs.ec.europa.eu/carbon-border-adjustment-mechanism_en

[S13] Service-Public.fr / impots.gouv.fr, French household employee and services à la personne tax credit materials. <https://www.service-public.gouv.fr/particuliers/vosdroits/F12> and <https://www.impots.gouv.fr/particulier/emploi-domicile>

[S14] Reuters, "Tax enforcement weakened after Trump job cuts, IRS data shows," April 15, 2026. <https://www.reuters.com/legal/government/tax-enforcement-weakened-after-trump-job-cuts-irs-data-shows-2026-04-15/>

[S15] Bill Gates' robot tax proposal (2017 Quartz interview); European Parliament Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics, draft articles on robot tax (subsequently rejected); Republic of Korea, 2018 reduction of tax incentives for automation equipment under the 2017 tax revision.
https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html

[S16] Internal Revenue Service, "IRS releases tax inflation adjustments for tax year 2026, including amendments from the One Big Beautiful Bill," 2025/2026 tax-year parameters including standard deduction and brackets. <https://www.irs.gov/newsroom/irs-releases-tax-inflation-adjustments-for-tax-year-2026-including-amendments-from-the-one-big-beautiful-bill>

[S17] Internal Revenue Service, Topic No. 560, "Additional Medicare Tax"; IRS, "Understanding employment taxes," employer withholding and no employer match for Additional Medicare Tax. <https://www.irs.gov/taxtopics/tc560> and <https://www.irs.gov/businesses/small-businesses-self-employed/understanding-employment-taxes>

[S18] China Briefing, "Pre-Tax Super Deduction of R&D Expenses in China," 2023; Tax Foundation, "U.S. Must Fix R&D Treatment to Compete with China," 2024. <https://www.china-briefing.com/news/china-rd-expenses-pre-tax-super-deduction-explainer/> and <https://taxfoundation.org/blog/us-china-competition-technology-rd-tax-treatment/>

[S19] Ayush Chopra, Santanu Bhattacharya, DeAndrea Salvador, et al., "The Iceberg Index: Measuring Skills-centered Exposure in the AI Economy," MIT and Oak Ridge National Laboratory, October 2025. arXiv:2510.25137. <https://arxiv.org/abs/2510.25137>

[S20] Congressional Budget Office, "The Distribution of Household Income, 2022," January 2026, including effective federal tax rates by income quintile and historical effective tax rate tables for 1979–2022. <https://www.cbo.gov/publication/61911>

****Data Cautions for Next Draft / 下一版数据注意事项****

Item	English caution	中文注意事项
Re-taxation window	Must be estimated by occupation, age, region, wage replacement, tax-quality ratio,	应按职业、年龄、地区、工资替代率、税基质量系数和再就业时

Item	English caution	中文注意事项
	and time-to-reemployment; no single national constant.	间估算，不能假设统一全国常数。
Corporate-side recovery	Requires effective tax rate, deductions, depreciation, R&D expensing, profit allocation, buybacks/dividends, and timing.	需纳入有效税率、扣除、折旧、研发费用化、利润分配、回购/分红和时间安排。
IRS paradox	Use Yale Budget Lab and GAO as model/estimate sources; phrase all revenue effects as estimates, not audited actual losses.	使用 Yale Budget Lab 和 GAO 作为模型/估算来源，所有收入影响应表述为估算，不表述为已审计实际损失。
Falk-Tsoukalas model	Use as market-failure support under model assumptions, not as universal empirical proof.	作为模型假设下的市场失败支撑，不作为普遍经验证明。
OBBBA capital bias	Avoid saying human training is never deductible; the stronger claim is timing and structural asymmetry.	避免说人力培训从不允许扣除；更强论点是税收时间性和结构性不对称。
§5.5 worked examples (v0.3)	First-year illustrations only; do not mix first-year values with multi-year restoration. Dynamic audits require time-indexed calculation.	仅为第一年示例；不要把第一年变量与多年期恢复值混用。动态审计需按时间展开计算。
Payroll details (v0.3)	High-income examples must include Additional Medicare Tax on the employee side only; no employer match.	高收入算例需纳入雇员端 Additional Medicare Tax ；该项没有雇主匹配。
AI expenditure treatment (v0.3)	Do not assume all AI spending qualifies for Section 168(k); distinguish depreciable property, Section 174A R&E, cloud/vendor/API/SaaS payments, and capitalization.	不能假设所有 AI 支出都适用 Section 168(k) ；需区分可折旧财产、 Section 174A 研发支出、云/供应商/API/SaaS 付款和资本化处理。
§11.2/§11.5 thresholds and coefficients	Illustrative calibration placeholders; require empirical calibration through pilots before statutory use.	示例性校准占位值；进入法定使用前需通过试点进行经验校准。
§13.3 China adaptation (v0.3)	Discusses institutional feasibility, not political endorsement; R&D super-deduction treatment must be verified against current central and local rules.	讨论制度可行性，不代表政治认可；研发费用加计扣除处理需根据中央和地方现行规则核验。
Non-profit/PBC structures	PBC is a governance form, not a tax-exempt status by itself; TDL should target downstream	PBC 本身是治理形态，不是免税身份； TDL 应针对下游替代行为，而非组织形式。

Item	English caution	中文注意事项
	substitution behavior, not organizational form.	
§5.6 exposure envelope (v0.5)	Exposure-envelope illustration, not a revenue forecast. Layer 1/2 rely on a counterfactual full-exposure assumption; the 28–35% combined rate is a synthetic assumption, not a directly observed national rate; the 5–15% realization rate is illustrative, not an adoption forecast. MIT figures are subject to revision.	暴露包络线示意，不是税收收入预测。第一层/第二层依赖反事实全暴露假设；28–35%的综合率是假设，不是直接观测到的全国税率；5–15%的实现率是示例，不是采纳预测。MIT 数据可能修订。