the budget lab

# Methodology: How potential AI futures interact with the current tax system

Published: July 20, 2026

This document describes the methodology behind The Budget Lab's report *How potential AI futures interact with the current tax system*. We combine projections of the economic effects of AI over the medium term from Karger et al. (2026) with the same microdata that powers our tax microsimulation model to produce a counterfactual estimate of each tax unit's capital and labor income under different AI scenarios. The analysis isolates the mechanical revenue impact of the three interacting features of any potential AI growth scenario: increased GDP growth, a shift in the labor- and capital-share of factor income, and a change in the distribution of labor income.

## Data and baseline construction

### The tax microsimulation file

All tax-side computations except the estimation of the corporate income tax use The Budget Lab's tax microsimulation. The base file is the 2015 IRS Public Use File (PUF, ~150,000 records), supplemented with non-filer imputations and aged to the baseline year through SSA population weights and CBO economic scaling. The Tax-Simulator calculator handles individual income tax (AGI, standard and itemized deductions, ordinary and preferential rates, AMT, NIIT, EITC, CTC, CDCTC, education credits, Saver's, QBI) and the payroll-tax module. It does not handle the corporate income tax or the estate tax; we address the CIT through the wedge described below.

The tax microsimulation alone does not carry asset-level wealth, which we need to allocate AI-driven capital income across households. To address this, we match tax units in the PUF with similar units in the Federal Reserve's Survey of Consumer Finances (SCF). This process involves imputing asset values onto tax units using a random-forest machine-learning approach. This means that for every tax unit, we have an estimate of cash, equities, bonds, retirement balances, life insurance, annuities, trusts, real-estate funds, primary and other home equity, and pass-through equity, as well as active/passive splits for S-Corp and partnership profit. The PUF-SCF match is validated against aggregates from NIPA, SOI, and the Distributional Financial Accounts (DFA) upstream in the data pipeline. This project uses the PUF-SCF matched file as the ground truth.

We specify a tax unit's income as $Y^X_{j,i}$, where X specifies the type of income (labor, capital, etc.), $j \in \{0, 1\}$ specifies before or after the shock, and i is a tax-unit index. The aggregate income values lack the i subscript (i.e.,

$Y_j^X$). Each tax unit has a PUF weight $w_i$ that has been adjusted to match the SSA and CBO targets discussed above.

## Per-unit labor and capital aggregates

Once the merged file is loaded, we group a tax filer's pre-shock income into two buckets:

- **Labor income** ($Y_{0,i}^L$): wages, sole-proprietor and farm income, and the labor portion of each unit's S-Corp and partnership net profit.
- **Capital income** ($Y_{0,i}^K$): interest (taxable and exempt), dividends (ordinary and preferential), realized capital gains, other gains, net rent and estate income, taxable IRA and pension distributions, plus the capital portion of S-Corp and partnership net profit.

These values do not flow into the tax simulation but are only used to design the AI shock.

## Splitting pass-through profit between labor and capital

S-Corp and partnership profit blends true entrepreneurial labor ("a partner is paid for working") with returns on invested capital ("a partner is paid for owning"). The split matters: a shift from wages to capital income carries different tax consequences than a shift from wages to pass-through active income, even when the dollar amount is the same. To address this we follow Saez and Zucman (2020) in allocating profits across labor and capital income differentially across the income distribution. For each tax unit, let $W^\star$ be the 99.99th percentile of positive pre-shock wages ($Y_{0,i}^w$) and $Y_{j,i}^{PT}$ be their pass-through income of type $PT$ (S-Corp active, partnership active, and passive).

First, a tax unit's passive labor pass-through income is $Y_{j,i}^{L,Pass} = Y_{j,i}^{Pass} \cdot 0.25$ with, $Y_{j,i}^{K,Pass} = Y_{j,i}^{Pass} \cdot 0.75$. Second, for both S-Corp and partnership active income, we divide up income between capital and labor differentially based on income relative to $W^\star$:

$$Y_{j,i}^{L,PT} = 0.75 \cdot \min\left(Y_{j,i}^{PT}, W^\star\right) + 0.25 \cdot \max\left(Y_{j,i}^{PT} - W^\star, 0\right)$$

And

$$Y_{j,i}^{K,PT} = Y_{j,i}^{PT} - Y_{j,i}^{L,PT}$$

The intuition is that an active partner earning at or below typical top-end wages looks like they are being paid mostly for their labor, while an active partner whose draw vastly exceeds top wages is more plausibly receiving rents on firm capital.[1]

# Macro shock parameterization

Table 1 reports the values that are implied by our implementation of Karger et al. (2026) economic forecasts. First, the key parameters are:

- an annualized real GDP growth rate $(r_{ai})$ over the same horizon, again as a median across forecasts;
- a post-shock labor / capital share of factor income ($\theta_1^L$, $\theta_1^K$) at the policy horizon, reported as a 5-year-ahead median across analyst forecasts;
- a labor-income inequality parameter (λ) that alters the distribution of positive labor income; and
- for each variant, we estimate a version with and without the shift in the capital-labor share (denoted as $\chi \in \{R, F\}$).

We therefore have a policy shock defined by $\{r_{ai}, \theta_1^L,\ \theta_1^K, \lambda, \chi\}$.

We use three intensity variants: Slow, Moderate, and Rapid (labeled S, M, R). We take 2030 as the baseline year to match Karger's 5-year horizon and pin every variant from Tables 19 (growth rates) and 39 (labor shares).

Karger et al. (2026) surveys several different groups of AI experts; we use estimates for the economist subsample. Our microsimulation already grows in line with CBO's baseline level of cumulative GDP growth through 2030 of 9.7%, or 2.2% for 2026 and 1.8% for all following years $(G_{CBO})$. We use g to denote five-year cumulative growth rates, and r to denote annual growth rates.

APPENDIX TABLE 1

## Key Parameters

| | AI Adoption Scenarios via Karger et al. (2026) | | |
|---|---|---|---|
| | **Slow** (S) | **Moderate** (M) | **Rapid** (R) |
| **AI-adoption inputs** *(5-year horizon, 2025-2030)* | | | |
| Annual GDP growth under AI ($r_{ai}$) | 2.0% | 2.6% | 3.3% |
| 2030 capital share ($\theta_1^K$) | 45.0% | 46.2% | 48.7% |
| 2030 labor share ($\theta_1^L$) | 55.0% | 53.8% | 51.3% |
| **Derived cumulative growth** *Cumulative over the 5-year horizon, above the CBO baseline path.* | | | |
| GDP ($g_Y$) | 0.59% | 3.58% | 7.17% |
| Capital ($g_K$) | 1.72% | 7.54% | 17.28% |
| Labor ($g_L$) | -0.32% | 0.41% | -0.94% |
| **Labor-income inequality parameter ($\lambda$)** *Compressive < 1, Expansive > 1* | | | |
| Compressive | 0.994 | 0.964 | 0.928 |
| Proportional | 1.000 | 1.000 | 1.000 |
| Expansive | 1.006 | 1.036 | 1.072 |

**Source:** Karger, Buehler, Cox, Saint-Jacques, and Bjorkegren (2026), CBO (2025), and Budget Lab

Data Image

## From shares and growth to growth rates by factor

Let $\theta_0^L$ and $\theta_1^L$ be the baseline and post-shock labor shares of factor income, with capital shares $\theta_t^K = 1 - \theta_t^L$. Let $H = 5$ be the Karger horizon and

$$G_{CBO} \;=\; \left(1 + r_{CBO}^{2026}\right) \cdot \left(1 + r_{CBO}^{2027+}\right)^{H-1}$$

be cumulative CBO growth from 2025 to 2030. Then

$$g_Y \;=\; \frac{\left(1 + r_{ai}\right)^H}{G_{CBO}} - 1$$

is the additional 5-year growth in nominal factor income caused by AI. For the three variants this deviation works out to roughly +0.6% (S), +3.6% (M), and +7.2% (R) above the CBO baseline at year 2030.

The growth rates of capital $(g_K)$ and labor $(g_L)$ follow from $g_Y$ and the pre-and post-shock factor income shares of capital and labor; $(\theta_0^K, \theta_1^K, \theta_0^L, \theta_1^L)$:

$$g_K = \frac{\theta_1^K\left(1 + g_Y\right) - \theta_0^K}{\theta_0^K}, \qquad g_L = \frac{\theta_1^L\left(1 + g_Y\right) - \theta_0^L}{\theta_0^L}$$

The microsimulation's measured labor share differs from the NIPA share Karger reports — the microsim aggregates exclude employer FICA, undistributed C-corp profits, and imputed housing rent — but the growth rate of each factor is comparable across data sources, so the shock translates cleanly.

# Labor-side redistribution

The labor-side step computes a counterfactual labor income aggregate

$$Y_1^L = Y_0^L \cdot (1 + g_L)$$

and redistributes it across tax units. We offer three redistribution rules — proportional (S0), compressive (S2), and expansive (S3). None of the rules imply any extensive-margin adjustment; employment is held constant. The omitted S1 is a planned future expansion involving occupation-level AI exposure.

## Proportional (S0)

We hold constant the labor income for tax units with negative baseline labor income and scale the labor income of all tax units with positive baseline labor income by a single multiplier $\rho$ so the aggregate equals $Y_1^L$:

$$Y_{1,i}^L = \begin{cases} \rho \cdot Y_{0,i}^L & \text{if } Y_{0,i}^L > 0, \\ Y_{0,i}^L & \text{otherwise.} \end{cases}$$

By construction, $\sum_i w_i \cdot Y_{1,i}^L = Y_1^L$ (limited to positive values). The proportional scenario preserves the baseline distribution of positive labor income.

## Compressive (S2) and Expansive (S3)

It is empirically uncertain whether the labor-side adjustment induced by AI will compress (S2) or stretch (S3) the distribution of labor income.[2] We consider both and implement them via opposing dispersion-shift specifications using a log-affine transformation applied to tax units with positive labor income:

$$Y_{1,i}^L = \mu_1 + \lambda \cdot \left(\ln Y_{0,i}^L - \mu_0\right),$$

where $\mu_0$ is the weighted log-mean of positive baseline labor income, $\mu_1$ is solved so the positive-subset aggregate matches $Y_1^L$ after scaling, and $\lambda$ is defined as the ratio of post- to pre-shock standard deviation of $\ln Y_{j,i}^L$ on the positive subset. We tie $\lambda$ to the productivity shock $g_Y$ through

$$\lambda_{S2} = 1 - k \cdot g_Y \quad \text{(compressive)}, \qquad \lambda_{S3} = 1 + k \cdot g_Y \quad \text{(expansive)}.$$

For now we assume $k = 1$, which results in a proportional mapping between growth and dispersion. The code supports a scalar multiplier if a different calibration becomes appropriate.

The compressive scenario captures the world in which AI's labor-side cost falls hardest on the top end of the wage distribution (substitution at the high end, e.g. cognitive professional work). The expansive scenario captures the opposite — AI raises the productivity (and pay) of high-skill labor while displacing or stagnating mid- and bottom-wages. We do not take a view on which is more likely.

# Capital-side allocation

This section details the allocation of the excess capital income generated by the shock to individual tax units and — within tax units — across streams of capital income. Given how we define the baseline capital-income flow as a function of taxable income reported on tax returns, we are growing an already-realized stream of taxable income. For this initial version of the model, we assume that realization rates are held constant, meaning that capital income is not realized (and therefore taxed) at a different rate due to the productivity shock or the capital-labor shift.

## Two share modes: reallocation vs. fixed share

We run every variant under two different factor-share settings ($\chi$) to let readers separate the productivity and reallocation channels:

- **Reallocate (R)** — use $\theta_1^K$ from Karger so capital grows faster than labor.
- **Fixed share (F)** — pin $\theta_1^K := \theta_0^K$ so the post-shock factor split equals the baseline. Here, we use the same $g_Y$, but $g_K = g_Y$ so labor and capital both grow at the productivity rate.

The fixed-share option isolates the productivity increase from the factor reallocation. Together, the two share modes let us decompose "AI's effects on revenue and distribution" into "more taxable output" and "the mix of who earns that output has shifted."

## From the macro shock to dollars flowing to households

The aggregate capital income flow added by the shock is

$$Y_1^K = \ Y_0^K \cdot (1 + g_K)$$

For ease of notation, we define the excess capital income as:

$$X = Y_1^K - Y_0^K = Y_0^K \cdot g_K$$

Corporate income tax operates upstream of household realizations, so the AI capital flow $X$ reaches tax units in full. In order to capture the effect of an AI shock on corporate income tax revenue, we need to estimate what growth in C-corporation income is implied by $X$. We proceed simply by multiplying the change in capital income by the ratio of CBO's baseline level of CIT revenue $(R_{CBO}^{CIT})$ to baseline (2030 in this case) realized capital income ($Y_0^K$), implying that the AI CIT delta then scales linearly with $X$:

$$\Delta R^{CIT} = X \cdot \frac{R_{CBO}^{CIT}}{Y_0^K}$$

We can back-out the implicit assumptions implied by this set-up as follows. Let $\kappa$ be the share of $Y_0^K$ that flows through C-Corporations (versus entities that are not exposed to the corporate income tax, like partnerships and S-Corporations), $\tau_C^{stat}$ be the statutory corporate tax rate $(21\%)$, and $\eta$ be an adjustment factor absorbing the statutory-vs-effective gap (avoidance, NOLs, credits, profit shifting, etc..). Then:

$$R_{CBO}^{CIT} = Y_0^K \cdot \tau_C^{stat} \cdot \eta \cdot \kappa$$

Our formulation above then assumes that $\tau_C^{stat}, \eta, \kappa$ are constant across AI-induced shocks.

$$\Delta R^{CIT} = X \cdot \frac{R_{CBO}^{CIT}}{Y_0^K} = X \cdot \tau_C^{stat} \cdot \eta \cdot \kappa = X \cdot \tau_C^{eff} \cdot \kappa$$

Where $\tau_C^{eff} := \tau_C^{stat} \cdot \eta$. The absence of change in the statutory tax rate is true by assumption, and future work will investigate how AI-induced capital growth might differentially impact the take-up of various adjustments that make up $\eta$ and vary across corporate form, impacting $\kappa$.

With $\tau_C^{stat} = 21\%$ (TCJA), $\kappa \approx 0.50$ (via NIPA), and CBO's 2030 CIT anchor of $\sim \$486$B (or a CIT-to-GDP equaling $0.013$ times GDP), $\eta$ falls out to roughly one in practice; its exact value is reported on the parameters sheet of every run.

The corporate-tax delta enters aggregate revenue as a single number layered onto the micro-simulation revenue total; it is not attributed back to individual households. As a result, every distributional figure in our output reflects the full capital flow $X$ that reaches households. This is a deliberate choice: we do not adopt a per-household corporate-incidence assumption.

We distribute $X$ across tax units in proportion to each unit's share of pre-shock total household wealth:

$$X_i \;=\; X \cdot \frac{A_{0,i}}{\sum_k w_k A_{0,k}},$$

where $A_{0,i}$ is the sum of the unit's SCF-imputed wealth columns (cash, equities, bonds, retirement balances, life insurance, annuities, trusts, real-estate funds, primary and other home equity, pass-through equity, and miscellaneous non-financial assets). Allocating to total wealth assumes the AI shock raises returns proportionally across asset classes; pinning to a narrower base would push the distributional incidence further toward the top and is an area we plan to explore in the future.

## Within-unit allocation across income types

The asset base contains both income-bearing and non-income-bearing classes. Of the income-bearing classes, four matter for tax purposes: taxable equities, bonds, pass-through equity, and retirement balances. Let $A_{j,i}^{Inc}$ be their sum for a given tax unit. We allocate each unit's $X_i$ across the four classes proportional to its baseline holdings within $A_{0,i}^{Inc}$.

Units with zero income-bearing holdings route their entire $X_i$ to retirement, preserving the aggregate identity $\sum_i w_i X_i = X$. This fallback affects roughly 33% of tax units in 2030 but only 6.7% of $X$; most of the affected mass sits in tax units whose wealth is only in housing. If we redefine the asset base to exclude housing, only 2.4% of $X$ is routed through retirement via this fallback.

Each class then flows through to specific taxable-income types:

- **Taxable public-equity flow** splits between qualified dividends and gross long-term capital gains using a fixed share table (30% dividend / 70% gross LTCG).

- **Fixed-income flow** splits between taxable and tax-exempt interest at each unit's baseline ratio.
- **Pass-through equity flow** becomes pass-through ordinary income, further allocated across S-Corp and partnership active/passive sub-buckets using baseline positive capital holdings as weights.
- **Retirement flow** runs through the procedure described in the next section.

# Building the counterfactual and running the tax calculator

For every cell in the scenario grid (identified by $\{r_{ai}, \theta_1^L,\ \theta_1^K, \lambda, \chi\}$) we construct a tax-unit file with the same schema as the baseline and write it to a sibling scenario folder under the pinned Tax-Data vintage. Income is adjusted following the processes described above, and the Tax-Simulator then reads the counterfactual file as a sibling Tax-Data ID and runs the full IIT + payroll calculation on it.

## Decomposing the revenue change into labor and capital contributions

Two separate channels move revenue in a typical scenario: shifting labor income and adding capital income. The interaction between them is non-linear (because the tax schedule itself is non-linear: brackets, phase-outs, AMT), so a clean "labor versus capital" attribution requires care.

Each cell produces three Tax-Simulator runs:

- **Both** — labor and capital shocks both applied (the standard scenario).
- **Labor-only** (LO) — labor side shocked; capital flows held at baseline.
- **Capital-only** (CO) — capital side shocked; labor flows held at baseline.

We then decompose the total revenue change as

$$\Delta T_{both} \;=\; \Delta T_{LO} + \Delta T_{CO} + \text{interaction},$$

where

$$\text{interaction} \;=\; T_{both} - T_{LO} - T_{CO} + T_{base}.$$

The interaction term captures non-linear bracket effects (the combined shock can push a unit across an AMT or NIIT threshold that neither side would have triggered alone). Payroll tax counts mechanically as labor: the capital-only run leaves wages at baseline, so its payroll delta is zero.

# Aggregation and the revenue-to-GDP anchor

The final step reads Tax-Simulator output and produces the publishable aggregates: revenue deltas per instrument, pre-tax and after-tax income deltas, Gini deltas on pre-tax and after-tax income, decile shares, and top-1, top-0.1, top-0.01 shares. We also layer the macro CIT delta onto the microsim revenue total to produce the bottom-line revenue change.

## Connecting to the macro model

We use the Budget Lab Small Macro Model (BLSMM) to estimate the effect of these shocks on the debt-to-GDP ratio. To do this, we need to provide the model with the change in the revenue-to-GDP ratio. Because the modeling above omits various forms of tax revenue not directly estimated by either our macro wedge or our individual tax simulations, the model's internal revenue / GDP ratio understates the published level by roughly 1.3 percentage points in 2030. For the publishable revenue-to-GDP comparison we therefore convert the baseline to CBO's published 2030 ratio (17.7% from CBO publication 62105) into dollars and adjust this dollar amount by the estimated change in revenue from our simulations.

Represent the revenue-to-GDP ratio as $Z_j = \frac{R_j}{Y_j}$, where j indicates the source of the data, with $j \in \{0, 1\}$ following the notation above. We want to estimate $\Delta Z$, or the change in the revenue-to-GDP ratio due to a given AI shock. As noted above, defining $R_1$ based solely on the change in revenue estimated above under-estimates total revenue, so we define:

$$R_1 \;=\; Z_0 \cdot Y_0 + \Delta R_{TBL}$$

The estimated change in revenue ($\Delta R_{TBL}$) includes both the Tax-Simulator change in income and payroll tax (net of refundable credits) plus the corporate-tax computed outside the microsim. The baseline revenue-to-GDP ratio $Z_0$ and the GDP level $Y_0$ come from the 2026 Budget and Economic Outlook.

Given this, we can estimate the change in the revenue-to-GDP ratio:

$$\Delta Z = Z_1 - Z_0 = \frac{R_1}{Y_1} - \frac{R_0}{Y_0} = \frac{Z_0 \cdot Y_0 + \Delta R_{TBL}}{Y_0(1 + g_Y)} - \frac{R_0}{Y_0}$$

We can re-write this as:

$$\Delta Z = \underbrace{\frac{\Delta R_{TBL}}{Y_0(1 + g_Y)}}_{\text{Revenue}} + \underbrace{Z_0 \frac{g_Y}{(1 + g_Y)}}_{\text{Macro}}$$

Where the first term represents the change in revenue, and the second represents the shock to GDP. Even with no revenue change, GDP growth alone would push revenue / GDP down. Anchoring to CBO applies that drag to a realistic revenue base, so the published response number is meaningful relative to the forecast level. The maintained assumption is that revenue streams the model omits (notably the baseline CIT level and any pieces of "other" revenue not captured by Tax-Simulator) are unchanged in the scenario except through channels already in $\Delta R$.

With this input, we work through the following process:

1. Linearly scale the revenue-to-GDP ratio so that by the baseline year, the full $\Delta Z$ is applied. The adjustment is zero in 2025, 0.2 of the target in 2026, and reaches full level in the baseline year, and held at full level thereafter.
2. Solves for a constant per-year productivity bump (added to BLSMM's potential-output path) such that the annualized 2025–baseline-year real GDP growth in BLSMM matches the Karger variant's $r_{ai}$ target.
3. Reads off BLSMM's 2030 federal debt and debt/GDP

BLSMM's CBO-style fiscal feedback (the model's $\psi_1, \psi_2$ elasticities) mechanically lowers outlays as potential output rises. Under an AI shock policymakers may instead raise outlays for displaced-worker support; the BLSMM number therefore reads as a baseline-feedback estimate of debt/GDP rather than an all-things-considered projection.

# Future work

This initial version of the model has several areas where there is room for improvement. We identify several areas here.

## Labor-side AI exposure

A natural fourth labor-incidence rule scales each unit's labor income by an occupation's exposure to AI: $Y_{1,i}^L = (1 + \beta e_i)\, Y_{0,i}^L$ where $e_i$ is an occupation-level AI-exposure index, and $\beta$ is solved so the aggregate matches $Y_1^L$. The mechanic is straightforward; the missing piece is an occupation-level exposure imputation onto the tax-unit file. Once that imputation exists, the AI-exposure scenario slots into the labor-scenario axis alongside the existing proportional / compressive / expansive rules.

## Intensive versus extensive labor-side responses

The current specification assumes that the AI effect is applied on the intensive margin: all tax units have their labor income adjusted by some parameter. However, it is highly likely that some of the effect of AI on labor income will be felt on the extensive margin via job loss or creation. A future extension would allow individual income to adjust along this margin to match $Y_1^L$.

## Moving away from the 1040-derived baseline

Currently, the AI shock is applied to both labor and capital as a function of observed income on the PUF and imputed assets via the SCF. However, this bakes in a number of assumptions, including the realization rate of capital, the shares of corporate income that is held by C-Corporations, the treatment of assets held in retirement accounts, and the way in which we implement the corporate income tax effect. Ideally, we would begin with a model of the macro-economy that could be subjected to an AI shock. That shock would flow down to different types of income (capital, labor, etc.) via a series of parameters, allowing a far greater degree of flexibility in defining the structure of a potential shock. We discuss a few examples of these sorts of extensions below:

- If the macro-shock is applied in such a way to allow us to observe an impact of AI on capital income that flows through C-Corporations, then we could specify a more extensive corporate income tax model, rather than our simple macro-wedge approach detailed above.
- The current release treats AI capital gains as recognized immediately at accrual, because the capital income shock is built from the on-1040 realized base and re-applying a realization discount would double-count. A natural extension is to begin with a definition of $X$ that includes unrealized capital gains and apply an explicit realization rate — the share of the annual gain stock realized each year — as a separate parameter: This requires a setting where $X$ captures unrealized as well as realized gains and an estimate of the realization rate from the literature. Further work would condition the realization rate on individual or economy-level

characteristics: the marginal tax rates a unit faces, the composition of its wealth and income, and the relative size of its capital versus labor base.

- For retirement income specifically, we would ideally model not just the realization rate but the share of realizations that are taxable. Most retirement wealth sits in tax-deferred vehicles (401(k), IRA, traditional pensions); the marginal AI dollar flowing into those accounts would not be taxed in-year.

## Lifetime present value

The single-year flow treatment misses a longer-term dimension of capital-gains taxation: gains can compound unrealized, get realized at some hazard rate, or escape tax entirely via step-up at death. A present-value formulation — unrealized pool compounding at $\rho$, realized at hazard rate $r$, discounted at $\delta$, with a step-up haircut $\varphi$ on the residual at horizon $T$ — captures this. Tax-Simulator is an annual-flow object and cannot natively consume a lifetime-PV input, so this extension requires either an out-of-model aggregation step or a redesign of the Tax-Simulator handoff.

## Behavioral response

The current specification is mechanical. Two behavioral channels are worth modeling: labor-supply responses to the change in net-of-tax returns (Elasticity of Taxable Income $\sim 0.25$ per Saez-Slemrod-Giertz), and realization-rate responses to capital gains tax exposure.

## General equilibrium

The current model holds prices, wages, and the asset stock $A_i$ at baseline and distributes the flow $X$ without revaluing the stock. A full GE pass would let asset prices re-equilibrate as AI capital returns rise, pushing the distributional results further toward the top.

# References

Karger, E., Bredemeier, C., et al. (2026). *AI and the Macroeconomy: Productivity, Income, and Distribution*. NBER working paper w35046.

Smith, M., Yagan, D., Zidar, O., and Zwick, E. (2019). “Capitalists in the Twenty-First Century.” *Quarterly Journal of Economics* 134(4).

Saez, E., and Zucman, G. (2020). “The Rise of Income and Wealth Inequality in America: Evidence from Distributional Macroeconomic Accounts.” *Journal of Economic Perspectives* 34(4).

Saez, E., Slemrod, J., and Giertz, S. (2012). “The Elasticity of Taxable Income with Respect to Marginal Tax Rates.” *Journal of Economic Literature*.

Piketty, T., Saez, E., and Zucman, G. (2018). “Distributional National Accounts.”

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.”

Felten, E., Raj, M., and Seamans, R. “How will Language Modelers like ChatGPT Affect Occupations and Industries?”

Chodorow-Reich, G., et al. (2024). “The Effects of the 2017 Tax Cuts and Jobs Act on US Corporations.”

Congressional Budget Office (2026). *Budget and Economic Outlook* (publication 62105).

Federal Reserve Board. *Distributional Financial Accounts*.

The Budget Lab at Yale. *Tax Microsimulation at The Budget Lab*.

The Budget Lab at Yale. *Estimating the Distributional Impact of Policy Reforms*.

## Footnotes

1 We assume that active losses are allocated 25 percent to capital and 75 percent to labor.

2 Code S1 is purposely omitted for future work. Our goal is to have an AI-exposure based distribution of labor income in this place.