

UBI is a damper, not a cure

A heterogeneous-agent simulation of AI-driven labor displacement, and what it tells us about the policies that actually flatten Gini

1. Executive Summary

The macroeconomic policy debate around AI-driven labor displacement is largely conducted in models that cannot answer the question being asked. Representative-agent DSGE — still the workhorse of central-bank policy analysis — collapses the population to a single optimising household. By construction, it cannot distinguish a routine-cognitive worker whose task is automatable and unaugmented from a STEM-technical worker whose tasks are augmented and complementary. It cannot tell you what UBI does to a Gini coefficient because the model has no Gini coefficient — every agent is the average agent.

The newer heterogeneous-agent extensions (HANK, TANK, two-agent NK) improve on this by allowing two or three income classes, but the tractability that makes them solvable analytically also smooths over the structural details that drive the redistributive question: the lagged firm-level adoption of new capital, the long-tenure unemployment exit channel, the Pareto-tailed wealth concentration that determines who actually owns the capital being redistributed.

There is also an alpha-eating problem in the policy literature itself. The redistributive 'alpha' of any flow-based policy — UBI funded by a capital-income tax, automation taxes, payroll-rebate schemes — is proportional to the fraction of national income that policy taps. As the labor share falls and the capital share rises, the same policy alpha shrinks against a growing structural transfer it is supposed to offset. The lever weakens precisely as the problem worsens. This is the gap an agent-based model is built to close.

In this blog, we simulated 4,000 heterogeneous workers across seven occupational classes, 200 firms, and a Pareto-distributed capital pool over a 10-year horizon, in three calibrated AI-capability scenarios, with five policy regimes, and 60 Monte Carlo seeds per cell.

Three findings emerged that materially change how the UBI debate should be framed:

1. UBI compresses inequality but cannot reverse its direction in a high-displacement world.

Across three calibrated AI scenarios, a 30%-of-median-wage UBI funded by a 20% capital tax lowers the terminal Gini coefficient by 0.04, 0.05, and 0.07 respectively — real compression, statistically robust across 60 seeds. In the Aggressive scenario, however, terminal Gini still climbs from 0.35 to 0.43, ending above the US 2024 disposable-income baseline of 0.40 and still rising at month 120. **So what:** fixed-rate UBI works as advertised in low-and-moderate-displacement regimes, but in a world of rapid AI capability gains it buys time rather than solving the problem. Policy debates that frame UBI as *the* answer to AI-

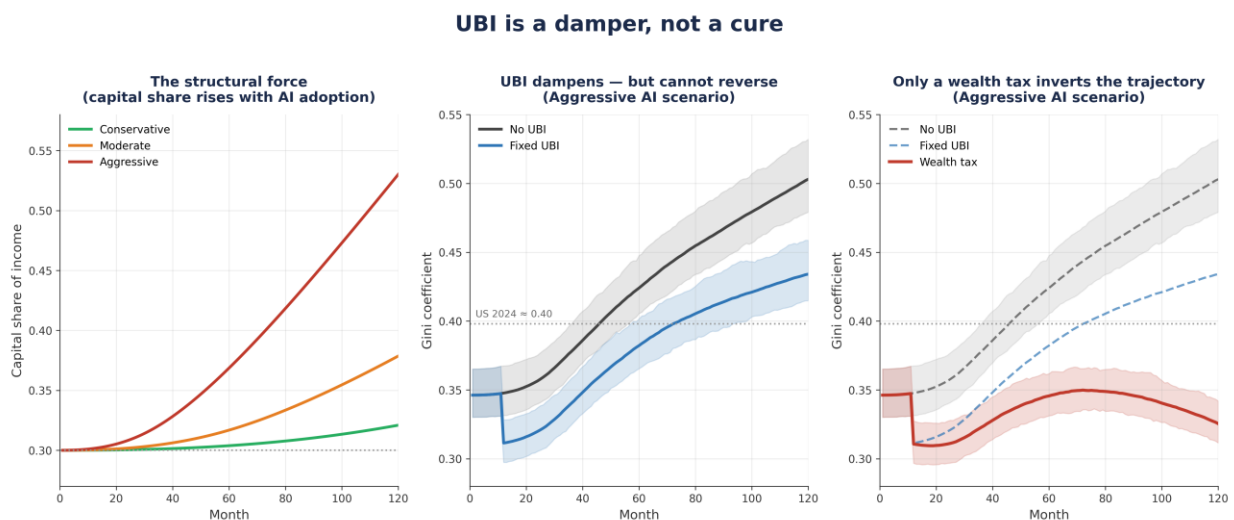
driven inequality are mis-specifying the question; in the high-displacement scenarios where the policy matters most, fixed UBI is necessary but insufficient.

2. The structural engine is the rising capital share, and a fixed-rate flow tax cannot catch it.

Capital's share of national income drifts up by 2 percentage points in Conservative, 8 in Moderate, and 23 in Aggressive — from 30% to 53% in a single decade. That transfer from labor to capital compounds independently of any per-capita redistribution. A flat 20% capital tax offsets roughly half of it; the rest is beyond the lever's reach. The funded-share metric makes this concrete: the same UBI is 37% self-financed in Conservative and 73% in Aggressive — policy capacity grows with the problem but never quite catches it. **So what:** stop debating UBI generosity numbers in the abstract. The binding constraint isn't how much you promise to pay, it's how much the funding base can raise against a moving target. Any redistributive policy designed for the AI age must scale with capital share, or the lever weakens precisely as the problem worsens.

3. Of three policy fixes tested, only the wealth tax on the ownership stock inverts long-run Gini direction.

A capital-share-scaled tax is a stronger damper but still doesn't reverse direction in Aggressive (terminal Gini 0.41 vs 0.43 under baseline UBI). Indexing UBI generosity to capital income is *mathematically identical* to baseline UBI in this model, because the funded-share constraint binds — switching the target rule moves a number on the policy fact sheet, not the cash flow. Only the wealth tax inverts: top-10% capital ownership erodes from 64% to 41% over the decade, and Aggressive Gini peaks at month 71 before falling to 0.33, below the US baseline. **So what:** the public debate routinely conflates flow-side and stock-side redistribution. Taxing capital income (a flow) dampens inequality but cannot reverse it once ownership is concentrated, because the flow is itself proportional to the stock. Only stock-side intervention — wealth taxes, sovereign wealth funds with broad-based dividends, ownership-redistribution schemes — moves the steady state.



Three panels, three findings: capital share rises with AI adoption (left); fixed UBI dampens but doesn't reverse the resulting inequality (centre); only a wealth tax inverts the trajectory (right). Aggressive AI scenario, P10–P90 across 60 seeds.

2. Definitions

The six core concepts in this paper interlock more than the surrounding literature typically admits. We use an **Agent-Based Model (ABM)** — a stochastic simulation in which a heterogeneous population of workers, firms, and capital owners interact under explicit behavioural rules — because the question we're asking is intrinsically distributional. We want macroeconomic outcomes (the income distribution, employment dynamics, inequality) to *emerge* from those interactions rather than be imposed by the representative-agent assumption that closed-form models rely on. An ABM is the right method when the heterogeneity is the point.

The policy lever we test is **Universal Basic Income (UBI)** — an unconditional, flat-rate transfer paid to every adult, funded here by a tax at rate τ on capital income, with generosity set as a fraction g of the initial median wage. Both τ and g are dials policy can move; everything that follows traces what those dials do under different AI-displacement trajectories.

The structural force the policy is fighting against is the **capital share s** — the fraction of national income flowing to capital rather than to labor. The empirical labor-share decline literature documents s rising from roughly 0.30 to 0.40 across advanced economies since the 1980s; our Aggressive scenario continues that trajectory to $s = 0.53$ over a single decade. As s rises, two things happen at once: the base on which τ operates grows, *and* the labor income UBI is meant to supplement shrinks. The policy lever and the problem it offsets move together.

The **funded share $\phi_t \in [0,1]$** is the diagnostic that ties the two together — the fraction of full-generosity UBI that the τ tax can actually cover in any given month. When the capital base is small relative to the promise, $\phi_t < 1$ and the per-capita transfer is scaled down proportionally; when the base is large, ϕ_t saturates at 1 and the policy is fully self-financed. ϕ_t is the model's running answer to the question "*can this policy pay for itself, given what's happening to the economy beneath it?*" — and how it evolves across scenarios is half the headline finding.

The *shape* of the initial capital distribution matters too, and we govern it with the **Pareto α** parameter: smaller α means a heavier right tail and more concentrated wealth. We use $\alpha \approx 1.5$ as the baseline because it approximately matches the empirical US wealth distribution. This parameter doesn't move with policy, but it determines the magnitude of every lever the policy gets — stock-side interventions like wealth taxes are proportional to the concentration they redistribute against, so a heavier tail amplifies them.

The headline outcome we track at every tick is the **Gini coefficient G** — the standard 0-to-1 inequality measure of the cross-sectional income distribution. Every other quantity in the paper is ultimately a story about what's happening to G . The reference benchmark is the US 2024 disposable-income Gini ≈ 0.40 (World Bank 2022, latest available). When we say UBI "compresses inequality," we mean it lowers G against a no-UBI counterfactual at every t . When we say it "doesn't reverse" the trend in Aggressive, we mean G still rises over the horizon despite that compression — because the structural force in s , working through ϕ_t and weighted by the wealth tail at α , outpaces what a flat τ can fund.

3. Methodology

The model is governed by seven equations that together determine the simulation. Each tick is one calendar month; the horizon is 120 months.

3.1 The accounting identity

Each worker's total income is the sum of three components — labor (wage \times employment indicator), capital dividend (their share of the post-tax capital pool), and UBI (when active):

$$y_{i,t} = w_{i,t} \mathbb{1}[e_{i,t} = 1] + \kappa_{i,t} + u_{i,t}$$

3.2 The displacement hazard

Each tick, every employed worker faces a per-tick displacement hazard $h_{\{i,t\}}$ that scales with their occupational AI exposure e_o , the complement of task complementarity $(1 - c_o)$, the scenario-specific displacement intensity \mathcal{J}_s , current AI capability A_t , and the average firm AI adoption \bar{a}_t . The leading constant calibrates the monthly base rate.

$$h_{i,t} = 0.036 \cdot e_{o(i)} (1 - c_{o(i)}) \mathcal{J}_s A_t \bar{a}_t$$

3.3 The funding rule

UBI generosity at full payout is $U^*_t = g \cdot \tilde{w}_0 \cdot n$. Tax revenue is collected at rate τ on the gross capital pool, which is sized so the capital share of total income equals s_t . The funded share ϕ_t is $\min(1, R_t / U^*_t)$; the actual UBI per worker is $\phi_t \cdot g \cdot \tilde{w}_0$. This is the binding policy constraint of the model — and, as Section 5 Case 2 shows, the constraint that determines which policy variants actually move outcomes.

$$U_t^* = g \tilde{w}_0 n; \quad R_t = \tau \left(\frac{s_t}{1 - s_t} \right) \sum_{i: e_{i,t} = 1} w_{i,t}$$

3.4 Monte Carlo design

Three AI scenarios — Conservative (intensity $\mathcal{J} = 0.6$, capability growth 0.008/mo), Moderate ($\mathcal{J} = 1.0$, 0.015/mo), and Aggressive ($\mathcal{J} = 1.6$, 0.028/mo) — are each crossed with five policy regimes (no UBI, fixed UBI, and three Section 8.2 variants). Each cell is a 60-seed Monte Carlo ensemble of 4,000 workers over 120 monthly ticks. Cross-seed mean and P10–P90 bands are reported for every output. Per-seed noise enters through initial wage draws (lognormal), displacement (Bernoulli per tick), reskill success (Bernoulli conditional on >18 months unemployed), and Gaussian noise on capability growth.


5. Case Studies

Three findings from the Monte Carlo. We treat each as a case study with a setup, a result, and an interpretation.

Case Study 1: The headline — UBI compresses, but cannot reverse

The first question is the simplest one, and the rest of the paper depends on its answer: does a fixed-rate, fixed-generosity UBI funded by a capital tax actually compress the income distribution under AI displacement, and does that compression scale with how aggressive the displacement turns out to be? Everything we test afterwards — the policy variants and the sensitivity sweeps — assumes a definite answer here, so this is where we set the foundation.

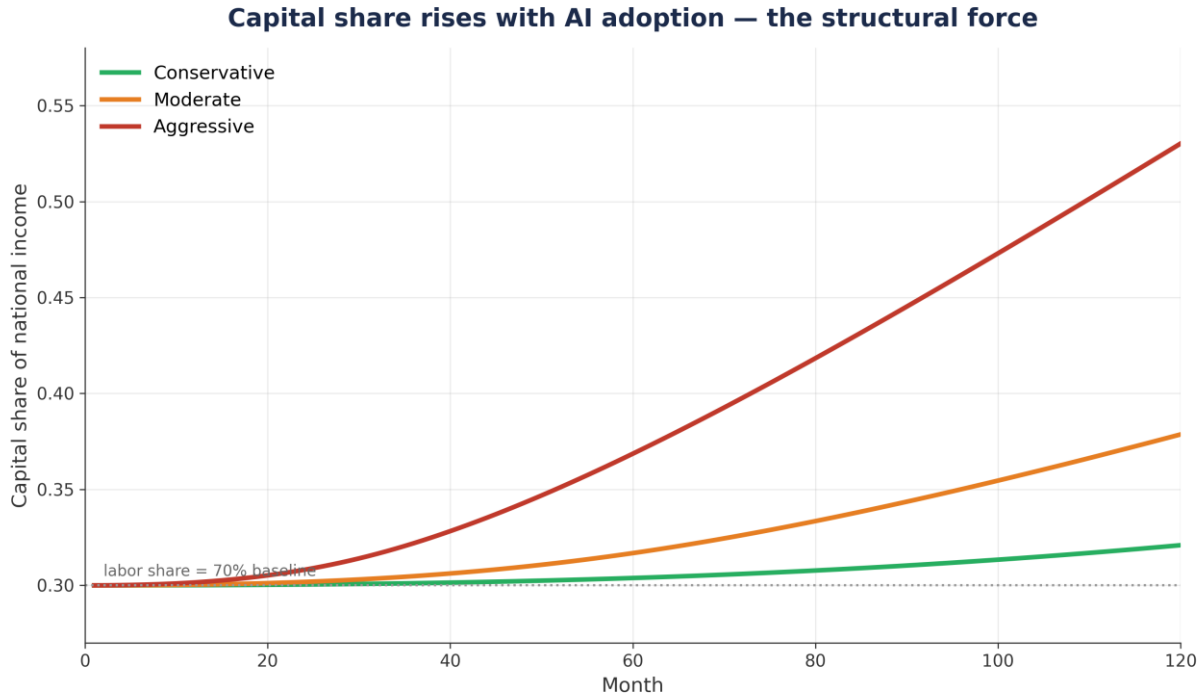
Setup. 60-seed ensembles under each scenario, with and without baseline UBI ($g = 0.30$, $\tau = 0.20$, start month = 12). Outcome: terminal Gini at month 120.



Gini coefficient over the 10-year horizon, by scenario. Black line: no UBI. Blue line: with baseline UBI. Shaded regions are P10–P90 across 60 seeds.

Result. Δ -Gini = -0.037 (Conservative), -0.045 (Moderate), -0.069 (Aggressive). The compression is real and statistically robust. But in Aggressive, terminal Gini still ends at 0.43 — above the US 2024 disposable-income baseline of 0.40, and still rising at month 120.

Interpretation. The capital share of national income is rising sharply. Conservative ends roughly where it starts; Moderate adds 8 percentage points; Aggressive adds 23. This is a structural transfer of national income from labor to capital that operates independently of the per-capita transfer.



Capital share of national income over the horizon. Conservative ends roughly where it starts. Moderate adds 8 percentage points. Aggressive adds 23.

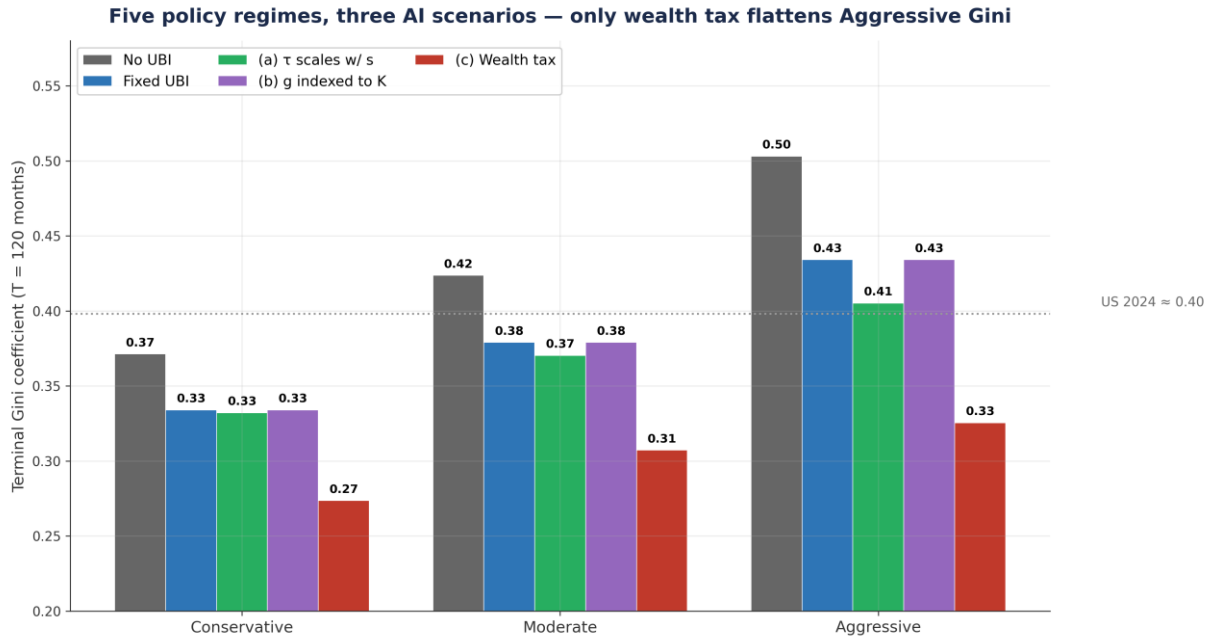
A 20% capital tax can redistribute about a fifth of that growing pool back to workers. In Aggressive, that's 10.6 percentage points of national income — enough to offset roughly half the labor-to-capital transfer, but not the whole thing. Fixed UBI is therefore a damper on a capital-share-rise inequality regime, not a cure.

Case Study 2: Three policy variants, one that works

Case Study 1 establishes that baseline UBI is a damper, not a cure, in the Aggressive regime: Gini falls relative to the no-UBI counterfactual but still rises in absolute terms over the decade. The natural next question is whether any modification of the basic design can do better. Earlier work (v1.0) sketched three candidate fixes that each target a different dimension of the problem; v1.1 simulates all three under the same protocol. The experiment doubles as a comparative test of where the binding constraint on flow-based redistribution actually sits.

Setup. We tested three modifications to the basic UBI design:

- **(a) τ -scaled tax** — $\tau_t = \tau_0 + \alpha(s_t - s_0)$: the capital-tax rate rises automatically with capital share. By month 120 in Aggressive, τ_t reaches 31% (from base 20%).
- **(b) capital-indexed generosity** — $U^*_t = g_K \cdot (\text{capital_pool} / n)$: generosity grows with the productive economy rather than with stationary median wage.
- **(c) wealth tax on stock** — each tick, 0.5%/month of every worker's capital share is redistributed equally — eroding the underlying ownership concentration directly.

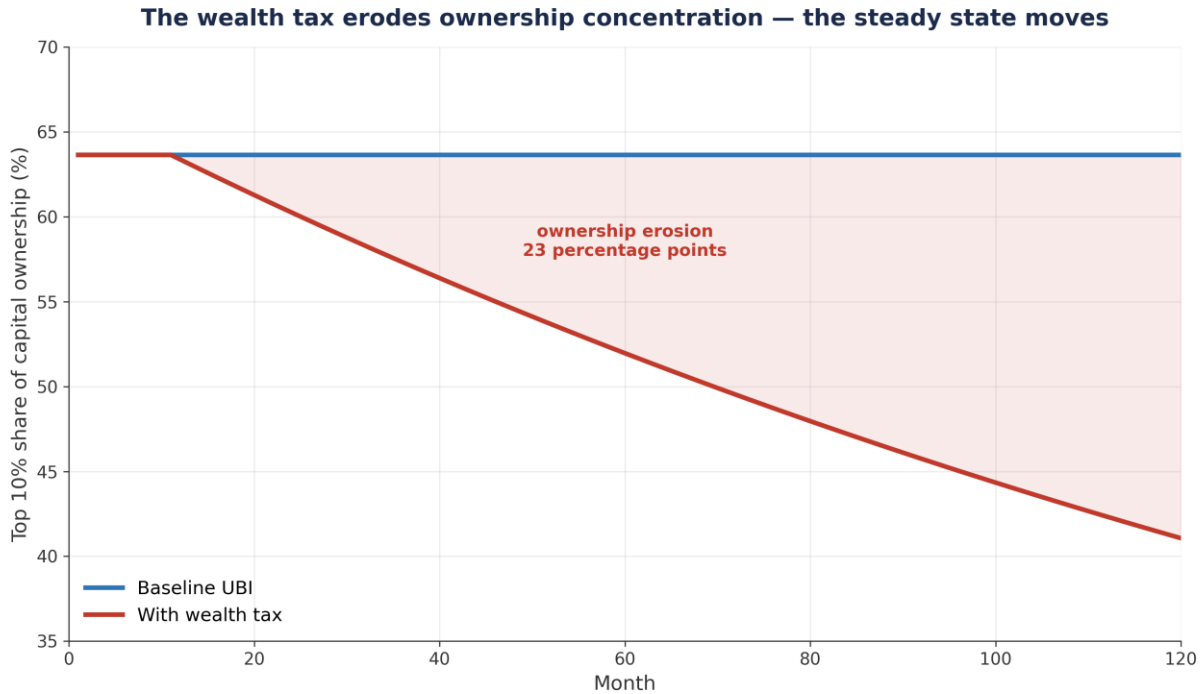


Terminal Gini at $T = 120$ across all five policy regimes and three AI scenarios. The wealth tax (red) is the only intervention that brings Aggressive-scenario Gini below the US 2024 baseline of 0.40.

Result. Variant (a) is a stronger damper but doesn't invert direction. Variant (b) is identical to baseline UBI within Monte Carlo noise. Variant (c) is the only policy that inverts long-run Gini — Aggressive ends at 0.33, with the trajectory peaking at month 71 and declining thereafter.

Interpretation. Variant (b) was the surprise: indexing target generosity to capital income looks structurally elegant on paper, but in this model it produces zero change in equilibrium. The reason is the funding rule in Section 3.3 — tax revenue, not the generosity target, is the binding constraint. Switching the index changes a number on the policy fact sheet but doesn't move the cash flow.

The deeper distinction is flow versus stock. Variants (a) and (b), and baseline UBI, all tax the flow of capital income while the underlying ownership concentration compounds independently. Variant (c) erodes the ownership concentration directly, and the flow follows.



Top 10% share of capital ownership in the Aggressive scenario, with and without the wealth tax. Under the wealth tax the concentration erodes by 23 percentage points over the decade.

Case Study 3: How robust is the redistributive payoff?

The headline numbers in Cases 1 and 2 sit on top of a single calibration: one Pareto α governing initial wealth concentration, one capital-share-drift trajectory per scenario, one set of behavioural UBI parameters. Before drawing any policy conclusions from those numbers, we need to know whether the rankings and the sign of the redistributive payoff are robust to plausible perturbations — or whether the finding is an artifact of the specific baseline values. This case study is the stress test, focused on the two parameters most exposed to calibration uncertainty.

Setup. 20-seed ensembles under $\pm 50\%$ perturbations on the two parameters most exposed to calibration uncertainty: capital-share drift δ_κ ($\pm 50\%$ of base value) and Pareto α governing initial wealth concentration ($\alpha \in \{0.75, 2.25\}$ bracketing the baseline 1.5).

Scenario	Baseline Δ -Gini	$\delta_\kappa \times 0.5$	$\delta_\kappa \times 1.5$	$\alpha = 0.75$	$\alpha = 2.25$
Conservative	-0.037	-0.035	-0.038	-0.058	-0.029
Moderate	-0.045	-0.038	-0.050	-0.068	-0.035
Aggressive	-0.069	-0.051	-0.072	-0.099	-0.057

Result. Δ -Gini sign is preserved in every cell. $\pm 50\%$ on capital-share drift moves the magnitude by ≤ 0.018 — the ranking and direction are robust. The Pareto α is the more consequential parameter: a heavier ownership tail ($\alpha = 0.75$) inflates Δ -Gini by $\sim 40\%$; a lighter tail ($\alpha = 2.25$) attenuates it by $\sim 20\%$.

Interpretation. The redistributive lever is proportional to the concentration against which it operates. Heavier wealth tails make every redistributive policy bigger, including the wealth tax. This is direct quantitative evidence for the Case Study 2 finding — stock-side intervention dominates flow-side intervention because concentration itself is the relevant state variable.

6. Conclusion

Three takeaways for anyone designing a redistributive response to AI-driven labor displacement: (i) UBI works as advertised in low-and-moderate displacement regimes, compressing Gini by 0.04–0.05 against the no-UBI counterfactual; (ii) in high-displacement regimes the structural transfer of income from labor to capital outpaces what any flow-based redistribution can offset, and a fixed UBI is mathematically a damper, not a cure; (iii) of the three policy modifications we tested, only the wealth tax on the underlying ownership stock inverts long-run Gini direction in our Aggressive scenario — and the model says it does so by 23 percentage points of top-10% wealth concentration over a decade.

The full white paper, the validation report (12/12 checks pass at Grade A), and the interactive dashboard are all at [AI Labor Simulation](#).

6.1 Future outlook

Three extensions are queued. First, a behavioural-parameter sensitivity pass — the reskill / reservation-wage / LFP-exit parameters are currently treated as designs rather than fitted moments, and tightening them matters for any quantitative policy claim. Second, a parameter sweep across the (α , δ_κ , intensity, τ , g) space — roughly 100,000 cells — to construct policy iso-curves rather than three calibrated points. Third, extending the model to capture cross-country migration and political-economy effects on the tax rate itself, which are explicit limitations of the current closed-economy formulation.

6.3 About Simudyne

Simudyne builds high-fidelity agent-based simulations for clients in financial services, climate finance, and policy research. Our platform lets quants and policy analysts build, validate, and deploy agent-based models without sacrificing the rigour expected by regulators. The work in this post — the AI-Labor ABM, its validation report, the architecture above, and the interactive dashboard — was built end-to-end on the Simudyne platform.

If you're working on a question where heterogeneous-agent dynamics matter — supply chains, energy markets, financial stability, climate-driven migration, labor-market policy, or anything else where the average agent isn't the relevant agent — we'd like to hear from you. Find me at ilan@simudyne.com.

About the author



Ilan Gleiser is Chief Strategy Officer at Simudyne, where he builds agent-based simulation models on the Simudyne platform, specializing in climate-financial risk, systemic contagion in financial and commodities networks, nonlinear economic dynamics, and — as of this paper — the labor-market and inequality effects of AI-driven displacement. His work focuses on questions where the macro-level answer depends on heterogeneity that representative-agent models smooth over: the wealth distribution that determines who owns the capital being redistributed, the occupational mix that determines who is exposed to displacement, the funding base that determines whether a policy can pay for itself in equilibrium.