AI Scenario Economics Explorer (Beta)

Romeo Dean

April 2026

Abstract

This document describes the AI Scenario Economics Explorer, a scenario-driven macroeconomic model for analyzing the impact of AI and robots on economic growth, factor prices, income distribution, and structural change over 2025–2040. The model is exogenous/scenario-driven rather than forecasting: the user specifies annual trajectories for AI and robot quantities, efficiencies, and capability frontiers, and the model solves the implied full-economy equilibrium each year. An accompanying interactive web application at `https://scenario-growth-model.vercel.app/` lets users construct and explore scenarios in real time. Acknowledgments: Tom Houlden, Tom Cunningham.

Contents

1. Introduction

This model is a largely exogenous (i.e., driven by user-specified inputs) scenario-driven economic growth model with additional extensions, including a model of AI and robot hardware production costs and other economic and societal outcomes. The core of the model is a growth model where user-specified trajectories for AI and robot quantities (number of copies deployed), efficiencies (human equivalents on automated tasks), and capabilities (% of 2024 non-automated tasks they can automate) combine with human cognitive and physical labor through production functions to produce economic output.

There are several limitations and problems with modeling the economy in this way, both inherent to growth models and—given the key author's lack of expertise in economics—probably with the specific approaches chosen here. Nevertheless, we find this a useful exploratory tool to add color to some of our scenario forecasting in ways that form a nice baseline. We don't expect the numbers here to be right, but we do think the model behaves reasonably on the presets. The implementation of the model and the website were done using Claude Code, and therefore it's very possible that there are bugs and problems, particularly with exploring custom scenarios.

Overall, our views about the future from an economics perspective are largely informed by other higher-level views and arguments about AI capabilities improvements and diffusion, rather than through any model, including this one.

Model inputs

	AI (AI Input tab)	**Robots (Robot Input tab)**
Quantity	AI copies deployed = public deployment H100e × copies per H100e	Robot units deployed
Efficiency	Human-equivalents per copy	Human-equivalents per robot
Capability frontier	% of 2024 non-automated cognitive tasks AI *can* do	% of 2024 non-automated physical tasks robots *can* do

Quantity × efficiency gives the effective labor supply from each technology. The capability frontier caps what tasks machines can perform; the model then determines how much automation *actually* occurs based on relative cost. Everything else—output, wages, factor prices, capital, distribution—is solved endogenously.

2. Model Presets

The app ships with two presets representing very different trajectories for AI deployment and regulatory response.

2.1 AI 2030 Plan A Scenario

A world where the US and allied nations coordinate to manage the AI transition through international regulation beginning around 2030. Key features:

- **International deal on R&D titration.** Under mutual transparency and auditing, the US, China, and other countries titrate R&D to increase safety and the likelihood that humans will be able to control the systems.[1] In the model this limits an out-of-control feedback loop of R&D reinvestment, so the fraction of effective labor allocated to software

[1] This might also reflect pursuing safer / human-interpretable / inspectable, albeit less efficient, directions—for example, macro-scale traditional-style robot designs as opposed to self-replicating or micro-scale robots.

and hardware R&D decreases, limiting Jones-channel cost improvements and preventing runaway recursive self-improvement.

- **Cap-and-trade on production quantities.** AI and robot deployment quantities are capped via a cap-and-trade permit system that limits scaling to roughly 6-month doubling times after 2032. Without such a constraint, the cost model's implied dynamics (rapidly falling production costs, massive surplus) would in theory allow reinvestment and exploding production quantities, which might be too destabilizing.
- **Human monitoring and auditing carveout.** Certain tasks are kept "unautomated" by policy—reserved for human oversight, auditing, and accountability—even where AI and robots dominate them on pure capability.[2] This preserves a non-automated task bucket that humans must fill, maintaining a wage floor and keeping humans structurally relevant to the economy in the Plan A scenario, despite being dominated capability-wise after 2035 in cognitive tasks and 2036 in physical tasks.
- **Citizen-dividend redistribution.** Cap-and-trade permit revenue plus taxation funds citizen dividends, redistributing the surplus from AI/robot deployment to households.

This is not a prediction: it's a scenario with strong international coordination to slow things down.

2.2 AI 2030 Default World

A world with no slowdown governance: no titrated R&D and no cap-and-trade regulation on AI or robot production. AI and robot quantities are determined endogenously by the model. Key features:

- **Endogenous quantities.** Each year the savings pool is invested proportionally between capital, AI, and robot hardware based on the ROI of each (marginal product per dollar invested), as determined jointly by the growth and cost models. AI and robot stocks accumulate from this investment.
- **No R&D throttling.** The R&D allocation that drives Jones-channel cost improvements is held flat rather than titrated down after 2032, so hardware and design costs continue to fall.
- **No cap-and-trade.** No permit system limits deployment. Surplus from deploying AI/robots feeds directly back into production, producing singularity-like dynamics. The simulation truncates at a singularity cutoff once output grows more than 1000× in a single year.
- **No regulatory capability pause.** The capability frontier continues improving rather than being held at a post-2035 plateau.

This is closer to what we would predict for the economic outcomes of the default capability progression we expect in the AI 2030 world absent the Plan A governance interventions.

[2] The most notable examples for the regime itself are AI researchers working on R&D under the regime (e.g. probably mostly working on alignment and control techniques), and auditors and inspectors for the verification regime. These are tasks that *could* be handed off to AIs (AIs are capable of these tasks) but they aren't sufficiently trusted yet.

3. The Production Function

3.1 Top-level production function

Output is produced by combining capital and effective labor:

$$Y_t = A \cdot \mathrm{CES}(K_t,\ L_{\mathrm{eff},t};\ \alpha,\ \sigma_Y) \quad \textit{output = TFP × CES of capital and effective labor} \tag{1}$$

where Y_t is total output (GDP), A is total factor productivity (calibrated once at the base year and held constant), K_t is the capital stock (machinery, buildings, data centers, all physical capital), α is the capital share of output, and $L_{\mathrm{eff},t}$ is effective labor. σ_Y is the top-level elasticity of substitution between capital and labor. With $\sigma_Y = 1$ (Cobb-Douglas, the default) this reduces to $Y_t = A \cdot K_t^{\alpha} \cdot L_{\mathrm{eff},t}^{1-\alpha}$.

Effective labor combines cognitive and physical labor via CES:

$$L_{\mathrm{eff},t} = \mathrm{CES}(L_{\mathrm{cog},t},\ L_{\mathrm{phys},t};\ \theta,\ \sigma_L) \quad \textit{effective labor = CES of cognitive and physical labor} \tag{2}$$

where $L_{\mathrm{cog},t}, L_{\mathrm{phys},t}$ are cognitive and physical labor aggregates, θ is the cognitive task weight (US 0.68, China 0.50, World 0.60), and σ_L is the elasticity of substitution between cognitive and physical labor ($\sigma_L = 1$ by default).

On σ_Y: empirical estimates for historical economies point below 1. Gechert et al. (2022), in a meta-analysis of 3,186 estimates, find $\sigma_Y \approx 0.3$ after correcting for publication bias; Oberfield & Raval (2021) estimate 0.5–0.7 for US manufacturing; Karabarbounis & Neiman (2014) estimated ∼1.25. We keep Cobb-Douglas ($\sigma_Y = 1$) as a neutral default; there is also a case for $\sigma_Y > 1$ in AI-specific contexts, since AI capital (compute, robots) substitutes more directly for cognitive and physical labor than traditional capital goods did.

On σ_L: direct estimates are rare. The closest proxies from the skilled/unskilled labor literature are generally above 1: Katz & Murphy (1992) $\sigma \approx 1.4$; Card & Lemieux (2001) $\sigma \approx 2.0$–2.5; Ciccone & Peri (2005) $\sigma \approx 1.5$. Kording & Marinescu (2025, Brookings) use $\sigma < 1$ at the cross-sector "physical vs. intelligence" level to generate saturation dynamics. We default to 1.0 as a neutral Cobb-Douglas benchmark.

3.2 Task-based automated vs. non-automated aggregation

Task universe, year-t. The set of economically relevant tasks in year t. Pre-2024 automated work (factory automation, CNC, software pre-2024) doesn't count — it's embedded in capital K.

Automatable $f(t) \in [0, 1]$. Exogenous capability frontier. The fraction of the year-t task universe that AI/robots are both *able* to perform (capability) and *allowed* to perform (policy). Within this fraction, one AI unit (measured in H100e compute) is treated as interchangeable with $e(t)$ human-equivalent labor-hours on average — a simplifying assumption, since real per-task efficiency varies widely.

Non-automatable $1 - f(t)$. The complement: tasks AI/robots cannot do at all, or are not allowed to do. (If AI could do them at any relevant efficiency, they'd be in the automatable bucket.)

Within the automatable bucket, humans and machines are **perfect substitutes** (additive in the leaf). Within non-automatable tasks, only humans can contribute. The two buckets aggregate via task-weighted CES:

$$L_{\mathrm{cog},t} = \mathrm{CES_{task}}(h_{\mathrm{auto}} + A_{\mathrm{eff}},\ h_{\mathrm{non}};\ f_c,\ \sigma_c) \quad \textit{cognitive nest: humans + AI in auto, humans only in non-auto} \tag{3}$$

$$L_{\mathrm{phys},t} = \mathrm{CES_{task}}(h_{\mathrm{auto}} + R_p,\ h_{\mathrm{non}};\ f_p,\ \sigma_p) \quad \textit{same structure for physical, with robots} \tag{4}$$

where $h_{\text{auto}}, h_{\text{non}}$ are human hours allocated to the automatable and non-automatable buckets, $A_{\text{eff}} = \text{H100e} \times e_c(t)$ is effective AI labor in human-equivalent units, R_p is effective robot labor (analogous), f_c, f_p are the cognitive and physical capability frontiers, and σ_c, σ_p are the across-task elasticities (default 0.8).

Automated $a(t) \in [0, f(t)]$. Endogenous. The share of automatable-task labor (in human-equivalent units) that AI/robots actually supply in equilibrium; the rest is supplied by humans working alongside them. The model determines this by allocating humans across the automatable and non-automatable buckets until their marginal product is equal in both — humans don't care which side they work on, so they spread until indifferent:

$$h_{\text{auto}} + h_{\text{non}} = H \quad \text{s.t. } \text{MP}_{\text{auto}} = \text{MP}_{\text{non}} \qquad \textit{humans allocate to equalize marginal product across buckets} \tag{5}$$

where H is total human labor (from eq. 7 below). When AI/robots are scarce, AI's rental is pulled up by that scarcity and the wage on automatable tasks is high enough that humans profitably "undercut" by working there too $\rightarrow a(t) < f(t)$. As AI scales and its rental falls, eventually it drops below the non-automatable wage even with no humans on the auto side; humans abandon automatable entirely and $a(t) = f(t)$.

The gap $a(t) < f(t)$ is the "AI could do this task, but humans can still undercut because AI is scarce" regime. $a(t) = f(t)$ is the "AI is so abundant it has fully crowded humans out of these tasks" regime.

3.2.1 Across-task elasticity σ_c, σ_p

How complementary are automatable and non-automatable tasks? To what extent can you compensate for missing labor on non-automated tasks with additional labor on automated tasks? This controls the bottleneck strength on non-automated tasks.

A concrete way to see it: if you lose one hour of non-automated labor, how many hours of automated labor do you need to add to keep output constant? The answer depends on both σ and the current automation level. Analytically, the marginal rate of technical substitution (MRTS) is

$$\text{MRTS} = \frac{1-f}{f} \cdot \left(\frac{L_{\text{auto}}}{L_{\text{non}}}\right)^{1/\sigma}$$

where f is the automation share and $L_{\text{auto}}/L_{\text{non}}$ is the ratio of labor in the two buckets. Plugging in representative ratios (≈ 0.25 at 20% auto, ≈ 1 at 50%, ≈ 10 at 90%, ≈ 100 at 99%, ≈ 1000 at 99.9%):

	$\sigma = 0.3$	$\sigma = 0.5$	$\sigma = 0.8$	$\sigma = 1.0$	$\sigma = 2.0$	$\sigma = 10$
20% automation ($L_{\text{auto}}/L_{\text{non}} \approx 0.25$)	0.04	0.25	0.71	1.0	2.0	3.5
50% automation (balanced)	1.0	1.0	1.0	1.0	1.0	1.0
90% automation ($L_{\text{auto}}/L_{\text{non}} \approx 10$)	239	11.1	2.0	1.1	0.35	0.14
99% automation ($L_{\text{auto}}/L_{\text{non}} \approx 100$)	4.7×10^4	101	3.2	1.0	0.10	0.02
99.9% automation ($L_{\text{auto}}/L_{\text{non}} \approx 1000$)	1.0×10^7	1001	5.6	1.0	0.03	0.002

Hours of automated labor needed to compensate for losing 1 hour of non-automated labor. The 50% row is 1.0 for every σ because both buckets are balanced (same task weight, same labor) — there's no marginal-product asymmetry for σ to amplify. Above 50%, auto labor is abundant so it has low marginal value: replacing a non-auto hour requires many auto hours, and low σ amplifies that. Below 50% the reverse holds: auto labor is the scarce bucket with high marginal value, so a fraction of an auto hour suffices, and low σ amplifies *that*. At $\sigma = 1$ (Cobb-Douglas) spending shares are constant so the MRTS reduces to $(1-f)/f$ times the labor ratio, the geometric average. At $\sigma \to \infty$ (perfect substitutes) the labor-ratio factor drops out and MRTS approaches $(1-f)/f$; we never reach this regime in our defaults.

3.2.2 Factor prices

All factor prices are marginal products derived through the nested-CES chain rule:

$$r = \alpha Y/K, \qquad w_c = \partial Y/\partial h_{\text{non},c}, \qquad q_c = \partial Y/\partial A_{\text{eff}} \quad \textit{factor prices = marginal products} \tag{6}$$

where r is the interest rate, w_c is the human cognitive wage (marginal product of humans on non-auto cognitive tasks; $h_{\text{non},c}$ denotes h_{non} in the cognitive nest), and q_c is the AI rental rate (marginal product of effective AI labor). Analogous definitions for the human physical wage w_p (marginal product on non-auto physical tasks) and the robot rental rate q_r (marginal product of effective robot labor). We have $w_c \gg q_c$ when AI is abundant, because non-automatable tasks are the scarce bottleneck.

3.3 Labor supply (endogenous LFP)

Total human labor supplied each year responds to wages and citizen dividends:

$$H_t = \text{WAP}_t \cdot \text{LFP}_t \quad \textit{human labor = working-age pop × participation rate} \tag{7}$$

$$\text{LFP}_t = \frac{1}{1 + \exp(-z_t)} \quad \textit{logistic participation rate} \tag{8}$$

$$z_t = z_{\text{baseline}} + \beta_w \ln(\bar{w}_t/\bar{w}_0) - \beta_u(\text{CD}_t/\bar{w}_0) \quad \textit{wages push LFP up, citizen dividends push down} \tag{9}$$

$$z_{\text{baseline}} = \ln(\text{LFP}_{\text{target}}/(1 - \text{LFP}_{\text{target}})) \quad \textit{pins LFP = target at base-year wage, zero citizen dividends} \tag{10}$$

where WAP_t is working-age population, LFP_t is the labor force participation rate, z_t is a logistic shift parameter, $\bar{w}_t$ is the best available wage that year, $\bar{w}_0$ is the base-year best wage, CD_t is the per-person citizen dividend that year, and β_w, β_u are sensitivities. Defaults: $\text{LFP}_{\text{target}}$ is region-specific (US 0.62, China 0.648, World 0.686); $\beta_w = 0.7$ in the Plan A preset (citizen dividends and the international regime cushion the wage→participation link) and $\beta_w = 1.5$ in Default World (no buffers, labor supply tracks wages more strongly); $\beta_u = 0.1$ (very low, calibrated for extreme citizen-dividend levels in Plan A, assuming inherent desire to work keeps enough people participating even under very generous dividends).

H_t feeds $h_{\text{auto}} + h_{\text{non}} = H$ in equation (5), closing the feedback loop: wages → LFP → H → task allocation → Y → next-period wages.

3.4 Capital accumulation

Capital evolves via the standard law of motion:

$$K_{t+1} = s(r_t) \cdot Y_t + (1 - \delta) \cdot K_t \quad \textit{Solow–Swan: investment + undepreciated stock} \tag{11}$$

where s is the investment rate, δ is the depreciation rate (default 5.3% for US), and r_t is the interest rate (the marginal product of capital, $r_t = \alpha Y_t/K_t$, so r_t = capital income / capital stock).

With endogenous savings (on by default), the investment rate responds to the interest rate via a constant-elasticity response on the gross return $1 + r$, following the log-linear specification used by Boskin (1978) in his empirical work on savings elasticity:

$$s(r) = s_{\text{base}} \cdot \left(\frac{1 + r}{1 + r_{\text{base}}}\right)^{\eta_s} \quad \textit{Boskin (1978) style endogenous savings} \tag{12}$$

s_{base} is the baseline investment rate (US 0.19, China 0.428, World 0.262). r_{base} is the interest rate in the base year (calibrated from the base-year production function). η_s is the savings elasticity (default 0.8).

Boskin (1978) estimated empirical savings elasticities of 0.2–0.4 for postwar US data, but those estimates come from economies where consumption is dominated by necessities and saving is constrained by subsistence needs. Our scenarios depart from that regime in two ways, both pointing toward a higher elasticity:

- **Vertically-integrated AI/robot producers.** In our preset scenarios, we think it is very possible that firms producing AI and robots will, instead of publicly deploying or selling them, plow their output and labor back into expanding their own production capacity rather than selling to outside consumers. This would effectively function as a firm-level savings rate near 1. In the extreme, we could see fully self-replicating robot factories, or a takeover scenario in which AI and robots scale themselves (an economy-wide savings rate close to 100%, with no human consumption at all). Even in a more normal case where output routes through distributed household income (even if very unequally), we think savings elasticity could still be higher in high-abundance scenarios (see next).
- **Luxury-shifted consumption.** As output and incomes grow drastically, as they do in our preset scenarios, we expect consumption to shift toward discretionary and luxury goods. We guess that discretionary consumption is more deferrable than necessities, so a larger share of output can be redirected into investment even if it is routing through households.

We chose $\eta_s = 0.8$ so endogenous savings can rise meaningfully with r and approximate such regimes in the limit, without saturating the 100% savings cap too aggressively at moderate takeoff-r values. This is a judgment call, not a direct empirical finding. Numerically, with $s_{\text{base}} = 19\%$ and $r_{\text{base}} \approx 14\%$: if returns hit 10× baseline ($r \approx 140\%$), savings rises to ∼35%; at $r \approx 500\%$ savings reaches ∼72%; the 100% cap binds around $r \approx 810\%$.

3.5 Limitations of the growth model

The production function and growth dynamics make several deliberate simplifications. Known limitations include:

Production function structure.

- **The task-based nested CES is a choice we are uncertain about.** AI substitutes for cognitive labor, robots for physical, with σ_c, σ_p controlling how complementary automatable and non-automatable tasks are. Alternative structures (more sectors, different nesting, non-CES functional forms) could give meaningfully different takeoff dynamics. The σ elasticities in particular are the most consequential parameters for the model's behavior in extreme scenarios.
- **Factor prices are marginal products, not forecasts of pay.** All prices (w_c, w_p, q_c, q_r, r) are derived through the nested CES chain rule. They are indicators of *productive contribution*, not forecasts of actual compensation: real pay depends on bargaining power, institutions, policy, labor-market frictions, and norms that the model does not capture.
- **Humans perfectly reallocate across tasks.** The task-bucket optimization assumes humans costlessly shift between automatable and non-automatable tasks each year. We probably overstate employment and labor-force participation as a result; in reality, frictions would slow reallocation and produce displacement even when aggregate employment could stay high.

Real vs nominal / composition of growth.

- **Real-output focused, 2025-dollar anchored.** The model tracks real output and real factor prices, with everything expressed in 2025-dollar terms. It abstracts from price-level dynamics and the composition of growth. In reality, a lot of the gain from AI may show up as *consumer surplus* (cheaper goods and services) rather than higher nominal GDP. Nominal growth during an AI takeoff could plausibly be much lower than the real output numbers this model produces.

Investment and how AI/robot quantities appear.

- **AI/robot quantities come out of thin air in Plan A.** In the exogenous mode used by the Plan A preset, AI and robot production quantities are user-specified trajectories: they appear without requiring capital investment to build them. The cost model performs a post-hoc accounting of how much investment *would have been required* to produce those quantities, letting the user sanity-check the numbers, but nothing in the growth loop enforces consistency with the economy's actual savings pool. In Default World (endogenous mode) this is partially fixed: the savings pool is split between capital, AI, and robots by ROI, and stocks only grow when investment flows in.
- **Ad-hoc savings response.** The Boskin-style power law is a reduced-form approximation, not derived from household preferences.
- **No forward-looking consumption decisions.** Households don't anticipate future returns or smooth consumption over time; they just save a fraction of current income based on the current interest rate. A proper optimizing model would have savings respond to *expected future* returns too, which could shift the timing of investment.
- **No capital markets, credit, or money creation.** Investment each year is funded entirely from the previous year's savings ($s \cdot Y_{t-1}$). There is no borrowing against future output, no leverage, no monetary expansion. In reality, during an AI takeoff firms would likely leverage heavily against anticipated productivity gains, potentially pulling forward a lot of the investment our model stretches out over years. Real interest rates could also spike far above what our endogenous-savings response captures; see Chow, Halperin & Mazlish (2023) for a theoretical argument that rates above 100% annually are plausible under rational expectations of transformative AI.
- **No foreign investment flows between regions.** The US, China, and World simulations run independently; real-world capital flows between regions are not modeled.

4. Cost Model

Core question: How much does it cost to produce AI chips and robots?

Production costs decline through two multiplicative channels:

$$\text{Cost } \$ = (\text{Cost } \$)_{2025} \times \text{Wright } (\$/\text{unit produced}) \times \text{Jones } (\text{quality/unit}) \tag{13}$$

Wright's Law drives down the manufacturing cost per unit (learning-by-doing, economies of scale). Jones drives down the design cost per unit of quality-adjusted capability (R&D makes each unit you produce, even at fixed quantities, more capable). Both indices start at 1.0 in 2025 and decline over time, multiplicatively reducing cost.

Motivation

In AI chips there are two dynamics we want to capture separately:

1. **Pure chip design and quality** (Jones channel): how much compute performance can you extract per fixed unit of manufacturing, e.g. transistors/compute per wafer. Driven by R&D into chip architecture, process nodes, and design.
2. **Pure production cost** (Wright channel): how cheap is it to actually manufacture a wafer, e.g. AI and robots improving the cost of raw materials, manufacturing inputs, fab labor, etc. Driven by cumulative production experience and scale.

The same two-channel split applies to robots (Wright = cost per physical unit; Jones = capability per unit).

4.1 Wright's Law (manufacturing learning)

Each doubling of cumulative production reduces manufacturing cost by a fixed percentage (Wright 1936). One of the most robust empirical regularities in economics: it has held across semiconductors, solar panels, batteries, and many other technologies.

$$\text{Wright}_{i,t} = (Q_{\text{cum},t}/Q_{\text{cum},0})^{-b} \quad \textit{each doubling of } Q_{\text{cum}} \textit{ cuts cost by } (1-2^{-b}) \tag{14}$$

Unit cost falls by a fixed fraction $(2^{-b} - 1)$ with each doubling of cumulative production Q_{cum}. b is the learning-rate exponent (higher $b \to$ faster cost decline). $Q_{\text{cum},t}$ is cumulative production through year t (total units ever produced); $Q_{\text{cum},0}$ is its base-year value. Our rough estimates: AI chips $b = 0.10$ (post-28nm era, cost-per-transistor improvements have slowed sharply, about 7% per doubling); robots $b = 0.20$ (earlier on the experience curve, about 13% per doubling).

4.2 Jones (1995) design improvement

This is the model's key mechanism for how AI feeds back into its own cost decline: R&D labor generates new design ideas that reduce cost.

$$\dot{D}_{i,t} = r_{\text{base}} \cdot (L_{\text{design},t}/L_{\text{design},0})^{\lambda} \cdot D_{i,t}^{1-\phi} \quad \textit{Jones design improvement: } \lambda \textit{ toes, } \phi \textit{ fishing-out} \tag{15}$$

$\dot{D}_{i,t}$ is the time derivative of the design index (equivalently, the within-year change). r_{base} is the base-year improvement rate. $L_{\text{design},0}$ is base-year R&D labor. $\lambda < 1$ captures "stepping on toes" (duplicated effort, diminishing returns to more researchers: $10\times$ researchers $\to 10^{\lambda}$ faster progress; with $\lambda = 0.2$, only $1.6\times$). ϕ controls "ideas getting harder to find": as the design stock $D_{i,t}$ grows, each new idea contributes proportionally less. The key steady-state ratio is $\lambda/(1-\phi)$, the "OOMs of progress per OOM of researcher growth"; Bloom, Jones, Van Reenen & Webb (2020) find $\sim$5 for semiconductors.

$$L_{\text{design},t} = L_{\text{eff},t} \cdot \text{frac}_{\text{R\&D},t} \quad \textit{R\&D workforce} = L_{\text{eff}} \times \textit{exogenous allocation \%} \tag{16}$$

R&D workforce equals effective labor times the user-specified R&D allocation fraction. In the Plan A Scenario preset, governance titrates R&D heavily, so this fraction shrinks significantly after 2032. In Default World it rises and stays flat at high levels post-2030. **The feedback loop:** more AI deployed $\to L_{\text{eff}}$ grows $\to$ more R&D labor $\to$ faster design improvement $\to$ cheaper AI. But with $\lambda = 0.2$, the feedback has strong diminishing returns.

Combined equations

$$\text{MC}_{i,t} = \text{MC}_{i,0} \cdot \text{Wright}_{i,t} \cdot \text{Jones}_{i,t} \quad \textit{marginal cost = base-year MC scaled by two channels} \tag{17}$$

Marginal cost ($\text{MC}_{i,t}$) for sector $i \in \{\text{AI}, \text{robot}\}$ is the base-year $\text{MC}_{i,0}$ (user-specified) scaled by manufacturing learning and design improvement. Both indices start at 1.0 in 2025.

$$P_{i,t} = \text{MC}_{i,t} \cdot \mu_{\text{mfg},t} \cdot \mu_{\text{design},t} \quad \textit{price = cost} \times \textit{user-specified markups} \tag{18}$$

Price $P_{i,t}$ equals marginal cost times markups; $\mu_{\text{mfg},t}$ and $\mu_{\text{design},t}$ are user-specified manufacturing and design markup trajectories.

$$\text{Surplus}_{i,t} = (q_{i,t} - P_{i,t}) \cdot Q_{i,t} \quad \textit{buyer surplus = (value − price) × quantity} \tag{19}$$

Buyer surplus is value ($q_{i,t}$, the marginal product of the technology from the growth model) minus price $P_{i,t}$, times quantity $Q_{i,t}$ deployed. Total surplus is split three ways: *producer surplus* (markup above MC), *buyer surplus* (deployers' gap between value and price), and *government* (taxes + cap-and-trade permit revenue, which funds citizen dividends).

Parameter defaults & calibration

Param	Default	Description
b_{AI}	0.10	AI chip learning rate. Post-28nm era. Cost per transistor stopped declining after ∼2011; wafer costs rose sharply with EUV lithography. Historical (pre-28nm): $b \approx 0.42$.
b_{robot}	0.20	Robot learning rate. Higher than AI chips because humanoid robots are early on the experience curve (analogous to where semiconductors were in the 1970s). Uncertain due to minimal cumulative production to date.
$r_{\text{base,AI}}$	0.26	AI hardware design base rate: 26%/yr ≈ Moore's Law (1.35×/yr). Anchors where the Jones trajectory starts.
$r_{\text{base,robot}}$	0.05	Robot design base rate: 5%/yr. Rough estimate; industrial robots historically improved at 2–3%/yr, but humanoid robots may improve faster.
λ_{AI}	0.20	AI stepping-on-toes. From Sequeira & Neves (2020): $\lambda \approx 0.2$ across many domains. 10× researchers → 1.6× faster progress.
λ_{robot}	0.20	Same as AI; no strong reason to differentiate.
ϕ_{AI}	0.90	AI fishing-out. $\lambda/(1-\phi) = 0.2/0.1 = 2.0$, below the historical steady-state ratio of ∼5.
ϕ_{robot}	0.933	Robot fishing-out. $\lambda/(1-\phi) \approx 3.0$: slightly above agriculture/aggregate-TFP but below semiconductors.

How to think about $\lambda/(1-\phi)$. In the Jones (1995) semi-endogenous growth framework, this ratio is "OOMs of progress per OOM of researcher growth" in steady state. Bloom, Jones, Van Reenen & Webb (2020) found ∼5 for semiconductors: every 10× increase in researcher count sustained ∼100,000× cumulative progress. Our forward-looking parameter gives 2.0 for AI, reflecting an expectation that the historical pace will slow for semiconductors. We use 3.0 for robots, chosen as a middle point between aggregate-TFP/agriculture (∼3) and semiconductors (∼5). There is no direct Bloom-style estimate for robotics.

Cap-and-trade

The Plan A scenario includes a cap-and-trade regime where governments issue a certain number of permits to produce restricted goods (including AI chips and robots) in order to restrict overly destabilizing production levels. Compared to status-quo economics today, these caps are set very high, allowing robots to double every 6 months and AI chips to grow 5× per year in early years, but then flattening after 2035 when scaling is paused for other reasons in the Plan A scenario. Each new AI chip or robot requires a permit, and the model prices these permits as a fraction of its lifetime surplus implied by the growth and cost models (value produced minus cost paid, over its useful life). The fraction is an "auction efficiency" parameter (default 0.70, because we assume high levels of competition and market pricing in of the usefulness of the AIs/robots). Government revenue from the permits funds citizen dividends. With cap-and-trade enabled, the AI/robot quantities specified in the growth model are automatically treated as the quantities permitted by the cap-and-trade regulation.

5. Energy and Climate

We added an exploratory energy model that has (i) total primary energy demand, (ii) the bottom-up energy consumption of AI compute and robots, (iii) climate externalities (waste heat, CO_2 emissions), and (iv) naive energy cost modelling by type of generation source (CapEx, OpCost, LCOE) and an endogenous function to decide what energy gets built over time. It also implements optional carbon credit auctioning, where in the desired year, the cost to offset emissions of each generation source is added to the cost of that source. Like the cost model, it leaves the growth-model factor prices untouched, it is not a part of the core model.

5.1 Total primary energy (top-down, GDP-anchored)

Global energy consumption in 2025 was around 19.0 TW-yrs. In the Plan A scenario, there is $\sim$270× cumulative GWP growth by 2040 (on our default model assumptions). If the energy intensity of the economy (ratio of GDP to energy consumption) stayed constant, then energy consumption would also increase by 270× — but historically we have observed a decrease in energy intensity since 1920.

We expect this energy decoupling to continue, because the drivers of growth (AI, robots, and accompanying capital) should all become more energy efficient per unit of economic output (AI hardware in particular has a strong power-efficiency trend). In our energy model, we assume the energy-decoupling trend of the last few decades will continue, and energy consumption will grow at roughly 40% the pace of economic growth (we call this energy growth / GDP growth the *energy elasticity of GDP*, $\varepsilon = 0.4$). We are highly uncertain about this value, and find values between 20% and 70% plausible (80% CI, conditional on this economic growth trajectory).

$$E_t = E_{2025} \cdot \left(\frac{Y_t}{Y_{2025}} \right)^{\varepsilon} \tag{20}$$

An energy elasticity of GDP of 40% leads to energy consumption on Earth growing around 9× during Plan A, bringing total primary energy to 180 TW-yrs by 2040 (the 20% case lands at 60 TW-yrs, the 70% case at 960 TW-yrs).

5.2 AI compute and robot energy, bottom-up

For AI compute and robots we separately model their power efficiency and combine with the stocks from the growth model to get their total power draw. We assume AI compute draws roughly 1,000 W per H100-equivalent at 2025 efficiency, and a human-equivalent robot draws around 4,000 W. We then assume both per-unit numbers improve over time at a Jones-style rate driven by the same R&D allocations as the cost model — so AI-driven progress on chip cost and AI-driven progress on chip power efficiency are tied together.

In the Plan A scenario this produces AI compute energy that grows from $\sim$10 GW-yrs in 2025 to around 26 TW-yrs by 2040, with robot energy around 33 TW-yrs — so AI plus robots is roughly 33% of total energy by 2040, up from less than 1% today. "Other" (industry, transport, buildings, residential) is the rest.

5.3 Climate externalities

Depending on the energy source, this scale-up has implications for the Earth's surface temperature. There are two separate things to consider: (1) a net increase in waste heat on Earth, and (2) emissions — particularly CO_2 — leading to temperature increases through radiative forcing.

Waste heat. Solar panels convert energy that was arriving on Earth anyway from the sun, so they only change Earth's net energy balance if they meaningfully change the planet's albedo. Nuclear and fossil generation, on the other hand, create a direct increase in waste heat by releasing energy that was previously locked up in the Earth's crust. That said, even at the higher end of levels we might reach by 2040 in Plan A, waste heat on its own does not become a major consideration. A back-of-envelope: 50 TW of fossil and nuclear waste heat spread over Earth's surface is roughly $0.1\,\mathrm{W/m^2}$ of additional forcing, against $\sim 3\,\mathrm{W/m^2}$ from accumulated CO_2. We drop heat and albedo externalities from the model accordingly.

Carbon emissions. If fossil fuels were used as the sole source for this energy scale-up, carbon emissions would become a significant problem by 2040 absent significant mitigation (e.g. direct air capture), with equilibrium surface temperature rising to roughly +3.0°C over pre-industrial levels, up from +1.8°C today.

We therefore think there should be CO_2 emission mitigation policies agreed to globally. Direct air capture (DAC) provides an affordable path to mitigation, especially with help from AI and robot labor during the 2030s — the thermodynamic floor for separating CO_2 from the atmosphere is around 0.034 MWh per tonne, and engineered systems get within $\sim 15\times$ of that today (Climeworks Mammoth, ~\$600/tCO_2). We assume DAC costs \$30/tCO_2 by 2040.

The carbon policy the model can support is:

- **Private cap-and-trade market (no government subsidy needed).** Any fossil emitter globally is required to offset 100% of emissions with equivalent carbon-capture credits. DAC operators earn revenue by capturing carbon and selling credits to emitters. Reaches net-zero emissions but does not draw down legacy CO_2.

The cap-and-trade regime is priced into the energy mix by adding the per-source lifecycle emissions ($\times$ the CO_2 price) onto each source's marginal cost. By default, the model enacts the cap-and-trade regime in 2035 with credits priced at the DAC cost curve; in the Default World scenario, policy is assumed not to bind and no carbon price is charged.

5.4 Energy cost model

Each generation source has cost modelling for CapEx (\$/W-delivered) and OpCost (\$/MWh), naively estimated from today's levels and modelled forward. Levelized cost is the standard formula, with both capex and opex amortized at a real WACC of 7% (interpreted as the $r - g$ risk premium):

$$\mathrm{LCOE}_{i,t} \;=\; \mathrm{CapEx}_{i,t} \cdot \mathrm{CRF}_i \cdot 114.16 \;+\; \mathrm{OpCost}_{i,t} \;+\; \mathrm{externality}_{i,t} \tag{21}$$

where $\mathrm{CRF}_i = \mathrm{WACC}/(1-(1+\mathrm{WACC})^{-\mathrm{Life}_i})$ is the capital recovery factor and 114.16 converts \$/W-delivered·year to \$/MWh.

We feed in CapEx and OpCost separately rather than a single LCOE because they evolve differently. The 2025 anchors produce a stock-weighted average of around \$63/MWh, which matches the world energy bill divided by TWh consumed.

5.5 Endogenous mix dynamics

Each year, new capacity is allocated across sources by a softmax over LCOE: cheaper sources get most of the new builds, but not all of it (the temperature parameter governs how aggressively the cheapest source dominates). The market price each year is the stock-weighted average LCOE across the remaining fleet.

In Plan A with the 2035 cap-and-trade regime, this produces a fairly sharp transition: by 2032 solar and wind dominate new builds but the deployed mix is still mostly legacy fossil, which is already getting outcompeted by scaling solar/wind/nuclear on our naive cost model, then in 2035 carbon credits fire, fossil's marginal cost spikes from ∼$50 to ∼$116/MWh with the DAC caputre requirement, while renewables are at ∼$30/MWh, and this retires almost all fossil within a single year. By 2040 the mix is approximately 100% solar / wind / nuclear.

Total energy investment peaks at 6–8% of GWP in 2032–2034 (demand growth and fossil-replacement build-out fire simultaneously), then falls as the CapEx Jones drives solar and wind below $1 per W-delivered.

5.6 Limitations of the energy and climate model

- **Cost modelling is extremely naive.** Each generation source's CapEx and OpCost trajectories are propagated forward with Jones-style parameters that we don't have good data or well-justified values for. The cost trajectories should be read as illustrative rather than calibrated.
- **DAC cost is also not well justified.** The trajectory from ∼$600/t$CO_2$ today to $30/t$CO_2$ by 2040 is anchored on Climeworks today and a thermodynamic-floor BOTEC, with the rate of progress in between essentially guessed.
- **Top-level primary energy is naive.** The single-ε decoupling formula extrapolates a historical trend forward with no structural model of where energy is actually consumed. The implied non-AI / non-robot power consumption (industry, transport, buildings, residential) — which we just back out as the residual between total energy and the bottom-up AI+robot energy — might not be in a reasonable range for a singularity-era economy.
- **AI and robot power efficiency is uncertain but somewhat better grounded.** The Jones-style trajectories on W per H100-equivalent and W per HE-robot are also uncertain, but the AI compute power-efficiency trend has been clear and consistent for decades, so we at least have a sense of what to anchor to there.

6. Extensions (scenario exploration)

Exploratory add-ons to the growth and cost model that add more scenario color.

6.1 Endogenous AI and robot quantities (Default World)

Unlike the other extensions in this section (which are post-hoc layers that leave the core growth model untouched), this one *replaces* the exogenous AI and robot trajectories the user draws in the input tabs with an investment-allocation rule. It is what the *AI 2030 Default World* preset runs.

Investable savings pool. Each year the same pool $s(r_t) \cdot Y_t$ that funds capital in (11) now funds all three of capital, AI hardware, and robot hardware. Under *Plan A*, 100% of the pool goes to capital and AI/robot quantities follow the user trajectories, with a post-hoc display of what investment that requires in the cost model; under *Default World*, the pool is split across K, AI, and robots in proportion to a naive myopic **return on investment**, defined as the annual marginal product at current rates per $1 invested. Because of the explosive growth these scenarios experience, we think this is a decent proxy—the "useful life" of the new production being greater than 1 year may even be too long if anything, given the likely speed of change of the technology.

$$\mathrm{ROI}_K = r, \qquad \mathrm{ROI}_{\mathrm{AI}} = \frac{q_c}{\mathrm{MC}_{\text{per AI HE}}}, \qquad \mathrm{ROI}_{\mathrm{robot}} = \frac{q_r}{\mathrm{MC}_{\text{per robot HE}}} \qquad \textit{marginal product per \$1 invested} \tag{22}$$

$r = \alpha Y/K$ is the interest rate; q_c, q_r are the AI and robot rental rates from (6); $\mathrm{MC}_{\text{per AI HE}}$ and $\mathrm{MC}_{\text{per robot HE}}$ are the cost model's lifetime marginal costs per human-equivalent unit, which decline over time via Wright's Law (14) and Jones R&D (15).

Allocation rule. Shares are set proportionally to ROI:

$$\phi_K = \frac{\mathrm{ROI}_K}{\sum \mathrm{ROI}}, \quad \phi_{\mathrm{AI}} = \frac{\mathrm{ROI}_{\mathrm{AI}}}{\sum \mathrm{ROI}}, \quad \phi_{\mathrm{robot}} = \frac{\mathrm{ROI}_{\mathrm{robot}}}{\sum \mathrm{ROI}}.$$

The combined AI+robot share $\phi_{\mathrm{AI}} + \phi_{\mathrm{robot}}$ is capped at the average automation frontier $(f_c + f_p)/2$ of tasks in the economy — which we use as a proxy for the cap on what the economy can reinvest into AI and robot production.

Within-year equilibration. Pure start-of-year allocation is winner-take-all: whichever factor happens to have the highest ROI that year absorbs nearly all investment, and shares flip abruptly year-over-year as the leading factor saturates. To smooth this and better approximate market equilibrium, within each year we iterate: pick shares, advance stocks by the implied investment, re-solve the production function at the *post*-investment stocks, recompute ROIs, and update shares with damping. Five iterations typically reduce the spread between the three ROIs to within 5%.

Stock accumulation. After the shares converge, stocks update:

$$\mathrm{AI}_{t+1} = (1 - \delta_{\mathrm{AI}})\,\mathrm{AI}_t + \phi_{\mathrm{AI},t} \cdot s(r_t) Y_t \,/\, \mathrm{MC}_{\text{per AI HE},t} \qquad \textit{AI stock: prev stock less depreciation, plus new HEs bought} \tag{23}$$

$$R_{t+1} = (1 - \delta_R)\, R_t + \phi_{\mathrm{robot},t} \cdot s(r_t) Y_t \,/\, \mathrm{MC}_{\text{per robot HE},t} \qquad \textit{robot stock, same structure} \tag{24}$$

$$K_{t+1} = (1 - \delta) K_t + \phi_{K,t} \cdot s(r_t) Y_t \qquad \textit{capital: same law as (11), with share scaling} \tag{25}$$

In Plan A (exogenous quantities), $\phi_K = 1$ and (25) reduces to (11).

Singularity cutoff. If $Y_{t+1}/Y_t > 1000$ in any year, the simulation stops. This is a numerical safeguard against runaway takeoff: by the time the model is producing such growth, the underlying assumptions (Wright learning, fixed markups, factor-share CES structure) are well outside their domain of validity.

6.2 Goods vs. services decomposition

This add-on models the relative prices of goods and services by assuming a certain relative cognitive and physical labor intensity for each, together with a price-sensitive consumer demand system that splits spending between them as relative prices evolve over time.

Given the factor prices (w_c, w_p, q_c, q_r) and Y from the main solver, output is split into two categories with different cognitive/physical intensity: services (cognitively heavy, $\theta_S = 0.80$) and goods (more physical, $\theta_G = 0.40$). The module outputs **relative prices** P_S/P_G within each year (a pure within-year comparison, with the CES aggregate index normalized to 1 each year since the model does not track absolute price levels over time). The corresponding consumer expenditure shares come from a CES demand system with elasticity $\eta = 0.5$ (the default, at the upper end of the 0.2–0.5 range Comin, Lashkari & Mestieri (2021) estimate across specifications in postwar cross-country data, and consistent with the broader structural-transformation literature: Ngai & Pissarides 2007; Herrendorf, Rogerson & Valentinyi 2013).

This was added with the idea in mind that if AI automates cognitive work faster than robots automate physical work, services might become relatively cheaper than goods at some aggregate level. Notably we don't do any cost modelling, so this does not model actual price levels over time. Being a post-hoc add-on, the core growth-model outputs are unaffected.

Equations. Let p_{cog} be the unit cost of a cognitive labor hour (a CES blend of humans and AI), p_{phys} the unit cost of a physical labor hour, and $p_{\text{eff},j}$ the effective labor unit cost for sector j using that sector's cognitive weight θ_j. These come for free from the production function via Shephard's lemma (any nested CES quantity aggregator has a matching nested CES price aggregator).

$$\nu = \frac{a_S \cdot P_{G,0}^{1-\eta}}{a_S \cdot P_{G,0}^{1-\eta} + (1-a_S) \cdot P_{S,0}^{1-\eta}} \quad \textit{calibrate } \nu \textit{ so base-year service share} = a_S \tag{26}$$

$$P_S \propto (p_{\text{eff},S}/p_{\text{eff,agg}})^{1-\alpha}, \quad P_G \propto (p_{\text{eff},G}/p_{\text{eff,agg}})^{1-\alpha} \quad \textit{sector prices from relative labor costs} \tag{27}$$

$$S = \nu \cdot P_S^{-\eta} \cdot Y, \quad G = (1-\nu) \cdot P_G^{-\eta} \cdot Y \quad \textit{CES demand: cheaper sector gets more quantity} \tag{28}$$

where ν is calibrated once at the base year to match empirical services spending share a_S; P_S, P_G are rescaled each year so the CES aggregate index equals 1 (a pure numeraire choice); S, G are real quantities consumed. With $\eta = 0.5$, a 1% drop in relative service price raises service quantity by only 0.5%, so spending on services actually falls (services and goods are mostly complements).

6.3 Land

The land add-on models demand for four land categories (urban, rural, agricultural, and commercial), each with its own starting expenditure share and its own elasticities. Each category has an income elasticity β (how demand responds to rising incomes), a demand elasticity η^D (how demand responds to price), and a supply elasticity η^S. Whether land expenditure grows faster, slower, or in line with overall consumption depends mostly on β: $\beta > 1$ luxury (share rises with income), $\beta < 1$ necessity (share falls), $\beta = 1$ proportional. Defaults: urban and rural residential $\beta \approx 1.1$ (slight luxury), agricultural low β (necessity), commercial in between. A wilderness-deregulation option releases protected land after a configurable year, loosening the supply constraint. Each period the model clears each category's land market for its rental rate and land-use share.

We have land as a post-hoc add-on as opposed to a factor in the production function, because our view is that land is unlikely to be a binding constraint on AI/robot production except through regulation: land-use footprints even under extreme growth scenarios are small relative to available land on Earth. The more relevant and interesting thing about land would be the availability and pricing of land on the consumer side, which is what this extension explores.

Equations. For each endogenous category i (urban, rural, agricultural), let v_i be land rent per hectare, M_i acres supplied/demanded, R_i total rent spent on that category, X total consumer expenditure, with $v_{i,0}, M_{i,0}, X_0$ the base-year values.

$$X_t = (1 - s(r_t)) \cdot Y_t \quad \textit{consumer expenditure = output − savings} \tag{29}$$

$$M_i^D \propto X^{\beta_i} \cdot v_i^{-\eta_i^D}, \qquad M_i^S \propto v_i^{\eta_i^S} \quad \textit{demand and supply shapes per category} \tag{30}$$

$$\frac{v_i(t)}{v_{i,0}} = \left(\frac{X_t}{X_0}\right)^{\beta_i/(\eta_i^D+\eta_i^S)} \quad \textit{equilibrium rent, relative to base year} \tag{31}$$

$$\frac{M_i(t)}{M_{i,0}} = \left(\frac{v_i(t)}{v_{i,0}}\right)^{\eta_i^S} \quad \textit{acres from supply curve at new rent} \tag{32}$$

$$R_i(t) = v_i(t) \cdot M_i(t) \quad \textit{total rent = rent/acre × acres} \tag{33}$$

If total endogenous acreage exceeds the physically available land (after removing protected wilderness and commercial floors), all endogenous acres are scaled proportionally and prices are re-solved from the demand curve, so scarcity raises rents.

US defaults: urban $\beta = 1.1, \eta^D = 1.0, \eta^S = 0.5$, starting share 89%; rural $\beta = 1.1, \eta^D = 1.0, \eta^S = 0.8$, starting share 7%; agricultural $\beta = 0.1, \eta^D = 0.4, \eta^S = 0.6$, starting share 4%. Overall land expenditure share is 6.8% of household spending at base year. Commercial land is exogenous: its acreage is held fixed at the base-year value, functioning as a set-aside that reduces land available for the three endogenous categories rather than a market the model clears.

6.4 Income distribution and inequality

In this add-on, households are binned into percentiles (1, 25, 50, 75, 90, 99, 99.9) that each own different shares of the five income sources: wages, capital, AI, robots, and land. **Ownership shares across percentiles are exogenously set by the user** in the Income Distribution section of the Parameters panel; the defaults mirror real-world wealth concentration. The module does not endogenize ownership changes over time, it just treats the exogenous distributions as fixed.

Output charts.

- **Income Distribution** — per-person income by percentile over time.
- **Income Composition** — at each percentile, what fraction of income comes from wages vs. capital vs. AI vs. robots vs. land.
- **Net Worth Distribution** — accumulated wealth per percentile.
- **Gini Coefficients** — overall and by source, showing how AI shifts inequality.

Equations. For each factor $j \in \{\text{wage, capital, AI, robot, land}\}$, the user specifies a cumulative ownership CDF at 7 percentile breakpoints (p1, p25, p50, p75, p90, p99, p99.9) and a finer multiplier curve at 9 percentiles. The CDF sets coarse ownership shares; the multiplier curve is a smooth fill-in used to draw distribution charts.

$$\text{share}_j(p) = \text{CDF}_j(p + dp) - \text{CDF}_j(p) \quad \textit{ownership share per percentile, from CDF} \tag{34}$$

$$m_j(p) = \text{PCHIP interpolation of user-specified multiplier breakpoints} \quad \textit{within-bucket income multiplier} \tag{35}$$

$$I_j(p) = m_j(p) \cdot \text{share}_j(p) \cdot \text{TotalIncome}_j \quad \textit{per-person income at p from factor j} \tag{36}$$

$$I_{\text{total}}(p) = \sum_j I_j(p) + \text{CD}(p) + \text{FA}(p) \quad \textit{total income = factors + domestic citizen dividend + foreign aid received} \tag{37}$$

$$W_{t+1}(p) = W_t(p) + s(r_t) \cdot I_{\text{total},t}(p) \quad \textit{net-worth accumulation at economy-wide savings rate} \tag{38}$$

$$\text{Gini} = 1 - 2\int_0^1 L(p)\,dp \quad \textit{standard Gini from Lorenz curve L(p)} \tag{39}$$

where TotalIncome_j is the aggregate factor-j income from the growth model (rK for capital, $q_c \cdot A_{\text{eff}}$ for AI, etc.), $\text{CD}(p)$ is the per-person domestic citizen dividend funded from the region's own tax revenue, $\text{FA}(p)$ is the per-person foreign aid received from other regions' tax revenue (zero in donor regions), and $L(p)$ is the Lorenz cumulative income share.

Top-1% shares implied by the US defaults: wages 10%, land 20%, capital 39%, robots 65%, AI 75%. Wages are the most broadly distributed; AI and robots are the most concentrated.

Wage income — workers vs. non-workers. The wage component above conceals an additional split. Each year, households are sorted by their wage skill s_h (derived from the wage CDF) and the top LFP $\times N$ are designated workers. Workers receive wage income proportional to their skill, $w_h^{\text{worker}} = (s_h/\bar{S}_W) \cdot \text{wL}_{\text{net}}$, where $\bar{S}_W$ is the population-weighted total skill of workers and wL_{net} is total net wage income. Non-workers receive labour-equivalent income at a fraction ϕ_{nw} of the skill-matched worker rate (default $\phi_{nw} = 0.20$):

$$w_h^{\text{non-worker}} = \phi_{nw} \cdot s_h \cdot \frac{\text{wL}_{\text{net}}}{\bar{S}_W}.$$

This includes transfers. With $\phi_{nw} = 0.20$ the bottom-tail (p1) ends up around 1–3% of the median, matching US Census ACS individual-income data after accounting for transfers.

6.5 Taxation & Citizen Dividends

Two modes for funding citizen dividends:

- **Cap-and-trade-funded dividends** (Plan A default). The government revenue stream from the cost model (cap-and-trade permits + labor tax + corporate tax + sales tax) pools into a redistribution budget.
- **Naive flat taxes.** Alternative: user-set tax rates on capital income (τ_k), AI rental income (τ_{AI}), and robot rental income (τ_R). Simpler; lets you explore counterfactual tax structures without the full cost-model machinery.

Revenue is split into five configurable allocation shares of each region's tax base, each a 3-anchor time series (2032 / 2035 / 2040, linearly interpolated): domestic dividends for US, China, and Rest-of-World, plus foreign-aid contributions from the US and China. Defaults: 25% → 50% → 50% for all three domestic shares; 10% → 20% → 20% for the two foreign-aid shares. The per-person dividend amount feeds back into labor supply via β_u in the LFP equation (9).

7. Other model-wide limitations

- **Fixed TFP per region.** Each region's TFP is calibrated to 2025 GDP and held constant. Cross-country productivity differences persist unchanged. In reality, AI diffusion would likely narrow these gaps.
- **Technology trajectories mostly exogenous.** AI/robot quantity, efficiency, and capability are user-specified. The Jones feedback (AI → more R&D → faster cost decline) and the endogenous-investment mode (Default World) are the only endogenous channels.
- **Single composite good.** No new goods, quality improvements, or consumer surplus from variety. Much of AI's value may come through channels that never appear in GDP.
- **No trade between regions.** Each region is a closed economy. No cross-border capital flows, technology diffusion, or supply chain interdependence.
- **Regional breakdown is not very trustworthy.** The US, China, and World simulations each run independently with their own parameter files and their own base-year data, and we currently use global-average parameters as a proxy for Rest-of-World. Beyond the no-trade limitation above, the independent-simulation design means the three regional lines can diverge in ways that aren't internally consistent (e.g. the summed "World (sum)" line differs from the standalone global simulation). The regional decomposition is useful for getting a rough sense of how different parameter sets affect the trajectory, but should not be read as a careful multi-region model.

8. Solver Architecture & Implementation Notes

Time step and state. The model simulates 16 years (2025 to 2040) in 1-year increments. Each year has its own production-function solve plus a cost-model solve. State carried between years: capital stock K, AI stock, robot stock, cumulative production Q_{cum} for Wright, design indices D for Jones, and (in endogenous mode) the accumulated investment-funded stocks.

Within-year solve. Given K_t, $\text{AI}_{\text{cog},t}$, $R_{\text{phys},t}$, H_t, and the exogenous capability frontiers $f_{c,t}, f_{p,t}$, the solver:

1. Allocates humans across the auto and non-auto task buckets by bisection on h_{auto}/H until the marginal products equalize across buckets (5).
2. Builds $L_{\text{cog}}, L_{\text{phys}}$ from the task-based CES (3)–(4).
3. Combines them into L_{eff} via (2).
4. Computes Y from (1), calibrated via A at the base year.
5. Derives wages, AI/robot rentals, and the interest rate as marginal products (6) via the nested-CES chain rule.

Endogenous labor supply (LFP iteration). When LFP responds to wages and citizen dividends, the within-year solve is wrapped in a fixed-point iteration: guess H, solve for wages, recompute LFP from the new wages, update H, repeat until $\Delta H/H < 1\%$. Damping factor 0.5 to prevent oscillation. Typically converges in 5–8 iterations.

Between-year accumulation.

- **Capital:** $K_{t+1} = (1-\delta)K_t + s(r_t)Y_t - \text{(R\&D cost)}$; R&D cost is subtracted from the investable savings pool.

- **Wright/Jones:** Q_{cum} grows with production; design index D declines at the Jones rate (14)–(15).
- **Endogenous mode only:** capital, AI, and robot stocks all accumulate from the investable savings pool via ROI-weighted shares (22) and stock-update equations (23)–(25). See §6.1.

Cost model integration. Each year after the growth solve, the cost model runs with the current AI/robot quantities, Q_{cum}, R&D labor allocation ($L_{\text{eff}} \times$ user fraction), and markup trajectories. It produces MC, permit prices (if cap-and-trade is active), buyer prices, surplus decomposition, and tax revenues. Tax revenue feeds back into next year's citizen-dividend pool with a one-year lag.

Base-year calibration.

- TFP A calibrated once so Y_{2025} matches the region's base-year GDP given K_{2025} and $L_{\text{eff},2025}$.
- ν (goods/services) calibrated from base-year services spending share a_S (26).
- Land demand/supply scale coefficients calibrated from base-year rents $v_{i,0}$ and acres $M_{i,0}$.
- r_{base} for savings set to the base-year interest rate.

Regions run independently. The model runs a full simulation separately for US, China, and Rest-of-World when the multi-region toggle is on. No cross-region trade, capital flow, or technology diffusion. Each region has its own params file (α, θ, savings rate, base-year macro, ownership distributions, tax rates) and a separate cost-side state.

9. All Model Equations

Growth Model

$$Y_t = A \cdot \text{CES}(K_t, L_{\text{eff},t}; \alpha, \sigma_Y) \quad \textit{output = TFP} \times \textit{CES of capital and effective labor} \tag{1}$$

$$L_{\text{eff},t} = \text{CES}(L_{\text{cog},t}, L_{\text{phys},t}; \theta, \sigma_L) \quad \textit{effective labor = CES of cognitive and physical} \tag{2}$$

$$L_{\text{cog},t} = \text{CES}_{\text{task}}(h_{\text{auto}} + A_{\text{eff}}, h_{\text{non}}; f_c, \sigma_c) \quad \textit{cognitive task-CES: humans + AI in auto, humans only in non-auto} \tag{3}$$

$$L_{\text{phys},t} = \text{CES}_{\text{task}}(h_{\text{auto}} + R_p, h_{\text{non}}; f_p, \sigma_p) \quad \textit{same structure for physical, with robots} \tag{4}$$

$$K_{t+1} = s(r_t) \cdot Y_t + (1-\delta)K_t \quad \textit{Solow–Swan capital law of motion} \tag{5}$$

$$s(r) = s_{\text{base}} \cdot \left(\frac{1+r}{1+r_{\text{base}}}\right)^{\eta_s} \quad \textit{Boskin (1978) style endog. savings, } \eta_s = 0.8 \tag{6}$$

$$h_{\text{auto}} + h_{\text{non}} = H \;\; \text{s.t. } \text{MP}_{\text{auto}} = \text{MP}_{\text{non}} \quad \textit{humans allocate to equalize marginal product} \tag{7}$$

$$r = \alpha Y/K, \quad w_c = \frac{\partial Y}{\partial h_{\text{non},c}}, \quad q_c = \frac{\partial Y}{\partial A_{\text{eff}}} \quad \textit{factor prices = marginal products} \tag{8}$$

Labor supply (endogenous LFP)

$$H_t = \text{WAP}_t \cdot \text{LFP}_t \quad \textit{total human labor = working-age pop} \times \textit{participation rate} \tag{9}$$

$$\text{LFP}_t = 1/(1+\exp(-z_t)) \quad \textit{logistic} \tag{10}$$

$$z_t = z_{\text{baseline}} + \beta_w \ln(\bar{w}_t/\bar{w}_0) - \beta_u(\text{CD}_t/\bar{w}_0) \quad \textit{wages push LFP up, citizen dividends push down} \tag{11}$$

$$z_{\text{baseline}} = \ln(\text{LFP}_{\text{target}}/(1-\text{LFP}_{\text{target}})) \quad \textit{pins LFP = target at base-year wage, zero citizen dividends} \tag{12}$$

Cost Model

$$\mathrm{MC}_{i,t} = \mathrm{MC}_{i,0} \cdot \mathrm{Wright}_{i,t} \cdot \mathrm{Jones}_{i,t} \quad \textit{MC = base-year MC scaled by manuf. learning and design} \tag{13}$$

$$\mathrm{Wright}_{i,t} = (Q_{\mathrm{cum},t}/Q_{\mathrm{cum},0})^{-b} \quad \textit{Wright's Law: each doubling of } Q_{\mathrm{cum}} \textit{ cuts cost by } (1-2^{-b}) \tag{14}$$

$$\dot{D}_{i,t} = r_{\mathrm{base}} \cdot (L_{\mathrm{design},t}/L_{\mathrm{design},0})^{\lambda} \cdot D_{i,t}^{1-\phi} \quad \textit{Jones: } \lambda \textit{ stepping-on-toes, } \phi \textit{ fishing-out} \tag{15}$$

$$L_{\mathrm{design},t} = L_{\mathrm{eff},t} \cdot \mathrm{frac}_{\mathrm{R\&D},t} \quad \textit{R\&D workforce} = L_{\mathrm{eff}} \times \textit{exogenous allocation \%} \tag{16}$$

$$P_{i,t} = \mathrm{MC}_{i,t} \cdot \mu_{\mathrm{mfg},t} \cdot \mu_{\mathrm{design},t} \quad \textit{price = cost} \times \textit{user-specified markups} \tag{17}$$

$$\mathrm{Surplus}_{i,t} = (q_{i,t} - P_{i,t}) \cdot Q_{i,t} \quad \textit{buyer surplus = value − price, times quantity} \tag{18}$$

Energy and Climate

$$E_t = E_{2025} \cdot (Y_t/Y_{2025})^{\varepsilon} \quad \textit{total primary energy: GDP-anchored with constant decoupling elasticity } \varepsilon \textit{ (default 0.4)} \tag{19}$$

$$E_{j,t}^{\mathrm{tech}} = S_{j,t} \cdot W_{j,t}^{\mathrm{avg}} \quad \textit{bottom-up tech energy: stock} \times \textit{vintage-averaged power per unit } (j \in \{\textit{AI compute, robots}\}) \tag{20}$$

$$E_t^{\mathrm{other}} = \max\Big(0,\ E_t - E_{\mathrm{AI},t}^{\mathrm{tech}} - E_{\mathrm{robot},t}^{\mathrm{tech}}\Big) \quad \textit{residual “Other” = total minus bottom-up AI and robot energy} \tag{21}$$

$$\mathrm{LCOE}_{i,t} = \mathrm{CapEx}_{i,t} \cdot \mathrm{CRF}_i \cdot 114.16 + \mathrm{OpCost}_{i,t} + \mathrm{ext}_{i,t} \quad \textit{levelized cost per generation source } (i \in \textit{fossil/nuclear/hydro/solar} \tag{22}$$

Extensions

$$\nu = a_S P_{G,0}^{1-\eta}/[a_S P_{G,0}^{1-\eta} + (1-a_S)P_{S,0}^{1-\eta}] \quad \textit{calibrate } \nu \textit{ so base-year service share} = a_S \tag{23}$$

$$P_S \propto (p_{\mathrm{eff},S}/p_{\mathrm{eff,agg}})^{1-\alpha}, \quad P_G \propto (p_{\mathrm{eff},G}/p_{\mathrm{eff,agg}})^{1-\alpha} \quad \textit{sector prices from relative labor costs} \tag{24}$$

$$S = \nu P_S^{-\eta} Y, \quad G = (1-\nu) P_G^{-\eta} Y \quad \textit{CES demand: cheaper sector gets more real quantity} \tag{25}$$

$$X_t = (1 - s(r_t)) \cdot Y_t \quad \textit{consumer expenditure = output − savings (used in land demand below)} \tag{26}$$

$$v_i(t)/v_{i,0} = (X_t/X_0)^{\beta_i/(\eta_i^D + \eta_i^S)} \quad \textit{equilibrium land rent, base-year-relative} \tag{27}$$

$$M_i(t)/M_{i,0} = (v_i(t)/v_{i,0})^{\eta_i^S} \quad \textit{acres from supply curve at new rent} \tag{28}$$

$$R_i(t) = v_i(t) \cdot M_i(t) \quad \textit{total land rent = rent/acre} \times \textit{acres} \tag{29}$$

$$\mathrm{share}_j(p) = \mathrm{CDF}_j(p+dp) - \mathrm{CDF}_j(p) \quad \textit{ownership share per percentile bin, from CDF} \tag{30}$$

$$I_j(p) = m_j(p) \cdot \mathrm{share}_j(p) \cdot \mathrm{TotalIncome}_j \quad \textit{per-person income from factor } j \textit{ at percentile } p \tag{31}$$

$$I_{\mathrm{total}}(p) = \sum_j I_j(p) + \mathrm{CD}(p) + \mathrm{FA}(p) \quad \textit{factors + domestic citizen dividend + foreign aid received (FA = 0 in donor regions)} \tag{32}$$

$$W_{t+1}(p) = W_t(p) + s(r_t) \cdot I_{\mathrm{total},t}(p) \quad \textit{net-worth accumulation at economy-wide savings rate} \tag{33}$$

$$\mathrm{Gini} = 1 - 2\int_0^1 L(p)\,dp \quad \textit{standard Gini from Lorenz curve } L(p) \tag{34}$$

CES conventions.

$$\mathrm{CES}(x, y; \mu, \sigma) = \Big[\mu^{1/\sigma} x^{(\sigma-1)/\sigma} + (1-\mu)^{1/\sigma} y^{(\sigma-1)/\sigma}\Big]^{\sigma/(\sigma-1)}$$

$$\mathrm{CES}_{\mathrm{task}}(L_{\mathrm{auto}}, L_{\mathrm{non}}; f, \sigma) = \Big[f \cdot L_{\mathrm{auto}}^{(\sigma-1)/\sigma} + (1-f) \cdot L_{\mathrm{non}}^{(\sigma-1)/\sigma}\Big]^{\sigma/(\sigma-1)}$$

$\sigma = 1$: Cobb-Douglas. $\sigma < 1$: complements. $\sigma > 1$: substitutes.

A. Energy Cost Parameters (Claude research)

This appendix documents the source-based justification for the energy generation cost parameters used in §5. Each of the six generation sources (fossil, nuclear, hydro, solar, wind, other) has twelve parameters: a 2025 anchor for CapEx and OpCost plus Jones-curve dynamics ($r_{\text{base}}, \lambda, \phi$) for each side. Values below were assembled by Claude via web research in 2025; treat them as a defensible starting point, not a calibrated truth.

Conventions. $/W-delivered = installed cost $/kW ÷ (capacity factor × 1000). "Firmed" solar/wind values include a 4-hour battery storage adder. Build-mix weights reflect 2024 global capacity additions by region (China dominant for most renewables and nuclear). Capacity factors used: PV 17–25%, onshore wind 28–40%, CCGT 50–80%, coal 60%, nuclear 85–92%, hydro 35–50%, geothermal 85%, offshore wind 42%.

Starting CapEx ($/W-delivered, 2025 anchor)

Tech	Estimate	Math (build-mix weighted)	Sources
Fossil	$1.3	0.70×$1.40 (CN CCGT $700/kW ÷ 50% CF) + 0.20×$2.18 (US CCGT $1.2k/kW ÷ 55% CF) + 0.10×$2.00 (RoW) = $1.62/W. Rounded down to $1.3 leaning on pure-CCGT Asia-weighted mix.	Gas Turbine World; EIA AEO2025; IEA Renewables 2024
Nuclear	$4.5	0.55×$3.29 (CN Hualong $2.8k/kW ÷ 85% CF) + 0.20×$5.40 (RU VVER $5k/kW ÷ 88%) + 0.10×$5.00 (KR APR $4.5k/kW ÷ 90%) + 0.15×$9.00 (West blended) = $4.74/W. Asian-only would be $3.3/W; Vogtle final cost $10,784/kWe ÷ 92% ≈ $11.7/W. We use $4.5 as the current-build blend.	IAEA PRIS; OECD-NEA 2020; MIT CANES Next-AP1000; Vogtle final cost
Hydro	$5.0	0.65×$4.50 (CN $1.8k/kW ÷ 40% CF) + 0.15×$11.40 (OECD $4k/kW ÷ 35%) + 0.20×$6.67 (RoW $3k/kW ÷ 45%) = $5.97/W; rounded to $5.0.	IRENA RPGC 2024 (global avg $3,053/kW); Three Gorges $1,420/kW historical
Solar firmed	$4.0	0.55×$3.48 (CN $591/kW ÷ 17% CF) + 0.10×$2.63 (IN $525/kW ÷ 20%) + 0.10×$5.99 (EU $779/kW ÷ 13%) + 0.15×$5.72 (US $1.43k/kW ÷ 25%) + 0.10×$4.20 (RoW) = $4.05/W unfirmed. Firming adder: 4-hr BESS at global $165/kWh (CN $85/kWh) × 1.5 overbuild → +$1.5–2.0/W. $4.0 sits between the China-weighted firmed ($3.5) and full-global firmed ($5.5).	IRENA RPGC 2024; NREL ATB 2024 Util-Scale PV; BNEF Battery Storage Cost Survey 2024
Wind firmed	$3.5	0.70×$3.04 (CN $850/kW ÷ 28% CF) + 0.20×$3.50 (OECD $1.4k/kW ÷ 40%) + 0.10×$4.00 (RoW) = $3.23/W unfirmed. Firming adder +$0.5–1.0/W (wind needs less storage than solar).	IRENA RPGC 2024 (global $1,041/kW); NREL ATB 2024 Land-Based Wind; LBNL Wind Market Report 2024
Other	$5.0	0.55×$6.79 (offshore wind $2,852/kW ÷ 42% CF) + 0.20×$5.29 (Geo $4.5k/kW ÷ 85%) + 0.10×$11.10 (CSP $5k/kW ÷ 45%) + 0.15×$4.17 (Biomass $2.5k/kW ÷ 60%) = $6.53/W; rounded down to $5.0 acknowledging mix uncertainty.	IRENA RPGC 2024 (offshore, geo); NREL ATB 2024 Geothermal

Starting OpCost ($/MWh delivered, 2025 anchor)

Tech	Estimate	Math	Sources
Fossil	\$40	US CCGT 2024: fuel \$2.21/MMBtu × 6.5 Btu/kWh heat rate = \$14.4/MWh + O&M \$5/MWh = \$19/MWh. At more typical \$4/MMBtu: \$31/MWh. World blend: 0.50×\$30 (CN coal+CCGT) + 0.20×\$19 (US gas) + 0.15×\$80 (EU/Asia LNG) + 0.15×\$35 (India coal) ≈ \$36/MWh; we use \$40 to be slightly conservative on persistent Asian-LNG pricing.	EIA Henry Hub 2024; Lazard LCOE+ v17–v18; GNESTE techno-economic db
Nuclear	\$25	0.25×\$17 (CN) + 0.10×\$18 (KR) + 0.10×\$25 (FR) + 0.30×\$31 (US; NEI 2024 fleet non-capital \$20.48/MWh + provisions) + 0.25×\$25 (RoW) = \$24.1/MWh.	NEI Nuclear Costs in Context 2025; WNA Economics of Nuclear Power; Lazard LCOE+ v17
Hydro	\$10	IRENA O&M ∼1–4% of capex/yr. At \$3k/kW × 2.5% = \$75/kW-yr / (0.40 CF × 8760) = \$21/MWh, but most large hydro globally runs lower. CN Three-Gorges-class \$3–5/MWh. Weighted \$8–12/MWh.	IRENA RPGC 2024 (hydro O&M); NREL ATB 2024 Hydropower
Solar	\$8	NREL ATB 2024: fixed O&M \$22/kW-yr / (0.25 CF × 8760) = \$10/MWh (US). LBNL utility-scale levelized OpEx \$17/kW-yr ≈ \$8.8/MWh. CN \$4–5/MWh (lower labor). Blended \$6–9/MWh.	NREL ATB 2024 Util-Scale PV; LBNL Utility-Scale Solar 2024
Wind	\$12	LBNL 2024: ∼\$45/kW-yr / (0.35 CF × 8760) = \$14.7/MWh (US). Wood Mackenzie global ∼\$40/kW-yr; CN \$15–20/kW-yr / 30% CF = \$7.6/MWh. Blended \$10–13/MWh.	LBNL Land-Based Wind 2024; Wood Mackenzie Onshore Wind O&M H2 2024
Other	\$30	Offshore wind \$90/kW-yr / (0.42 CF × 8760) = \$24.5/MWh. Geo \$130/kW-yr / (0.85 × 8760) = \$17.4/MWh. CSP \$60/kW-yr / (0.45 × 8760) = \$15.2/MWh. Biomass fuel+O&M \$40–50/MWh (feedstock-dominated). Mix-weighted \$25–35.	IRENA RPGC 2024 (offshore, geo, bioenergy); NREL ATB 2024 Geothermal

Base improvement rate r_{base} (annual rate at $t = 0$)

The Jones-curve rate at year zero. Historical Wright's-law learning rates (% per doubling of cumulative capacity), divided by the current cumulative-capacity doubling time, give an empirically-anchored historical-rate estimate. Values below are forward-looking with mild AGI tailwinds.

Tech	CapEx	OpCost	Reasoning	Sources
Fossil	0.015	0.01	Rubin et al. coal Wright 5.6% (range −11 to 12%); CCGT −11 to 34% (central ∼0–5%/doubling). With slow fossil capacity growth: implied historical 0.001–0.005/yr. CapEx has risen since 2020 (gas-turbine OEM tightness). 0.015/yr is a modest AGI-tailwind anchor.	Rubin et al. 2015 review; GridLab Gas Turbine Costs Report 2025
Nuclear	0.07	0.07	*Bullish fusion-aware.* Lovering et al.: US negative learning, KR +20%/doubling. NREL ATB assumes 8% NOAK; DOE Adv Nuclear Liftoff targets 10–15%. We pick *above* the central case to capture option value from fusion (deuterium fuel + compact tokamaks like SPARC, \$0.5/MWh vs \$6 fission fuel) and AGI-driven factory production breaking the negative-learning trap.	Lovering et al. 2016; Eash-Gates et al. 2020 (Joule); DOE Adv Nuclear Liftoff 2024; CFS SPARC roadmap
Hydro	0.01	0.005	Junginger & Louwen 2020: 0.48–1.4%/doubling. Implied <0.05%/yr from cumulative growth alone. CapEx-side gets help from general productivity; OpCost stays slow (labor inflation often outpaces productivity). Site exhaustion (best sites taken first) dominates further reductions.	Junginger & Louwen 2020

Tech	CapEx	OpCost	Reasoning	Sources
Solar	0.15	0.10	Long-run Wright 20–24%/doubling; recent 2014–20 epoch 45%/doubling. Capacity doubled every ∼3 years → 7–17%/yr observed annual. 0.15/yr is central as doubling stretches by 2030. OpCost falls faster than baseline as automation hits installation/maintenance.	Our World in Data; Way et al. 2022; ScienceDirect 2022 levelized-cost learning
Wind	0.08	0.06	Rubin et al. onshore wind 12%/doubling (range 7–22%); 2010–20 epoch 40%/doubling. Implied 3–12%/yr. 0.08/yr is central for forward-looking.	Rubin et al. 2015; ScienceDirect 2022
Other	0.04	0.03	Offshore wind 5–10%/doubling; geo/biomass mature. ∼1–2%/yr from learning + general productivity. OpCost reduction limited by feedstock cost (biomass).	Junginger 2023 (offshore wind learning); IRENA RPGC 2024

Stepping-on-toes λ (researcher elasticity)

10× more R&D researchers gives 10^{λ} faster progress. Bloom et al. 2020 identify λ jointly with ϕ; for semiconductors $\lambda/(1-\phi) \approx 5$. Energy techs generally have lower λ than chips because they're more constrained by physical scale-up than by researcher headcount.

Tech	Estimate	Reasoning
Fossil	0.08	Mature; few OEMs (GE/Siemens/MHI) dominate gas-turbine R&D, limiting researcher-headcount elasticity. Bloom et al. mature-industry anchor.
Nuclear	0.25	*Bullish fusion-aware.* Large remaining design space (SMR, advanced reactors, factory production, fusion) plus AGI-accelerated R&D suggests stronger researcher-responsiveness than the standard 0.20 for moderate-tech. DOE Adv Nuclear Liftoff implies meaningful λ under modular regime.
Hydro	0.05	Civil-works dominated; very little R&D leverage. Junginger & Louwen 2020 imply very low researcher elasticity.
Solar	0.15	Chinese PV R&D scaled ∼50× from 2010–2020 with visibly proportional cost effects across modules, BOS, and inverters. Anchored on that empirical responsiveness.
Wind	0.10	Less researcher-responsive than solar (mechanical scale limits hit, e.g. rotor diameter → material).
Other	0.08	Mix includes early-stage techs (CSP, advanced geothermal, EGS) with R&D headroom.

Fishing-out ϕ (idea-depletion exponent)

$\phi < 1$ means new ideas get harder as design stock grows. Steady-state $\lambda/(1-\phi)$ is "OOMs of progress per OOM of researcher growth"; Bloom et al. 2020 find ~ 5 for semiconductors and ~ 3 for aggregate US growth. Energy techs should sit below those values — physical/scaling constraints dominate.

Tech	Estimate	$\lambda/(1-\phi)$	Reasoning
Fossil	0.85	0.53	Mature; fast fishing-out is appropriate. Anchored to Bloom et al. mature-industry tail.
Nuclear	0.90	2.50	*Bullish fusion-aware.* Multi-S-curve structure (fission → SMR → fusion) keeps slope shallow; each transition effectively resets the curve. Implied steady-state $\lambda/(1-\phi) = 2.5$ approaches Bloom et al. aggregate-US value (∼3).
Hydro	0.30	0.07	Even with strong site-exhaustion, general-productivity pass-through is steadier than 0.20 would imply.

Tech	Estimate	$\lambda/(1-\phi)$	Reasoning
Solar	0.70	0.50	Remaining design space (perovskites, tandems, bifacial, BOS) supports slightly slower fishing-out.
Wind	0.70	0.33	Captures ongoing turbine-size + offshore-specific design space.
Other	0.90	0.80	Mix has techs in different curve stages; slightly more fishing-out than full-mature assumption.

Source list

- IRENA — Renewable Power Generation Costs in 2024 (Summary): global TIC and LCOE for renewables.
- Lazard LCOE+ v17 (June 2024) and v18 (June 2025): Western LCOE benchmarks.
- NREL ATB 2024 (Utility-Scale PV, Land-Based Wind, Nuclear, Geothermal, Fossil): US fleet capex + O&M.
- BNEF Battery Storage Cost Survey 2024: BESS \$165/kWh global, \$85/kWh China (−40% YoY).
- NEI Nuclear Costs in Context 2025: US fleet 2024 non-capital ops \$20.48/MWh.
- EIA Today in Energy: 2024 Henry Hub \$2.21/MMBtu.
- EIA AEO2025 Capital Cost Study: US thermal plant capex.
- Wood Mackenzie Onshore Wind O&M Economics H2 2024: wind O&M by region.
- LBNL Utility-Scale Solar 2024 Edition: US PV fleet O&M.
- LBNL Land-Based Wind Market Report 2024.
- Vogtle final cost: US AP1000 \$10,784/kWe overnight.
- NextBigFuture: APR-1400 / Hualong One cost comparison ∼\$2–3k/kWe.
- Lovering et al. 2016: historical nuclear construction costs.
- Eash-Gates et al. 2020 (Joule): sources of cost overrun in nuclear construction.
- DOE Pathways to Commercial Liftoff: Advanced Nuclear (2024).
- MIT CANES: Overnight Capital Cost of the Next AP1000.
- Bloom, Jones, Van Reenen, Webb 2020 (AER): λ, ϕ anchors.
- Rubin et al. 2015: review of learning rates for electricity supply technologies.
- Way et al. 2022 / Volts: learning curves and cheap clean energy.
- Levelized cost-based learning analysis (ScienceDirect 2022): solar/wind recent-epoch Wright rates.
- Junginger & Louwen 2020: experience curves for renewable O&M.
- Junginger 2023: offshore-wind learning rates.
- Our World in Data: Wright's Law / learning curves.
- World Nuclear Association: Economics of Nuclear Power.
- GridLab Gas Turbine Costs Report 2025.

- Gas Turbine World capital costs.

References

- Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, 110(4), 1104–1144.
- Boskin, M. J. (1978). Taxation, Saving, and the Rate of Interest. *Journal of Political Economy*, 86(2), S3–S27.
- Chow, T., Halperin, B., & Mazlish, J. (2023). Transformative AI, existential risk, and real interest rates. Working paper.
- Comin, D., Lashkari, D., & Mestieri, M. (2021). Structural Change with Long-run Income and Price Effects. *Econometrica*, 89(1), 311–374.
- Herrendorf, B., Rogerson, R., & Valentinyi, Á. (2013). Two Perspectives on Preferences and Structural Transformation. *AEJ: Macroeconomics*.
- Jones, C. I. (1995). R&D-Based Models of Economic Growth. *Journal of Political Economy*, 103(4), 759–784.
- Ngai, L. R., & Pissarides, C. A. (2007). Structural Change in a Multisector Model of Growth. *American Economic Review*, 97(1), 429–443.
- Sequeira, T., & Neves, P. (2020). Stepping on Toes in the Production of Knowledge. *Oxford Economic Papers*, 72(4), 1036–1061.
- Wright, T. P. (1936). Factors Affecting the Cost of Airplanes. *Journal of the Aeronautical Sciences*, 3(4), 122–128.