UNIVERSITÉ DU QUÉBEC EN OUTAOUAIS

DÉPARTEMENT DES SCIENCES ADMINISTRATIVES

WORKING PAPER NO. 24

From Labor to Capital

Artificial General Intelligence, General Equilibrium, and Asset Valuation

Simon-Pierre Boucher

Département des sciences administratives
Université du Québec en Outaouais

simon-pierre.boucher@uqo.ca

Gatineau – Pavillon Alexandre-Taché
283, boulevard Alexandre-Taché
Gatineau, Québec, Canada J9A 1L8

October 2026

Version 1.0

Abstract

Restrepo (2025) shows that once artificial general intelligence lets compute perform every task, output becomes linear in compute, wages converge to the value of the compute a worker saves, and the labor share goes to zero. We extend his task economy by treating compute as a reproducible, priced, depreciating capital good that needs energy and works alongside structures and land, and by adding households that own the claims to income. The labor-share result survives any constant-returns production structure: it needs only within-task substitutability. The wage-level result does not: when physical complements are gross complements and compute becomes cheap, the price of compute and wages go to zero in level and the income that labor loses goes to the owners of the scarcest complement; with a floor on the compute price, the wage equals the user cost of the replacement compute. Income shares and asset values are different objects: compute is capitalized at the duration of chips, so capital's share can rise while the wealth-to-income ratio falls, and a growth boom raises the real rate and compresses valuation ratios when the elasticity of intertemporal substitution is below one. In a calibrated model solved along perfect-foresight paths to 2300 under conditional scenarios, the labor share falls to 0.283 by 2050, the real return peaks at 40.5%, the wealth-to-income ratio falls to 2.7 before land captures 0.605 of income and raises it to 31.0 by 2100, and the market value of compute never exceeds one fifth of GDP. A compute tax raises wages one-for-one with the tax-inclusive user cost; a land tax redistributes without distortion; a citizen fund keeps workers' consumption share bounded away from zero.

Keywords: artificial general intelligence, task-based automation, labor share, compute capital, asset pricing, wealth distribution, land rents, energy constraints

JEL Classification: E22, E24, E25, G12, J23, J24, O33, O41

1. Introduction

Restrepo (2025) describes an economy in which artificial general intelligence has arrived: every type of work that people do can also be done by computers, at a cost in compute that is finite and known, and the stock of compute grows without bound. His conclusions are stark and, within his premises, exact. All work that is essential for growth is eventually done by compute; output becomes linear in compute; wages converge to the value of the compute that a worker saves; and the share of labor in income converges to zero while labor's income does not fall. "The economy keeps growing; we stay in place." The chapter is deliberately silent about the other side of the national accounts. Compute is an endowment that nobody buys, depreciates or finances; nothing is said about what the income that leaves labor is worth as an asset, who owns the claims to it, or how that ownership evolves. This paper is about that other side.

We ask three questions. First, does the income that leaves labor go to the owners of compute, as the source concludes, or to the owners of what compute cannot reproduce — structures, land, energy sites? Second, what is that income worth: does "all income accrues to compute" mean that the value of compute capital, or of wealth in general, explodes? Third, how does the ownership of wealth evolve across households, and what do the policies that the source mentions in passing — universal dividends, treating compute "as a public or semi-public resource, akin to land" — do in general equilibrium?

To answer them we keep the source's task block intact and change one thing in production: compute is a reproducible capital good with a price that falls over time, a depreciation rate, an energy requirement, and complements — structures, land and energy — that it cannot replace and that enter output with an elasticity of substitution that may be above or below one. Everything on the household side is new: a two-class economy whose capitalists' Euler equation prices every claim, and a heterogeneous-agent economy in which households with different skills, discount factors and productivity histories hold the claims through a fund. We solve the model in three layers of increasing size: the exact replication of the source (Section 3); a limit economy with one complement in which eight propositions are proved (Section 7); and a quantitative economy with four capital stocks, energy and housing, calibrated to the United States in 2026 and solved along perfect-foresight paths to 2300 under a catalogue of conditional scenarios (Sections 8–10).

Findings. The labor-share result of the source is more robust than its author claims, and its wage result is less. The labor share goes to zero under any constant-returns production structure, with or without complements, with compute as an endowment or as a priced capital good, because the only thing it needs is that labor and compute be perfect substitutes within tasks (Proposition 1). The level of wages, by contrast, depends on compute remaining scarce. If the task composite and the physical

complements are gross complements ($\sigma < 1$) and compute becomes cheap, the price of compute goes to zero and so do wages — in level, not only as a share — and the income that labor loses goes to the owners of the scarcest complement (Propositions 2–4). With a floor on the compute price, wages converge to the user cost of the compute that replaces the worker, $w_s = (r+\delta_Q)\bar{p}\,\widetilde{\mathrm{CEU}}(s)$: "the rental price of your replacement" (Proposition 3). An energy bottleneck keeps the labor share positive (Proposition 5).

On valuation, the income share and the market value of a claim are different objects. On a balanced path the wealth-to-income ratio is $\sum_k s_k/[(1+r)(1+\pi_k)-(1-\delta_k)]+s_T/(r-g)$: a claim that depreciates at 25% a year and cheapens at 15% a year — compute — is capitalized at a price–earnings ratio of about two, against eight for structures and fifty for land (Proposition 6). Capital's share can rise while wealth-to-income falls. And a growth boom raises the real rate through the marginal investor's Euler equation, so that valuation ratios of long-duration claims *fall* during the boom when the elasticity of intertemporal substitution is below one (Proposition 7).

Quantitatively, in the baseline scenario — capability improving at 10% a year, compute prices falling at 15% a year to one hundredth of today's level, energy efficiency gains with a floor — the routine task is automated by 2029 and the expert task by 2057. The labor share falls from 0.56 to 0.283 in 2050 and below one in a thousand after 2075. The lower skill's wage falls in steps as the source's compute-binding transition predicts; the higher skill's wage rises to 3.51 times its 2026 level while its expert task is the last one that compute cannot do cheaply, then collapses — the source's technology-binding transition. Output rises 3.7-fold by 2050 and 233.0-fold by 2100, after which growth returns to its trend because land binds. The real return on wealth rises to 40.5% in 2060 and the wealth-to-income ratio *falls* from 4.1 to 2.7 in 2042 while the capital share rises; it then climbs to 31.0 by 2100 as land, which cannot be produced, captures 0.605 of income. The market value of compute never exceeds one fifth of GDP. The decomposition of 2026 valuations into cash-flow and discount-rate news shows the two channels to be of the same order, with opposite signs. In the heterogeneous economy the wealth Gini rises from 0.74 to 0.73 by 2100; a citizen fund holding one fifth of all claims keeps workers' consumption share bounded away from zero; a compute tax raises wages one-for-one with the tax-inclusive user cost at a cost in output; a land tax redistributes without changing allocations; an AI monopoly protects wages the way a compute tax does, with the difference that the rents go to the monopolist.

What this paper is not. It is not a forecast. Every scenario input — the rate of capability progress, the compute price floor, the energy-efficiency floor, the terminal compute requirement of each task — is documented as such, and the level of wages after 2060 is proportional to the product of two of them. It is not a model of the arrival of AGI under uncertainty: the paths are perfect-foresight equilibria, there are no risk premia, and the stochastic-arrival exercise of Section 10 is a labelled

approximation. And it does not model endogenous innovation, the creation of new human tasks beyond one reinstatement scenario, or nonhomothetic preferences; Section 12 lists what those would change.

Organization. Section 2 places the paper in the literature. Section 3 reproduces and audits the source. Sections 4–6 present the quantitative model, its households and its asset prices. Section 7 states and discusses the propositions, with proofs in Appendix A. Section 8 calibrates; Section 9 presents the baseline; Section 10 the scenarios, the heterogeneous economy, policy and welfare; Section 11 the empirical discipline and a falsifiable short-run design; Section 12 concludes. The replication package reproduces every number with one command.

2. Related literature and what is different here

Task models of automation. The work equation $X = L + Q/\alpha$ descends from the task framework of Autor et al. (2003), Acemoglu and Autor (2011) and Zeira (1998), in which technology replaces labor task by task. Acemoglu and Restrepo (2018) embed the task set in a growth model with endogenous creation of new tasks, Acemoglu and Restrepo (2019) separate displacement from reinstatement, and Acemoglu and Restrepo (2022) estimate that task displacement accounts for most of the rise in US wage inequality since 1980; Acemoglu and Restrepo (2020) and Hémous and Olsen (2022) provide the industrial-robot and automation–inequality counterparts. Restrepo (2025) takes the framework to its limit: when every task can be replicated by compute, growth is the growth of compute, wages are the value of the compute they save, and the labor share goes to zero. We keep his primitives and change one thing in the production block — compute is a reproducible, priced, depreciating capital good that needs energy and works alongside structures and land — and we add everything on the household side.

The economics of transformative AI. Aghion et al. (2019) show that Baumol-type bottlenecks determine whether automation produces a growth explosion or a sequence of level effects; our regime map (Table 3) is the asset-pricing counterpart of that observation. Trammell and Korinek (2023) survey growth under transformative AI; Korinek and Suh (2024) characterize the transition to AGI in task models with a distribution of task complexity and show that wages can collapse once compute crosses a threshold, which is the mechanism of Propositions 2–4 in a different parametrization. Korinek and Stiglitz (2019) and Korinek and Juelfs (2022) discuss distribution and the "future of work"; Acemoglu (2025) argues from Hulten's theorem that the near-term TFP gains from current AI are modest, a discipline we respect by starting the 2026 base point at the observed compute share. Nordhaus (2021) tests and rejects the singularity hypothesis on historical data; our

"Moore's law never ends" scenario is where the model itself produces one, and our baseline does not. Moll et al. (2022) is the closest antecedent on distribution: automation raises returns to wealth and so wealth inequality in an Aiyagari economy; we obtain their mechanism as one of three cells of Table 3 and add the valuation side. Hanson (1998) and Bostrom (2014) are the long-run speculative antecedents the source cites.

Growth, bottlenecks and networks. Jones (1995), Kortum (1997) and Segerstrom (1998) provide the semi-endogenous growth structure of the source's science block; Bloom et al. (2020) and Jones (2022) the evidence that ideas are getting harder to find. Baumol (1967) and Comin et al. (2021) are the references for unbalanced growth with gross complements, the force behind our scarce-complement regimes. Hulten (1978), Acemoglu et al. (2012), Jones (2011) and Baqaee and Farhi (2019, 2020) supply the production-network machinery; we use it only as the accounting discipline that gross output of compute and energy is intermediate, not final, output.

Heterogeneous households and incomplete markets. The household block is a Bewley (1986)–Huggett (1993)–Aiyagari (1994) economy with discount-factor heterogeneity as in Krusell and Smith (1998) and Carroll et al. (2017), solved by the endogenous grid method of Carroll (2006). Kaplan et al. (2018) and Auclert et al. (2021) are the references for transition dynamics with heterogeneous agents; our relaxation on the interest-rate path is the simplest member of that family. Krusell and Smith (1997), Constantinides and Duffie (1996) and Heaton and Lucas (1996) study asset prices with heterogeneous agents; because the two-class model nests the heterogeneous economy, the marginal investor in our pricing equations is always an unconstrained saver. On the dynamics of top wealth, Benhabib et al. (2011), Gabaix et al. (2016), Hubmer et al. (2021) and Fagereng et al. (2020) document the role of return heterogeneity, which we do not model and which is why the model under-predicts the top-1% share. Mian et al. (2021) is the reference for the demand-side consequences of wealth concentration.

Asset pricing and valuation. The stochastic discount factor is the one of Lucas (1978); we work under perfect foresight for the transition paths and discuss Mehra and Prescott (1985), Epstein and Zin (1989) and Bansal and Yaron (2004) when we interpret the discount-rate channel. The decomposition of valuation changes into cash-flow and discount-rate news follows Campbell and Shiller (1988a,b) and Campbell (1991), computed exactly along the deterministic path. Investment valuation uses Tobin (1969), Hayashi (1982), Abel and Blanchard (1983) and Cochrane (1991). Greenwald et al. (2025) attribute most of the post-1989 rise in US equity values to a reallocation of rewards from labor to shareholders; Farhi and Gourio (2018) decompose macro-finance trends into market power, intangibles and risk premia; Rachel and Summers (2019) and Eggertsson et al.

(2019) on the secular decline in real rates. Our Proposition 6 is a warning against reading the first of these results forward: the income that labor loses to compute is capitalized at the duration of chips, not of equities.

Factor shares, markups and intangibles. Karabarbounis and Neiman (2014), Elsby et al. (2013), Rognlie (2015), Barkai (2020), Autor et al. (2020), Kehrig and Vincent (2021) and Bergholt et al. (2022) document and dispute the decline of the labor share; the compute-equivalent wage gives a structural reading of a further decline. De Loecker et al. (2020), Crouzet and Eberly (2019) and Corrado et al. (2009) motivate the markup and intangible scenarios. Oberfield and Raval (2021) and Chirinko (2008) inform the elasticities σ_N and σ_T; Cooper and Haltiwanger (2006) the adjustment costs; Gomme et al. (2011) the return on capital.

Land, housing and wealth. Knoll et al. (2017) and Davis and Heathcote (2007) show that land, not structures, drives long-run house prices; our housing block separates the two. Piketty (2014), Piketty and Zucman (2014), Saez and Zucman (2016), Smith et al. (2023), Kuhn et al. (2020) and Jordà et al. (2019) provide the wealth and return facts. We do not identify $r > g$ with rising concentration: Proposition 6 shows that r and the composition of claims matter separately.

Policy. Guerreiro et al. (2022), Thuemmel (2023) and Costinot and Werning (2023) analyze robot taxes; our Proposition 9 gives the limit-economy version (a compute tax raises wages one-for-one with the tax-inclusive user cost). The land-tax and citizen-fund scenarios go back to George (2009), Meade (1964) and Atkinson (2015), and to the "universal dividend" the source mentions in its discussion.

Compute and energy. Sevilla et al. (2022) and Hobbhahn and Besiroglu (2022) document compute trends and the price–performance of hardware (doubling about every two years in the data we use); Koomey et al. (2011) the historical efficiency of computing; Masanet et al. (2020), Shehabi et al. (2024) and International Energy Agency (2025) the energy use of data centers (4.4% of US electricity in 2023). These are the sources behind the compute-price, energy-efficiency and energy-share inputs of Section 8.

3. The benchmark: Restrepo's economy, reproduced and audited

3.1. Primitives

Restrepo (2025) describes an economy that produces output $Y_t = F(\{X_t(\omega)\}_{\omega\in\Omega})$ by completing a finite set Ω of types of work, each of which can be done by people or by computers:

$$X_t(\omega) = L_t(\omega) + \frac{Q_t(\omega)}{\alpha_t(\omega)}, \tag{1}$$

where $\alpha_t(\omega)$ is the compute needed to replicate one worker and may be infinite. F is increasing, differentiable, concave and constant returns to scale. Skills $s \in \mathscr{S}$, in supply $H(s)$, can do the work in $\Omega(s)$, and the $\Omega(s)$ partition Ω. Total compute is Q_t. Two premises drive the results: AGI (Premise 1: every $\alpha_t(\omega)$ becomes finite and converges to $\alpha(\omega) \in (0,\infty)$) and abundant compute (Premise 2: Q_t grows without bound). Work ω is a *bottleneck* if output cannot grow without bound unless $X_t(\omega)$ or its marginal product does; it is *accessory* otherwise.

The chapter's results, with the page of the July 2025 draft, are: all bottlenecks are eventually automated while some accessory work may be left to humans (Proposition 1, p. 7); output converges to $A[Q_t + \sum_{s\in\mathscr{A}} \mathrm{CEU}(s)H(s) + \sum_{s\in\mathscr{N}} \widetilde{\mathrm{CEU}}(s)H(s)]$ with $\mathrm{CEU}(s) = \max_{\omega\in\Omega(s)} \alpha(\omega)$ and $\widetilde{\mathrm{CEU}}(s) < \mathrm{CEU}(s)$ for the skills whose most compute-intensive work is not automated (Proposition 2, p. 8); output grows at the rate of compute (Proposition 3, p. 9); the real price of compute converges to A and wages to $A\,\mathrm{CEU}(s)$ or $A\,\widetilde{\mathrm{CEU}}(s)$ (Proposition 4, p. 10); the compute share of GDP converges to one and the labor share to zero (Proposition 5, p. 10); and the sum of wages is higher with AGI than without (Proposition 6, p. 11). Section 1.2 of the source, labelled preliminary, describes two transitions: when compute binds, the price of compute falls monotonically and wages fall gradually around each automation date and otherwise grow (Proposition 7, p. 12); when technology binds, wages can jump down and the last-automated skills earn a transitory premium (Proposition 8, p. 13). No proofs are given; every proposition is followed by a heuristic argument.

3.2. Replication

We implement the economy with six types of work: four bottleneck classes aggregated by a CES with elasticity $\sigma_B = 0.5$ and two accessory classes entering output linearly, $Y = X_B + \sum_a \gamma_a X_a$. Skill L does work in $\{1,2,a_L\}$ with $\alpha = (1,2,3)$, skill H in $\{3,4,a_H\}$ with $\alpha = (4,8,6)$, supplies $(H_L,H_H) = (0.6,0.4)$, $\gamma_a = 0.05$. The competitive allocation is the solution of the concave program $\max Y$ subject to the labor and compute constraints, and factor prices are its multipliers (Appendix B). Figure 1 traces the equilibrium as the endowment Q grows from 10^{-2} to 10^5; Table 1 reports selected points.

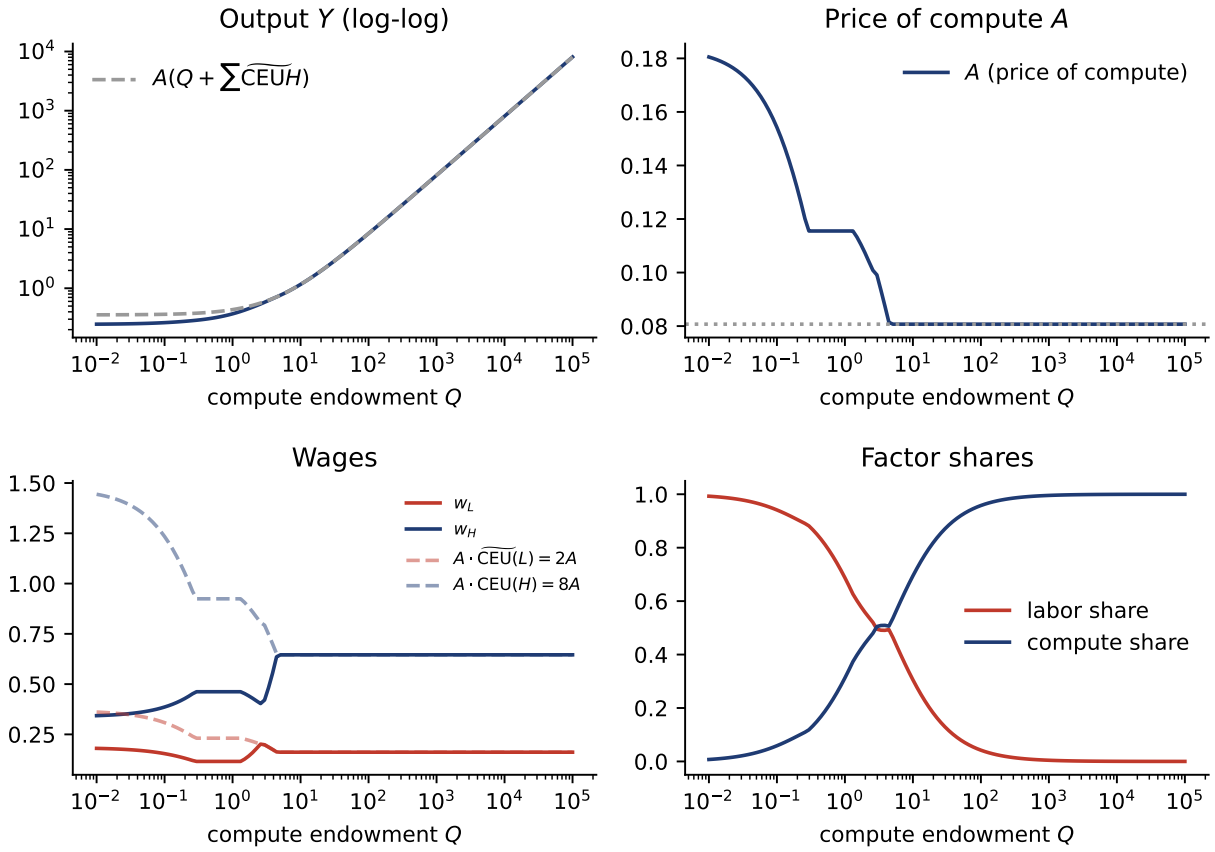


Figure 1. Replication of Restrepo (2025) in the finite-task CES economy. Output becomes linear in compute (dashed line, upper left) once every bottleneck task is automated ($Q \geq 5$); the price of compute falls in steps and converges (upper right); wages are pinned at A times the compute-equivalent of each skill (lower left); the labor share converges to zero (lower right). Source: `scripts/10_benchmark.py`.

Table 1. Selected points of the benchmark economy.

Q	Y	A	w_L	w_H	w_L/A	w_H/A	labor share
0	0.2455	0.0000	0.1838	0.3381	–	–	1.0000
0.01	0.2474	0.1805	0.1805	0.3431	1.00	1.90	0.9927
0.1	0.2624	0.1542	0.1542	0.3861	1.00	2.50	0.9412
1	0.3697	0.1155	0.1155	0.4621	1.00	4.00	0.6875
10	1.1624	0.0807	0.1614	0.6458	2.00	8.00	0.3056
100	8.4271	0.0807	0.1614	0.6458	2.00	8.00	0.0421
1000	81.0747	0.0807	0.1614	0.6458	2.00	8.00	0.0044
10000	807.5507	0.0807	0.1614	0.6458	2.00	8.00	0.0004
100000	8072.3102	0.0807	0.1614	0.6458	2.00	8.00	0.0000

Every statement of Propositions 1–6 holds exactly in the replication, and the shape of the transition is the compute-binding one of Proposition 7: A is non-increasing, wages fall around the two automation dates of each skill and then grow until the last date, after which they are pinned at A times the compute-equivalent. Two details that the source states only in a footnote (fn. 6, p. 9) are visible. First, skill H earns $A \cdot \mathrm{CEU}(H) = 8A$: its most compute-intensive task (task 4) is automated and labor concentrates there. Second, skill L earns $2A$, below $A \cdot \mathrm{CEU}(L) = 3A$: the accessory task with $\alpha = 3$ is never automated — its marginal value $\gamma_a = 0.05$ is below $A\alpha = 0.24$ — and, in our parametrization, it is not performed at all, because the same marginal value is below what L-workers earn in task 2. The "discount" of Proposition 4 arises whenever the max-α work of a skill is left unautomated, whether or not anyone does it; the source's description of group $\mathcal{N}$ as "skills that perform accessory work not automated" is one of the two cases.

3.3. What the benchmark cannot say

Three features of the source make it silent on the questions of this paper. Compute is an endowment: there is no price of compute capital, no investment, no saving, no interest rate, no depreciation, and therefore no answer to what compute owners' *wealth* is worth. There are no complementary inputs: energy appears only as a reinterpretation of Q_t (p. 4), and the source itself notes (p. 11) that Proposition 6 breaks when the AGI economy "uses some finite resource that is also needed by human labor". And there are no households: nothing is said about who owns what, how ownership evolves, or what the welfare of workers is beyond their wage bill. Sections 4–6 supply these elements; Section 7 shows which of Propositions 2–6 survive and which do not.

4. A task economy with priced compute and scarce complements

The quantitative model keeps the work equation and the skill partition of Section 3 and adds four capital stocks, energy, land and housing. Time is discrete and annual; $t = 0$ is 2026. All quantities are real in units of the 2026 final good. Notation is collected in Table 2 and in the model dictionary of the replication package.

4.1. Tasks and automation

There are four bottleneck task classes $i \in B = \{1,2,3,4\}$ — routine, lower cognitive, higher cognitive, expert — and two skills with $\Omega(L) = \{1,2\}$ and $\Omega(H) = \{3,4\}$. Workers supply H_s efficiency units inelastically; efficiency units, land and the energy resource grow at the trend rate g_Z (Harrod-neutral non-AI technical progress), so that the frozen-technology economy has an exact

balanced growth path and all statements about wages below are per efficiency unit. Task output is

$$X_{i,t} = L_{i,t} + \frac{Q_{i,t}}{\alpha_{i,t}}, \qquad X_{B,t} = \Big(\sum_{i\in B} \nu_i X_{i,t}^{\rho_B}\Big)^{1/\rho_B}, \quad \sigma_B = \tfrac{1}{1-\rho_B} < 1, \tag{2}$$

where $\alpha_{i,t}$ is the compute needed to replicate one efficiency unit of labor in task i. Automation feasibility improves along a scenario path: in the baseline $\log \alpha_{i,t}$ declines linearly at rate g_α from its 2026 value to a terminal value α_i^∞, ordered $(1,2,4,8)$ across tasks so that the expert task is the last to become cheap to replicate; in the "rapid AGI" scenario all $\alpha_{i,t}$ jump to their terminal values at a date. Whether a task is actually automated is an equilibrium outcome: firms use compute in task i if and only if $A_t \alpha_{i,t} \le w_{s(i),t}$, and labor of skill s flows to the task in $\Omega(s)$ with the highest marginal value. Automation can therefore be feasible but unprofitable, exactly as in the source.

4.2. Final goods, structures, land and energy

Final output is a nested CES in the task composite, structures and machinery K^N, production land T and energy E^Y:

$$V_t = \big[\nu_N (K_t^N)^{\rho_N} + (1-\nu_N) X_{B,t}^{\rho_N}\big]^{1/\rho_N}, \qquad Y_t = Z\big[\nu_T T_t^{\rho_T} + \nu_E (E_t^Y)^{\rho_T} + \nu_V V_t^{\rho_T}\big]^{1/\rho_T}, \tag{3}$$

with $\sigma_N = 1/(1-\rho_N)$ the elasticity between tasks and structures and $\sigma_T = 1/(1-\rho_T)$ the elasticity between the task–structure bundle and the two physical inputs that cannot be reproduced at will. The baseline uses $\sigma_N = 0.7$ and $\sigma_T = 0.6$; Section 10 varies both across the regimes of Table 3. Competitive final-goods firms rent structures at R_t^N, buy compute services at A_t, hire labor at $w_{s,t}$, rent land at R_t^T and buy energy at P_t^E; prices equal marginal products.

4.3. Compute as capital

Installed compute capacity K_t^Q delivers $Q_t = u_t K_t^Q$ units of services, $u_t \le 1$ the utilization rate, and each unit of service requires ε_t units of energy. Capacity is bought one period ahead at price p_t^Q in units of the final good and depreciates at δ_Q:

$$K_{t+1}^Q = (1-\delta_Q) K_t^Q + I_t^Q, \qquad p_t^Q = p_0^Q \max\{e^{-g_{pQ} t}, \underline{p}/p_0^Q\}, \qquad \varepsilon_t = \varepsilon_0 \max\{e^{-g_\varepsilon t}, \underline{\varepsilon}/\varepsilon_0\}. \tag{4}$$

The price decline is Moore's law for installed capacity, and the floor $\underline{p}$ is the date at which it ends; the energy-efficiency decline is Koomey's law with its own floor. The owner of a unit of capacity earns the net rental $R_t^Q = u_t(A_t - \varepsilon_t P_t^E)$; capacity is idle ($u_t < 1$) only when its energy cost exceeds its marginal value, which never happens on the baseline path and happens in the energy-constrained scenarios.

4.4. Energy, housing and the other capital stocks

Energy is produced from energy capital and a fixed resource (sites, fuel, grid rights), $E_t = Z^E (K_t^E)^{1-\eta} (R_t^E)^{\eta}$, and allocated between compute and final production, $E_t = \varepsilon_t Q_t + E_t^Y$. The resource earns the rent $\eta P_t^E E_t$. Housing services $H_t^S = (K_t^H)^{1-\nu_H} (T_t^H)^{\nu_H}$ are produced from residential structures and residential land and rented to households at R_t^H; households spend a Cobb–Douglas share θ_H of total expenditure on housing, so $R_t^H H_t^S = \theta_H/(1-\theta_H) C_t$ with C_t aggregate goods consumption. Structures, energy capital and housing structures are bought at price one, depreciate at $\delta_N, \delta_E, \delta_H$ and, like compute, are subject to quadratic adjustment costs $\frac{\phi_k}{2}(x_t^k - \delta_k)^2 K_t^k p_t^k$, $x_t^k = I_t^k / K_t^k$, so that Tobin's marginal q is $q_t^k = p_t^k [1 + \phi_k (x_t^k - \delta_k)]$ (Hayashi, 1982).

4.5. National accounting

Compute services and energy are intermediate inputs. Value added is $\text{GDP}_t = Y_t + R_t^H H_t^S$, and the income identity

$$\text{GDP}_t = \sum_s w_{s,t} H_{s,t} + R_t^N K_t^N + R_t^Q K_t^Q + R_t^E K_t^E + \eta P_t^E E_t + R_t^T T_t + R_{K,t}^H K_t^H + R_t^{TH} T_t^H \tag{5}$$

holds exactly (Euler's theorem on each constant-returns block), with R_K^H and R^{TH} the rentals of housing structures and residential land. Equation (5) is tested numerically at every date of every scenario (Appendix E); a deliberately broken version of the accounting — counting compute services as final output — fails the test, as it should. There is no double counting of compute services and compute capital: $A_t Q_t$ is the gross revenue of the compute sector, $R_t^Q K_t^Q$ its value added net of energy. Profits are zero under perfect competition; the AI-monopoly scenario adds a markup wedge on compute services that appears as a separate profit line accruing to the owners of compute firms.

4.6. Who owns what

Every stock — the four capital types, production land, residential land — is held through a single fund whose shares are the only asset households trade. Under perfect foresight all claims earn the same return r_{t+1} between t and $t+1$ (no arbitrage), and the fund's ex-dividend value is

$$W_t = \sum_{k \in \{N,Q,E,H\}} q_t^k K_{t+1}^k + P_t^T T + P_t^{TH} T^H. \tag{6}$$

Who holds the fund's shares is the subject of Section 5: capitalists alone in the analytical two-class economy, a continuum of households in the heterogeneous-agent economy. The citizen-fund scenario gives a public fund a fixed share θ_F of the claims and distributes its net payout equally; the

foreign-ownership scenario lets a share of compute equity be held abroad, so that compute payouts leave the economy and foreign investors finance the corresponding share of compute investment.

Table 2. Notation.

Symbol	Meaning	Symbol	Meaning
$X_{i,t}, L_{i,t}, Q_{i,t}$	task output, labor, compute in task i	K^N, K^Q, K^E, K^H	structures, compute, energy capital, housing
$\alpha_{i,t}$	compute per worker-equivalent in task i	p_t^Q, ε_t	compute price, energy per compute unit
$\widetilde{\mathrm{CEU}}(s)$	compute-equivalent unit of skill s	T, T^H, R^E	production land, residential land, energy res
$w_{s,t}$	wage per efficiency unit of skill s	q_t^k, x_t^k	Tobin's q, investment rate of capital k
A_t	price of compute services	P_t^T, P_t^{TH}	land prices
R_t^k	rental of capital k	W_t	aggregate wealth (ex-dividend value of all c
P_t^E, R_t^H	energy price, housing rent	r_{t+1}	real return on all claims between t and $t+1$
$\sigma_B, \sigma_N, \sigma_T$	elasticities (tasks; tasks–structures; top nest)	β, γ	discount factor, curvature of utility

5. Households, saving and the ownership of compute

Income can flow from labor to capital without anything happening to the distribution of wealth if everybody owns capital in proportion to their consumption, and it can concentrate wealth without any change in factor shares if only some households save. The model therefore tracks ownership explicitly, in two layers.

5.1. The two-class economy

The analytical layer has two types of households. Capitalists, of unit mass, own all the claims of Section 4, have no labor income, and choose expenditure x_t^K on goods and housing to maximize $\sum_t \beta^t u(c_t, h_t)$ with $u = (c^{1-\theta_H} h^{\theta_H})^{1-\gamma}/(1-\gamma)$. With the Cobb–Douglas inner aggregator the problem reduces to one over total expenditure with the ideal price index $P_t \propto (R_t^H)^{\theta_H}$, and the Euler equation is

$$\left(\frac{x_{t+1}^K}{x_t^K}\right)^{\gamma} \left(\frac{P_{t+1}}{P_t}\right)^{1-\gamma} = \beta(1+r_{t+1}). \tag{7}$$

Workers of skill s, of mass H_s, consume their wage income plus transfers each period. The two-class economy is the one in which Propositions 3–9 are proved and in which all scenarios of Section 9 are solved; it is also the limit of the heterogeneous economy in which a group of patient households holds almost all wealth. It has two purposes: the capitalists' Euler equation (7) is the stochastic discount factor that prices every claim, and the hand-to-mouth workers make the "demand compression" channel of Section 9 visible, because any fall in the wage bill is a one-for-one fall in workers' expenditure.

5.2. The heterogeneous-agent economy

The quantitative distributional layer replaces the two types by a continuum. A household is characterized by a permanent skill $s \in \{L, H\}$, a permanent discount factor $\beta_j \in \{\beta_1, \beta_2, \beta_3\}$ with masses $(0.5, 0.3, 0.2)$ within each skill (Krusell and Smith, 1998; Carroll et al., 2017), an idiosyncratic productivity e_t following a five-state Rouwenhorst approximation to an AR(1) in logs with persistence 0.95 and innovation standard deviation 0.20, and assets $a_t \geq 0$ invested in the fund. Income is $w_{s,t} e_t + \tau_t$ with τ_t the lump-sum transfer of the policy scenarios; expenditure x_t and savings a_{t+1} satisfy $x_t + a_{t+1} = a_t(1 + r_t) + w_{s,t} e_t + \tau_t$, and the Euler equation holds with equality whenever the borrowing constraint is slack. The problem is solved by the endogenous grid method (Carroll, 2006) on a 300-point grid and the distribution by non-stochastic simulation (Appendix B). The discount factors are set so that stationary savings at the 2026 prices equal the two-class wealth W_0, so that both layers start from the same balance sheet; the resulting wealth Gini and top-10% share are reported against the Distributional Financial Accounts in Section 8 (Board of Governors of the Federal Reserve System, 2026a). The model has no return heterogeneity and cannot reproduce the top-1% share (Gabaix et al., 2016; Hubmer et al., 2021); we report it as an untargeted moment.

General equilibrium in the heterogeneous economy is the requirement that the fund's shares are willingly held: $\int a_{t+1} \, \mathrm{d}\mu_t = W_t$ at every date, where W_t is the ex-dividend value of all claims. We solve for the path of returns by relaxation: given a path $\{r_{t+1}\}$, the production side (four capital Euler equations and the goods market) is solved with the banded Newton method of Appendix B, the household block is solved backward and simulated forward, and the return is lowered where households demand more claims than exist. The two-class solution is the starting point.

5.3. Wealth accounting

For both layers we decompose the change in aggregate wealth between any two dates into net saving, net payouts received and reinvested, and revaluation:

$$W_{t+1} - W_t = \underbrace{\left[(1 + r_{t+1}) W_t - W_t\right]}_{\text{return}} - \underbrace{\left[\textstyle\sum_k D^k_{t+1} + R^T_{t+1} T + R^{TH}_{t+1} T^H\right]}_{\text{payouts}} + \ldots, \tag{8}$$

where the return splits further into the payout yield and the capital gain $\sum_k (q^k_{t+1} K^k_{t+2} - q^k_t K^k_{t+1})$ and the land-price change. In the heterogeneous economy the same decomposition is applied household by household in simulation, which is what allows the paper to say whether wealth concentrates because the rich save more or because what they already own is revalued (Section 9).

6. Asset prices

6.1. The discount factor and the return on claims

Under perfect foresight the capitalists' Euler equation (7) defines a one-period discount factor $M_{t,t+1} = \beta\,(x^K_{t+1}/x^K_t)^{-\gamma}(P_{t+1}/P_t)^{\gamma-1}$ with $M_{t,t+1}(1+r_{t+1}) = 1$. Every claim is priced by $V_t = M_{t,t+1}(D_{t+1}+V_{t+1})$, and because there is no aggregate risk every claim earns r_{t+1}: the model has a term structure (the path of r_{t+1}) but no risk premia. The marginal investor is an unconstrained saver in both layers of Section 5: in the two-class economy by construction, in the heterogeneous economy because households with positive savings satisfy the Euler equation with equality and the return is the same for all of them. Borrowing constraints change who saves, not the pricing kernel (Krusell and Smith, 1997).

6.2. Capital goods

For each capital type k with purchase price p^k_t the firm maximizes the present value of net payouts $D^k_t = R^k_t K^k_t - p^k_t I^k_t - \frac{\phi_k}{2}(x^k_t - \delta_k)^2 K^k_t p^k_t$. The first-order conditions are $q^k_t = p^k_t[1+\phi_k(x^k_t-\delta_k)]$ and

$$q^k_t(1+r_{t+1}) = R^k_{t+1} + p^k_{t+1}\left[\phi_k(x^k_{t+1}-\delta_k)x^k_{t+1} - \frac{\phi_k}{2}(x^k_{t+1}-\delta_k)^2\right] + (1-\delta_k)q^k_{t+1}, \tag{9}$$

and by Hayashi (1982) the ex-dividend value of the firm is $V^k_t = q^k_t K^k_{t+1}$. A negative payout is a net issuance of shares; there is no dilution problem because V^k_t is the value of all shares after the issuance and the fund buys the new shares. Three valuation objects are reported: the ratio of market value to *replacement cost*, q^k_t/p^k_t, which exceeds one in investment booms; the price–earnings ratio $V^k_t/(R^k_t K^k_t)$, which for compute is dominated by depreciation and the expected price decline (Proposition 6); and the price–payout ratio V^k_t/D^k_t, which is undefined while the sector issues shares. Market-to-*book* ratios at historical cost are higher than market-to-replacement for compute, because the book value of a stock bought at yesterday's price overstates its replacement cost when prices fall; we report replacement-cost ratios throughout.

6.3. Land, housing and the aggregate balance sheet

Production land and residential land are priced by $P^T_t(1+r_{t+1}) = (1-\tau_T)R^T_{t+1} + P^T_{t+1}$ with the terminal value given by the Gordon formula at the terminal growth rate of rents, and the transversality condition is verified by the horizon check of Appendix B. A house is a bundle of a structure and a plot: its price is $q^H_t K^H_{t+1} + P^{TH}_t T^H$ and its price–rent ratio is that bundle divided by $R^H_t H^S_t$. The structure part is bounded by replacement cost and the land part is not, which is why land is where the value of an AI boom goes when the top nest has gross complements (Proposition 2c).

Aggregate wealth W_t is the ex-dividend value of all claims. It differs from the replacement cost of the physical stocks $\sum_k p_t^k K_{t+1}^k$ by the adjustment-cost premia $\sum_k (q_t^k - p_t^k) K_{t+1}^k$ and by the value of land, and from national wealth in the foreign-ownership scenario by the claims held abroad. We never add V_t^Q to $p_t^Q K_{t+1}^Q$: the first is the market value of the claim on the second.

6.4. Decomposing valuation changes

Along a deterministic path, the log price–payout ratio of any claim satisfies the Campbell–Shiller identity exactly, so the difference between two scenarios in $\log(V_t/D_t)$ can be written as the present value of differences in payout growth minus the present value of differences in returns (Campbell and Shiller, 1988a). We use this to separate, for each claim and each scenario, the *cash-flow* channel (AI raises or lowers what the claim pays) from the *discount-rate* channel (AI changes the return the marginal investor requires, Proposition 7), and we add the *supply* channel for capital goods (new capital is produced at replacement cost, so the value of existing claims is diluted by entry at $q^k = p^k$) and the *competition* channel in the AI-monopoly scenario (markup rents). Because the paths are deterministic there is no risk channel; the stochastic AGI-arrival extension of Section 10 shows what a hazard of arrival adds.

6.5. Bubbles

We impose the transversality condition on every claim, so there are no bubbles in the baseline. A rational bubble on compute equity is inconsistent with free entry at replacement cost (new capacity can always be produced at p_t^Q); the only claims that could carry a bubble in this economy are land and non-reproducible intangibles, and only if the interest rate falls to the growth rate. We check along every path that $r_t > g_t$ asymptotically, which is the condition under which the land price recursion converges.

7. Analytical results: when does income leave labor, and where does it go?

This section isolates the mechanisms in a version of the model that can be solved by hand. It keeps the task block of Section 3, adds one complementary input S and a price for compute, and asks three questions: whether the labor-share asymptote of Restrepo (2025) survives; whether the *level* of wages survives; and what the valuation of the claims to the income that leaves labor looks like. Proofs are in Appendix A. Every statement is conditional on the stated assumptions; the quantitative model of Sections 8–9 relaxes several of them numerically.

7.1. The limit economy

Assumption 1 (Limit economy)**.** Final output is $Y = ZF(X_B, S)$ with $F(X_B, S) = \left[\nu X_B^{\rho} + (1-\nu)S^{\rho}\right]^{1/\rho}$, $\rho = (\sigma-1)/\sigma$, $\sigma \in (0,\infty)$, $\nu \in (0,1)$, and the Cobb–Douglas limit at $\sigma = 1$. The bottleneck composite is $X_B = (\sum_{i\in B} \nu_i X_i^{\rho_B})^{1/\rho_B}$ with $\sigma_B < 1$ and $X_i = L_i + Q_i/\alpha_i$, $\alpha_i \in (0,\infty)$ for all i (Premise 1 has been reached), skills $s \in \{L, H\}$ with inelastic supplies H_s and a partition $\Omega(s)$ of B. $S > 0$ is either fixed or reproducible at constant price one with depreciation δ_S.

Lemma 1 (Compute-equivalent labor)**.** *Under Assumption 1 there is a finite $\bar{Q}$ such that for all $Q \geq \bar{Q}$ the output-maximizing allocation uses compute in every bottleneck task, and*

$$X_B = \mathscr{B}\widetilde{Q}, \qquad \widetilde{Q} \equiv Q + \sum_s \widetilde{\mathrm{CEU}}(s) H_s, \qquad \mathscr{B} = \Big[\sum_{i\in B} \nu_i^{\sigma_B} \alpha_i^{1-\sigma_B}\Big]^{-1/(1-\sigma_B)}, \tag{10}$$

where $\widetilde{\mathrm{CEU}}(s) = \max_{i\in\Omega(s)} \alpha_i$ is the compute-equivalent unit of skill s among its tasks. The price of compute services is $A = ZF_X\mathscr{B}$ and $w_s = A\,\widetilde{\mathrm{CEU}}(s)$.

Lemma 1 is Restrepo's Proposition 2 applied to the composite instead of to final output: once every bottleneck is automated, the composite is linear in effective compute. The constant $\mathscr{B}$ is the unit cost of the composite in compute units; it falls when any α_i falls. With the accessory block of Section 3 the same statement holds with $\widetilde{\mathrm{CEU}}(s)$ replaced by the discounted value of footnote 6 of the source; the quantitative model has no accessory tasks (Section 4).

7.2. The labor share result is robust; the wage result is not

Proposition 1 (Labor share)**.** *Under Assumption 1, for any $Q \geq \bar{Q}$ and any S,*

$$\frac{\sum_s w_s H_s}{Y} \leq \frac{\sum_s \widetilde{\mathrm{CEU}}(s) H_s}{Q + \sum_s \widetilde{\mathrm{CEU}}(s) H_s} \xrightarrow[Q\to\infty]{} 0. \tag{11}$$

The bound holds for every constant-returns F that is increasing in X_B, with or without complementary inputs, and whether compute is an endowment or a priced capital good.

The proposition says that the only ingredient behind Restrepo's Proposition 5 is the *within-task* perfect substitutability of labor and compute: labor is paid the value of the compute it saves, and the value of all compute cannot exceed output. Nothing about the rest of production matters. What the rest of production determines is the *level* at which wages settle, and who receives the income that labor loses.

Proposition 2 (Scarce complement: factor shares and wage levels)**.** *Under Assumption 1 with S fixed, as $Q \to \infty$:*

(a) *if* $\sigma > 1$*:* $A \to Z\nu^{1/\rho}\mathscr{B} > 0$*,* $w_s \to A\widetilde{\mathrm{CEU}}(s) > 0$*, the compute share* $s_Q \equiv AQ/Y \to 1$*,* $Y/Q \to A$*; Propositions 2–6 of Restrepo (2025) hold;*

(b) *if* $\sigma = 1$*:* $Y = Z(\mathscr{B}\widetilde{Q})^{\nu}S^{1-\nu}$*,* $A = \nu Y/\widetilde{Q} \to 0$*,* $w_s \to 0$*, factor shares are constant at* $(\nu, 1-\nu)$ *and output grows at* ν *times the growth rate of compute;*

(c) *if* $\sigma < 1$*:* $Y \to Z(1-\nu)^{1/\rho}S < \infty$*,* $A \to 0$ *at rate* $\widetilde{Q}^{\rho-1}$*,* $w_s \to 0$*, and the share of S converges to one.*

In cases (b) and (c) the total wage bill converges to zero, so the "no-AGI zone" argument behind Restrepo's Proposition 6 fails for Q large enough, as the source anticipates for "a finite resource that is also needed by human labor" (p. 11).

The economics is simple. Labor is paid its compute-equivalent, and the compute-equivalent is worth whatever compute is worth. With gross complementarity, abundant compute is worth nothing at the margin: a worker whose work can be replicated by compute that costs nothing is paid nothing, however skilled. The source's optimistic reading ("humans can still work . . . but their contribution becomes economically negligible") therefore depends on compute remaining scarce, which in the source is guaranteed by treating Q_t as an endowment that grows exogenously.

7.3. Compute as a reproducible capital good

Treating compute as capital changes which object is scarce. Let installed capacity K^Q be bought at price p_t^Q with depreciation δ_Q, no adjustment costs, and let the capitalists' Euler equation hold with discount factor β and curvature γ.

Proposition 3 (Balanced AGI growth with a floor on the compute price)**.** *Let S be reproducible at price one and let* $p_t^Q \to \bar{p} > 0$*. Under Assumption 1 with* $\sigma < 1$ *there exists a balanced growth path on which Y, S,* K^Q *and capitalists' consumption grow at a common rate g, the interest rate is* $r = R_S - \delta_S$ *with* R_S *the marginal product of S, and*

$$A = (r+\delta_Q)\bar{p}, \qquad w_s = (r+\delta_Q)\,\bar{p}\,\widetilde{\mathrm{CEU}}(s) > 0, \qquad s_Q = \frac{\nu(\mathscr{B}\widetilde{Q}/S)^{\rho}}{\nu(\mathscr{B}\widetilde{Q}/S)^{\rho} + (1-\nu)} \in (0,1). \quad (12)$$

Along this path the labor share converges to zero (Proposition 1), the compute share is constant and strictly below one, and the wage equals the user cost of the compute that replaces the worker.

Proposition 4 (Compute price converging to zero)**.** *Let* $p_t^Q \to 0$ *(Moore's law never ends) with S reproducible at price one.*

(a) *If* $\sigma > 1$*, no balanced growth path exists: the growth rate of compute and of output increases without bound (the economy is an AK economy whose capital price falls to zero).*

(b) If $\sigma \leq 1$, the economy converges to an AK path in S with $Y/S \to Z(1-\nu)^{1/\rho}$ (Cobb–Douglas: Y/S grows without bound but S's share stays $1-\nu$); $A \to 0$, $w_s \to 0$, and the income share and the market value of compute capital converge to zero.

Propositions 2–4 give the regime map summarized in Table 3. The vanishing labor share is common to all cells. The destination of the income that leaves labor is not: it is compute owners only when compute stays scarce (an endowment with $\sigma > 1$, or a priced capital good whose price has a floor), and it is the owners of the scarcest complement (structures, land, energy sites) otherwise. The level of wages is positive exactly in the cells in which compute keeps a positive price.

Table 3. Asymptotic regimes of the limit economy (all cells: labor share $\to 0$).

Compute	$\sigma > 1$ (substitutes)	$\sigma = 1$	$\sigma < 1$ (complements)
Endowment (source)	$w > 0$, $s_Q \to 1$, $g_Y = g_Q$	$w \to 0$, shares constant, $g_Y = \nu g_Q$	$w \to 0$, Y bounded by S
Priced, floor $\bar{p} > 0$	AK in compute, $w > 0$	BGP, $w > 0$, shares constant	BGP, $w > 0$, shares constant
Priced, $p^Q \to 0$	explosive	AK in S, $w \to 0$	AK in S, $w \to 0$, $s_Q \to 0$
Energy-bounded ($Q \leq \bar{E}/\bar{\varepsilon}$)	labor share $\to$ positive limit (Proposition 5)		

S fixed in the first row, reproducible at constant price in the next two. w is the wage per efficiency unit; with trend labor-augmenting progress at g_Z all wage statements hold per efficiency unit.

Proposition 5 (Energy bottleneck)**.** *Suppose compute services require $\varepsilon_t \geq \bar{\varepsilon} > 0$ units of energy per unit and energy available to compute is bounded by $\bar{E}$. Then $Q_t \leq \bar{E}/\bar{\varepsilon}$ for all t, the price of compute is bounded below by a positive constant whenever $F_X > 0$, and the labor share converges to a positive limit no smaller than $s_Q^\infty \sum_s \widetilde{\mathrm{CEU}}(s) H_s / (\bar{E}/\bar{\varepsilon} + \sum_s \widetilde{\mathrm{CEU}}(s) H_s)$, where s_Q^∞ is the limiting share of the task composite. Restrepo's Premise 2 fails and so does his Proposition 5.*

7.4. Valuation: capital's share rises, wealth-to-income can fall

Let s_k denote the gross income share of capital type k and $\pi_k \geq 0$ the rate at which its purchase price declines ($\pi_Q = g_{pQ}$ until the floor, zero for structures). On a balanced growth path with growth g and interest rate r, no adjustment costs and competitive rentals:

Proposition 6 (Wealth-to-income ratio)**.**

$$\frac{W}{Y} = \sum_{k \in \{N,Q,E,H\}} \frac{s_k (1+g)}{(1+r)(1+\pi_k) - (1-\delta_k)} + \frac{s_T (1+g_T)}{r - g_T}, \tag{13}$$

where s_T and g_T are the share and growth rate of land rents. A reallocation of income shares from a long-duration claim (land, structures) to compute capital lowers W/Y whenever $(1+r)(1+\pi_Q) - (1-\delta_Q) > (1+r) - (1-\delta_N)$, i.e. whenever compute depreciates faster or its price

falls faster than that of the claim it displaces. The price–earnings ratio of compute capital is $(1+g)/[(1+r)(1+\pi_Q)-(1-\delta_Q)]$*, which at* $r=0.035$*,* $\delta_Q=0.25$*,* $\pi_Q=0.15$ *is about* 2.3 *against* 8 *for structures and* $1/(r-g_T)$ *for land.*

The proposition separates three objects that the source and much commentary conflate: the share of income accruing to compute, the market value of compute capital, and aggregate wealth. Income can flow overwhelmingly to compute while the *capitalized* value of that income is small, because compute is a short-duration, rapidly cheapening asset. Whether the wealth-to-income ratio rises with AI depends on whether the income that labor loses is captured by claims that are long-lived (land, scarce sites, intangible moats) or short-lived (chips).

Proposition 7 (Discount-rate channel)**.** *With the capitalists' Euler equation and constant growth* g *of their expenditure,* $1+r=\beta^{-1}(1+g)^{\gamma}$*. For a claim whose cash flow grows at the common rate* g *(land under Cobb–Douglas, structures on the BGP), the price–cash-flow ratio is* $(1+g)/[\beta^{-1}(1+g)^{\gamma}-(1+g)]$*, which is decreasing in* g *if and only if* $\gamma>1$*. An AI-driven increase in the growth rate therefore* lowers *price–rent and price–earnings ratios of long-duration claims when the elasticity of intertemporal substitution is below one, and raises them otherwise; the level of prices rises in both cases if cash flows rise.*

Proposition 7 is the general-equilibrium reason why "higher profits" need not mean "higher valuations": the same boom that raises cash flows raises the rate at which the marginal investor discounts them. The quantitative section decomposes the valuation paths into these two channels.

7.5. Ownership

Proposition 8 (Two-class distribution)**.** *In the two-class economy, workers' expenditure equals* $\sum_s w_s H_s+\tau_t$ *where* τ_t *are transfers. With* $\tau_t=0$ *and in any regime in which* $Y_t\to\infty$*, the consumption share of workers converges to zero; their consumption* level *converges to* $A\sum_s \widetilde{\mathrm{CEU}}(s)H_s$*, positive in the cells of Table 3 with* $w>0$ *and zero otherwise. With a citizen fund paying workers a share* θ_F *of aggregate net payouts, the workers' consumption share is bounded below by* $\theta_F\,\bar{d}\,\mu_W$*, where* $\bar{d}$ *is the net-payout-to-GDP ratio on the balanced path and* μ_W *the workers' population share.*

Proposition 9 (Wages as the tax-inclusive rental price of the replacement)**.** *On the balanced path of Proposition 3 with a tax* τ_Q *on compute rental income rebated lump-sum, and an energy cost* εP_E *per unit of compute,*

$$w_s=\widetilde{\mathrm{CEU}}(s)\left[\frac{(r+\delta_Q)\bar{p}}{1-\tau_Q}+\varepsilon P_E\right]. \tag{14}$$

A compute tax raises wages one-for-one with the tax-inclusive user cost, at the cost of a lower compute stock and lower output; a land tax leaves all wages, allocations and prices except the land price unchanged.

The wage expression is the paper's one-line summary: in the AGI limit a worker earns what it costs, inclusive of taxes and energy, to rent the compute that would do the worker's job. Policies that make compute dearer raise wages and lower output; policies that tax the scarce complement (land) redistribute without changing production.

7.6. What the propositions do not say

They are statements about limits. They do not say how fast the labor share falls, whether wages rise before they fall (they can: Restrepo's Propositions 7–8 describe the transitory premium of the last-automated skills), how much of the wealth that households hold in 2026 is revalued, or how a heterogeneous population shares the income of the fund. Those are the questions for the quantitative model. The propositions also treat σ as a structural constant; if physical complements can themselves be economized over time (land through density, energy through efficiency), the relevant elasticity is the long-run one, and the quantitative model lets ε_t fall to a floor for that reason.

8. Calibration and identification

The base point is the US economy in 2026, normalized so that GDP per efficiency unit is one and the two skills together supply one unit of labor. It is constructed as a balanced growth path of the frozen-technology economy and then re-derived for each scenario so that the compute sector's 2026 rental is consistent with the price path the scenario anticipates. Table 4 lists every parameter with its value, its status — *data* (read from a retrieved series), *calibrated* (implied by base-point shares and the model's steady-state conditions), *literature*, *assumed* (a documented assumption with no series behind it) or *scenario* (a forward-looking input that no data identify) — and its source. The replication package records the provenance of each value in `results/data/parameters.json`; Appendix D documents the series.

8.1. What the data pin

Shares and stocks come from the national accounts: the labor share (PWT, 0.568 in 2023), residential structures (1.17 times GDP) and their depreciation (2.3%), nonresidential depreciation (8.9%), and the investment share of information-processing equipment (2.57% of GDP), which capitalized at $\delta_Q + g_Z + g_{pQ}$ gives a value of compute capital of 6.2% of GDP. The data-center share of electricity (4.4% in 2023) pins the energy intensity of compute. Given these, the model's steady-state conditions imply the rest of the base point: the stock of structures is the residual that makes factor payments add up to GDP (1.71 times GDP, against 1.22 for private nonresidential assets in

Table 4. Parameters of the baseline, with status and source.

Object	Symbol	Value	Status	Source
Labor share of GDP, 2023		0.568	data	PWT via FRED
Nonresidential fixed assets / GDP (comparison)		1.22	data	BEA via FRED
Structures / GDP implied by income shares	K_0^N	1.71	calibrated	residual of the income identity
Residential structures / GDP	K_0^H	1.17	data	BEA via FRED (derived)
Energy capital / GDP	K_0^E	0.30	assumed	assumption
Value of compute capital / GDP	$p_0^Q K_0^Q$	0.062	data-derived	IPE investment capitalized
Depreciation: structures, housing, compute, energy	δ_k	0.089, 0.023, 0.25, 0.05	data / assumed	BEA via FRED; accounting lives
Adjustment costs	ϕ_k	3, 3, 1, 3	literature	Cooper–Haltiwanger (2006) order
Real return in the frozen BGP	r^*	0.035	assumed	between TIPS (1.96%) and returns on wealth
Trend growth per efficiency unit	g_Z	0.015	assumed	real GDP growth 1990–2025 minus population
Curvature of utility	γ	2.0	literature	
Discount factor (implied)	β	0.9954	calibrated	$\beta = (1+g_Z)^\gamma/(1+r^*)$
Housing expenditure share (implied)	θ_H	0.162	calibrated	housing value added 12% of GDP
Land share of housing (implied)	ν_H	0.420	calibrated	housing Euler
Resource share in energy (implied)	η	0.353	calibrated	energy Euler
Energy sector output / GDP		0.040	assumed	assumption (electricity 1.7% + other)
Data-center share of electricity, 2023		0.044	data	Shehabi et al. (2024)
Production land rent / GDP		0.030	assumed	assumption
Elasticities: tasks; tasks–structures; top nest	$\sigma_B, \sigma_N, \sigma_T$	0.5, 0.7, 0.6	literature	Oberfield–Raval; Chirinko
Workers of skill L; skill premium	H_L; w_H/w_L	0.6; 1.7	assumed	
CES task weights (implied)	ν_i	0.262, 0.360, 0.189, 0.189	calibrated	task wage shares
Compute per worker, task 1, 2026 (implied)	$\alpha_{1,0}$	14.9	calibrated	compute Euler and task-1 tie
Implied compute share of task-1 input, 2026		0.2184	implied	
Terminal compute per worker by task	α_i^∞	1, 2, 4, 8	scenario	normalization $\alpha_1^\infty = 1$
2026 compute per worker by task (multiples of terminal)	$\alpha_{i,0}/\alpha_i^\infty$	15, 50, 400, 3000	calibrated / scenario	
Capability progress; compute price decline; efficiency gains	$g_\alpha, g_{pQ}, g_\varepsilon$	0.1, 0.15, 0.1	scenario	Epoch: 0.34 for GPU FLOP/$
Floors (fraction of 2026): compute price; energy per unit	$\underline{p}/p_0, \underline{\varepsilon}/\varepsilon_0$	0.01, 0.01	scenario	

the data — the difference is the model's broader definition, which includes public and intangible capital), the housing land share (0.42) and the energy resource share (0.35) make the housing and energy Euler equations hold, and the CES weights are the base-point cost shares. The price of compute services and the 2026 compute-equivalent of a routine worker ($\alpha_{1,0} = 14.9$) follow from the compute Euler equation and the task-1 tie $A = w_L/\alpha_1$; they imply that compute supplies 22% of the effective input of the routine task in 2026, an untargeted magnitude that we find plausible for the current penetration of software and AI in routine work.

8.2. What the data do not pin

Three groups of parameters are not identified by any historical observation. The elasticities between the task composite and physical complements ($\sigma_N = 0.7$, $\sigma_T = 0.6$) are literature values estimated at compute intensities far below those of the scenarios (Oberfield and Raval, 2021; Chirinko, 2008); every asymptotic statement of Section 7 changes sign at $\sigma_T = 1$, and Section 10 reports $\sigma \in \{0.4, 0.6, 1, 1.5\}$. The technology paths — the rate of capability progress g_α, the rate of decline of the compute price g_{pQ} and its floor, the energy-efficiency path — are scenario inputs. The Epoch AI hardware data imply a doubling of FLOP/$ every two years ($g_{pQ} = 0.34$); the baseline uses 0.15 because the price of *installed* capacity includes buildings, power and cooling that do not follow Moore's law, and treats 0.34 as the fast-Moore scenario. The floors matter most: Proposition 3 says

that the wage in the AGI limit is proportional to the product of the compute-price floor and the terminal compute per worker, so the baseline floor of 10^{-2} (compute capacity one hundred times cheaper than in 2026 once Moore's law ends, around 2057) is the single assumption on which the level of wages after 2060 depends. Finally, the terminal compute requirements $\alpha_i^{\infty} = (1, 2, 4, 8)$ are a normalization plus an ordering: one unit of 2026 capacity, which today does about one fifteenth of a routine worker's work, eventually does the whole of it, and the expert task needs eight times as much. The reader who finds this optimistic about algorithms should read the slow-capability scenario; the reader who finds it pessimistic, the fast one.

8.3. Heterogeneous households

The household block uses a productivity process with persistence 0.95 and innovation standard deviation 0.20, two assumed discount factors (0.93, 0.97) and a third calibrated so that stationary savings at the 2026 prices equal the two-class wealth (0.9885). The implied wealth Gini (0.74), top-10% share (0.57) and bottom-50% share (0.034) compare with 0.85, 0.69 and 0.02 in the Distributional Financial Accounts; the top-1% share (0.12 against 0.325) is under-predicted for the usual reason that the model has no return heterogeneity.

9. The baseline transition, 2026–2100

Every path in this section is a perfect-foresight equilibrium of the two-class economy under the baseline scenario of Table 4: capability progress at 10% per year, a compute price that falls by 15% per year until it reaches one hundredth of its 2026 level (around 2057), energy efficiency gains of 10% per year with a floor, and no policy. It is a conditional scenario, not a forecast; Section 10 shows how each statement changes with the inputs. Figure 2 shows the real side, Figure 3 the balance sheet, Table 6 selected years, and Table C1 in the appendix the residuals and the horizon check (maximum Euler residual $9e-10$; 50 extra years change the first 225 years of the paths by less than 10^{-7}).

9.1. Sequence of automation and the two Restrepo transitions

Tasks are automated in the order of their compute requirements: the routine task uses compute for more than half of its input from 2029, the two cognitive tasks from 2043 and 2045, the expert task from 2057. Automation is an equilibrium decision — it happens when $A_t \alpha_{i,t}$ falls below the wage, not when $\alpha_{i,t}$ reaches its floor — and the dates respond to both the price of compute (which falls with the capital price and rises with energy) and the wage (which rises when the other task of the same skill is automated and labor crowds in).

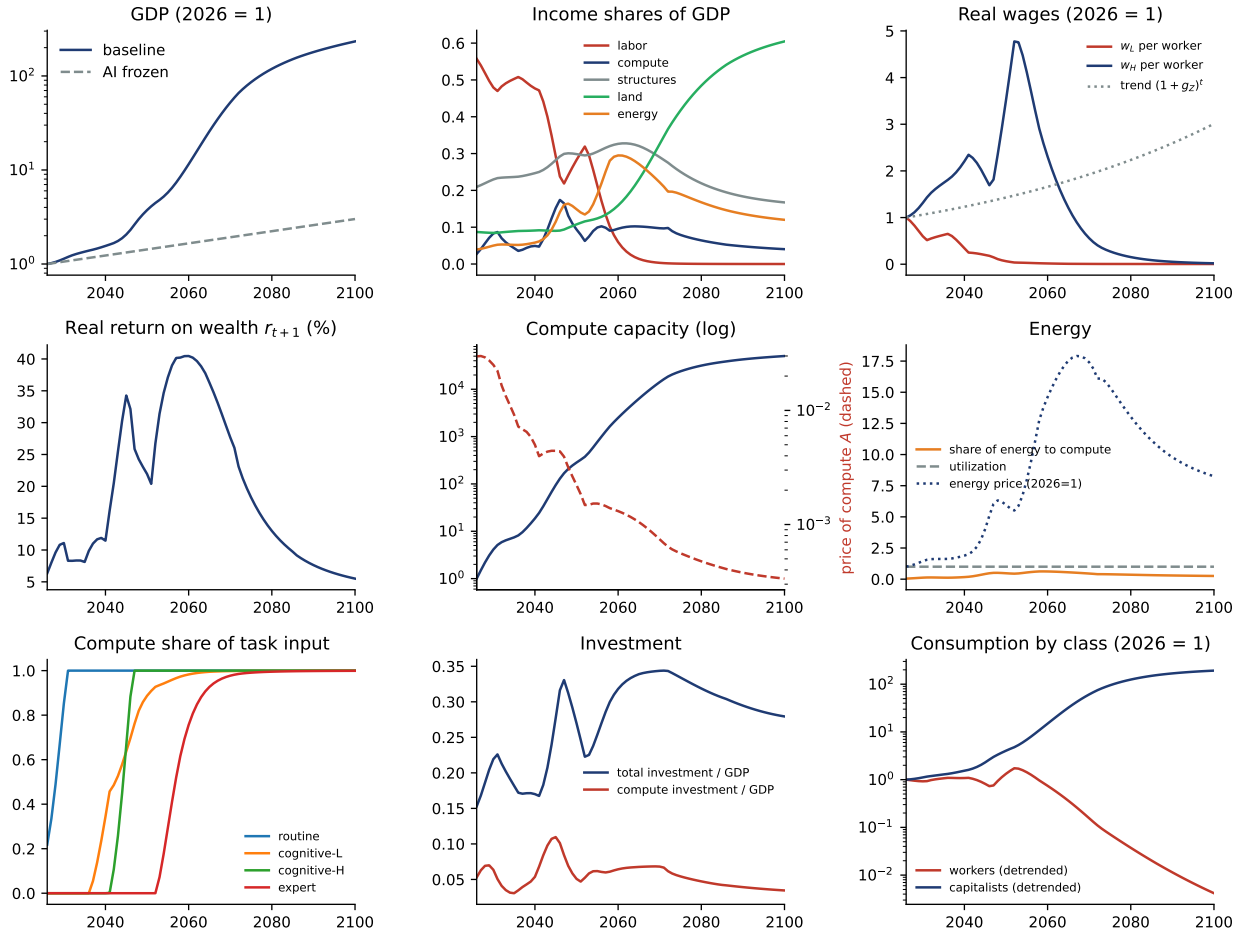


Figure 2. Baseline transition. GDP relative to 2026; gross income shares; real wages per worker relative to 2026 (the dotted line is trend growth); real return on wealth; compute capacity and the price of compute services; energy; the compute share of each task's input; investment; consumption of the two classes detrended by $(1+g_Z)^t$. Source: `results/scenarios/baseline_paths.npz`.

The wage paths display both transitions of the source in sequence. For the lower skill the transition is compute-binding (Restrepo's Proposition 7): the wage falls around each automation date and recovers between them (0.60 of its 2026 level in 2030, 0.34 in 2040, 0.06 in 2050), with no jump. For the higher skill the transition is technology-binding (Proposition 8): while its expert task is the last one that compute cannot do cheaply, the H wage rises to 3.51 times its 2026 level in 2050 — the transitory premium of the skill that "happened to be the last experiment to pan out" (source, p. 13) — and then collapses once the expert task is automated (0.27 in 2075, 0.02 in 2100). The labor share falls from 0.56 to 0.480 in 2030, 0.283 in 2050 and 0.001 in 2075 (Proposition 1); after 2075 it is below one in a thousand. Wages *per efficiency unit* converge to the compute-equivalent user cost of Proposition 3; wages *per worker* then grow again at the trend rate g_Z from a level that is a small fraction of 2026 (Section 10 shows how small depends on the floors).

9.2. Growth, the real rate and the news effect

Output grows by 1.6 times by 2040, 3.7 times by 2050 and 233.0 times by 2100, against three times on the trend; the growth rate peaks at 16.3% in 2062. The boom is a level effect: once compute is cheap, output is bounded by the physical complements and growth returns to g_Z, with land, structures and energy capturing 0.605, 0.17 and 0.12 of income in 2100 and compute 0.040 (Proposition 2c). The real return on wealth rises with the growth of capitalists' expenditure — 11.1% in 2030, 21.9% in 2050, a peak of 40.5% in 2060 — and returns to $r^* = 3.5\%$ by the end of the century. The jump already in 2026 (6.3% against 3.5% on the frozen path) is the perfect-foresight "news" effect: the entire path is anticipated, investment in compute surges (Tobin's q of compute is 1.5 in 2026), and the marginal investor demands a higher return to postpone consumption. The direction — anticipation of AI raises real rates and compute investment today — is a prediction of the structure; the size depends on γ and on perfect foresight.

9.3. The balance sheet: income to capital, value to land

Figure 3 contains the paper's main quantitative message. Between 2026 and the early 2040s the share of income accruing to capital rises and the wealth-to-income ratio *falls*, from 4.1 to 2.7 in 2042: the real rate is high, long-duration claims are discounted heavily (the land price–rent ratio falls from 17 to 12.5, the house price–rent ratio from 12 to 5), and the claim that captures the new income — compute — is capitalized at a price–earnings ratio of about 5.0 in 2026 and 2 in the 2040s, as Proposition 6 predicts for an asset that depreciates at 25% and cheapens at 15% a year. The market value of compute capital never exceeds one fifth of GDP and its share of wealth peaks near 10% around 2045. After 2050 the picture reverses: as the physical complements bind, income flows to land, the real rate normalizes, and the wealth-to-income ratio rises to 15.6 in 2075 and 31.0

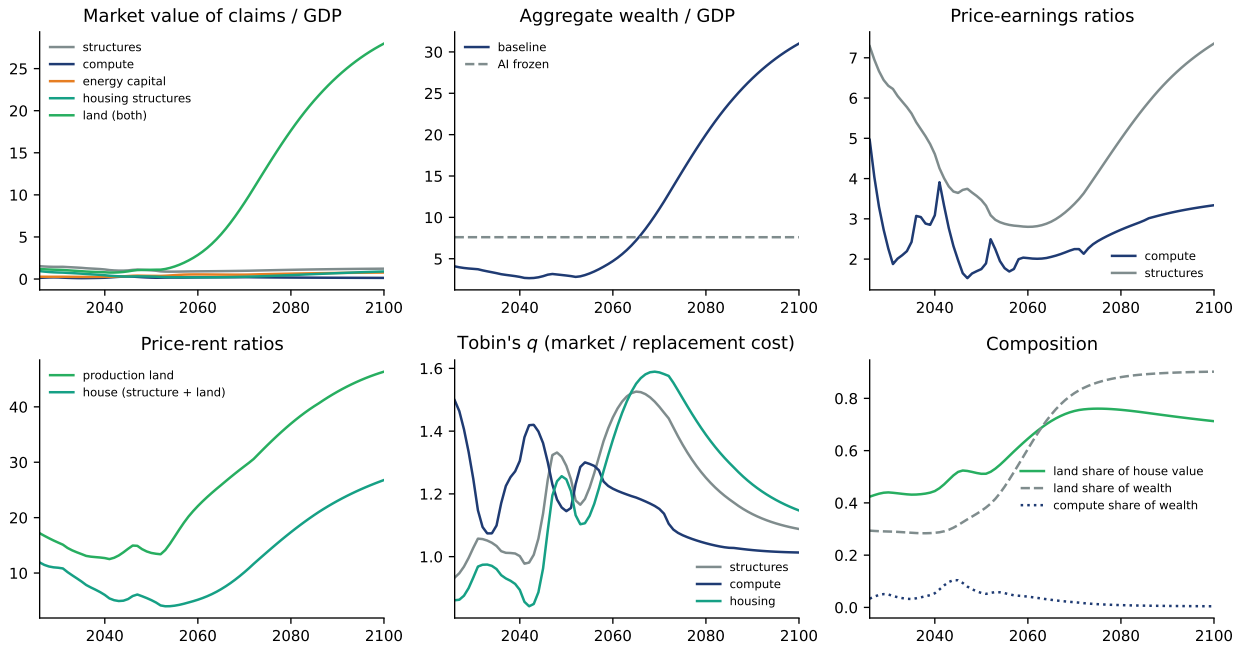


Figure 3. Baseline balance sheet. Market value of each claim relative to GDP; aggregate wealth; price–earnings ratios of compute and structures; price–rent ratios of production land and of houses; Tobin's q relative to replacement cost; composition of house value and of wealth.

in 2100, with land making up 90% of wealth and three quarters of house values. The economy ends up with a Piketty-like balance sheet, but by a route that has nothing to do with $r > g$ for reproducible capital: it is the scarcity of what compute cannot make.

Table 5 decomposes the 2026 value of each claim — baseline versus the frozen economy — into the cash-flow and discount-rate channels (exact under perfect foresight, Section 6). The cash-flow channel is large and positive for every claim, since the boom raises every rent and rental; the discount-rate channel is large and negative for every claim; and the net effect on 2026 *values* is -108% for production land, -20% for structures and 78% for compute. The two channels are of the same order of magnitude, which is why the sign of the valuation response to AI is a quantitative question, not a theorem.

9.4. Energy, utilization and the physical constraint

The share of electricity used by compute rises from 4.4% to a peak of about 60% around 2060 and settles near 25%; the energy price rises eighteen-fold by the late 2060s relative to 2026 and declines afterwards as energy capital catches up. Utilization of compute capacity stays at one throughout the baseline: firms anticipate the energy price and do not build capacity they cannot power. The energy-constrained scenarios of Section 10 are the ones in which the constraint binds on the margin.

Table 5. Decomposition of the 2026 value of each claim, baseline relative to the frozen economy (log points $\times 100$; exact Campbell–Shiller decomposition along deterministic paths).

Claim	Cash-flow channel	Discount-rate channel	Total
Production land	634.7	-743.0	-108.3
Residential land	347.9	-476.6	-128.8
Structures	375.3	-395.1	-19.8
Compute capital	542.5	-464.2	78.3
Housing structures	325.0	-350.7	-25.6
Houses (structure + land)	347.9	-476.6	-128.8

9.5. The two classes

Workers' consumption — their wage bill — rises in level until the early 2050s (the higher skill's premium outweighs the lower skill's losses), then falls to a few percent of its 2026 level by 2100 (Figure 2, bottom right), while capitalists' consumption rises more than a hundredfold. Workers' share of GDP goes from 0.47 to 0.24 in 2050 and essentially zero in 2100. In the two-class economy this is Proposition 8: with no claim on the fund, workers end up with the compute-equivalent of their labor, which is what cheap compute is worth. Section 10 reports the heterogeneous-agent economy, in which workers do own claims, and the policies that change this allocation.

10. Scenarios, heterogeneous households, policy and welfare

Table 7 summarizes the scenario catalogue in 2050 and 2100; Figures 4–8 show the paths. Every scenario re-calibrates the 2026 base point to the same shares, so that differences are differences in the anticipated future, not in the starting balance sheet.

10.1. Technology: how fast, how cheap, how constrained

The *timing* of the transition is governed by the rate of capability progress and by the compute price path; the *destination* by the floors. With capability progress at 20% a year (fast) the labor share is below 0.07 by 2050 and the real return peaks above 50%; at 5% a year (slow) the labor share is still 0.27 in 2050 and the wage of the lower skill is three times higher than in the baseline at that date, but by 2100 the two paths are within a factor of two of each other. The rapid-AGI scenario, in which every task reaches its terminal compute requirement in 2032, is the source's technology-binding transition in its purest form: output rises forty-fold by 2050, the real return exceeds 100% during the 2030s, the labor share is one in a thousand by 2050, and the expert skill's premium never

Table 6. Baseline transition, selected years (two-class economy).

	2026	2030	2040	2050	2075	2100
GDP (2026 = 1)	1.00	1.14	1.56	3.74	85.77	232.96
GDP growth (%)	2.2	4.5	3.5	9.7	7.8	2.3
Labor share	0.557	0.480	0.477	0.283	0.001	0.000
Compute income share	0.027	0.083	0.049	0.094	0.079	0.040
Structures share	0.210	0.230	0.247	0.299	0.250	0.167
Land share (both)	0.087	0.085	0.091	0.108	0.417	0.605
Energy share	0.039	0.050	0.058	0.149	0.190	0.120
Wage L per worker (2026 = 1)	1.000	0.600	0.342	0.065	0.005	0.005
Wage H per worker (2026 = 1)	1.000	1.314	2.206	3.515	0.275	0.020
Real return r (%)	6.3	11.1	11.5	21.9	18.3	5.5
Compute capacity (2026 = 1)	1.0	4.1	18.7	283.0	23,934	50,330
Value of compute / GDP	0.136	0.187	0.152	0.164	0.193	0.135
Wealth / GDP	4.08	3.78	2.81	2.97	15.63	31.01
Land price–rent ratio	17.2	15.5	12.7	13.6	33.1	46.3
House price–rent ratio	11.9	11.0	6.0	5.1	13.9	26.8
P/E compute	5.0	2.3	3.1	1.8	2.4	3.3
P/E structures	7.3	6.3	4.6	3.5	4.1	7.4
Tobin's q compute	1.50	1.25	1.31	1.14	1.07	1.01
Share of energy used by compute	0.044	0.121	0.166	0.474	0.390	0.256
Investment / GDP	0.153	0.219	0.170	0.270	0.334	0.279
Workers' consumption / GDP	0.467	0.402	0.400	0.237	0.001	0.000
Tasks using compute (of 4)	1	1	2	3	4	4

Table 7. Scenario catalogue: selected outcomes in 2050 and 2100 (two-class economy). w_L per worker relative to 2026; GDP relative to 2026; W/GDP aggregate wealth to GDP; land P/R price–rent ratio of production land; V^Q/GDP market value of compute capital; C_W/GDP workers' consumption share.

	2050								2100							
Scenario	Labor share	w_L (2026=1)	GDP (2026=1)	r (%)	W/GDP	Land P/R	V^Q/GDP	C_W/GDP	Labor share	w_L (2026=1)	GDP (2026=1)	r (%)	W/GDP	Land P/R	V^Q/GDP	C_W/GDP
Baseline	0.283	0.065	3.7	21.9	3.0	13.6	0.164	0.237	0.000	0.005	233.0	5.5	31.0	46.3	0.135	0.000
AI frozen (everything at 2026)	0.569	1.431	1.4	3.5	7.6	50.8	0.063	0.475	0.569	3.013	3.0	3.5	7.6	50.8	0.063	0.475
Capability frozen, hardware cheapens	0.494	0.343	2.0	9.1	5.3	47.9	0.168	0.414	0.387	0.149	13.4	3.9	10.1	49.7	0.458	0.324
Rapid AGI (all tasks in 2032)	0.001	0.029	44.7	25.3	12.6	31.0	0.095	0.001	0.000	0.004	243.9	3.8	34.8	49.8	0.132	0.000
Slow capability ($g_\alpha = 0.05$)	0.274	0.213	2.4	29.0	3.4	17.7	0.297	0.229	0.002	0.007	134.1	9.9	18.5	36.3	0.312	0.002
Fast capability ($g_\alpha = 0.20$)	0.066	0.029	7.5	54.2	3.8	20.8	0.169	0.056	0.000	0.004	246.3	4.0	34.1	49.1	0.132	0.000
Terminal compute requirements $\times 5$	0.274	0.071	3.6	21.6	3.1	14.5	0.182	0.229	0.000	0.022	203.5	5.0	29.4	46.6	0.221	0.000
Fast Moore ($g_{pQ} = 0.34$), cheap energy	0.080	0.019	9.7	37.2	4.4	19.3	0.377	0.068	0.000	0.004	244.3	5.0	32.6	47.5	0.136	0.000
Moore's law ends 2040 (floor 0.12)	0.262	0.124	3.5	18.0	3.8	16.5	0.443	0.219	0.000	0.036	169.3	8.2	24.2	42.8	0.430	0.000
Compute price floor 10^{-3}	0.282	0.065	3.7	22.1	2.9	13.0	0.164	0.236	0.000	0.002	244.1	5.3	32.0	46.6	0.021	0.000
Strong complements ($\sigma_N = \sigma_T = 0.4$)	0.137	0.063	3.2	21.7	7.9	44.9	0.121	0.115	0.000	0.003	21.2	4.0	34.8	49.6	0.033	0.000
Scarce land (share 8%, $\sigma_T = 0.4$)	0.193	0.060	3.1	18.0	11.3	33.8	0.149	0.163	0.000	0.003	13.2	3.8	41.5	50.0	0.078	0.000
Severe chip depreciation ($\delta_Q = 0.45$)	0.284	0.077	3.8	22.8	2.9	13.3	0.150	0.238	0.000	0.006	234.8	5.5	31.0	46.3	0.110	0.000
Expert task never automated	0.323	0.057	4.4	14.8	6.3	39.5	0.171	0.271	0.308	0.003	23.3	5.5	10.6	44.0	0.486	0.258
$\gamma = 1$	0.351	0.044	5.0	16.6	13.3	136.0	0.241	0.294	0.000	0.004	245.2	4.2	34.7	50.7	0.139	0.000
$\gamma = 4$	0.217	0.088	2.6	39.2	1.8	4.1	0.150	0.182	0.000	0.006	201.6	11.1	18.8	29.9	0.120	0.000

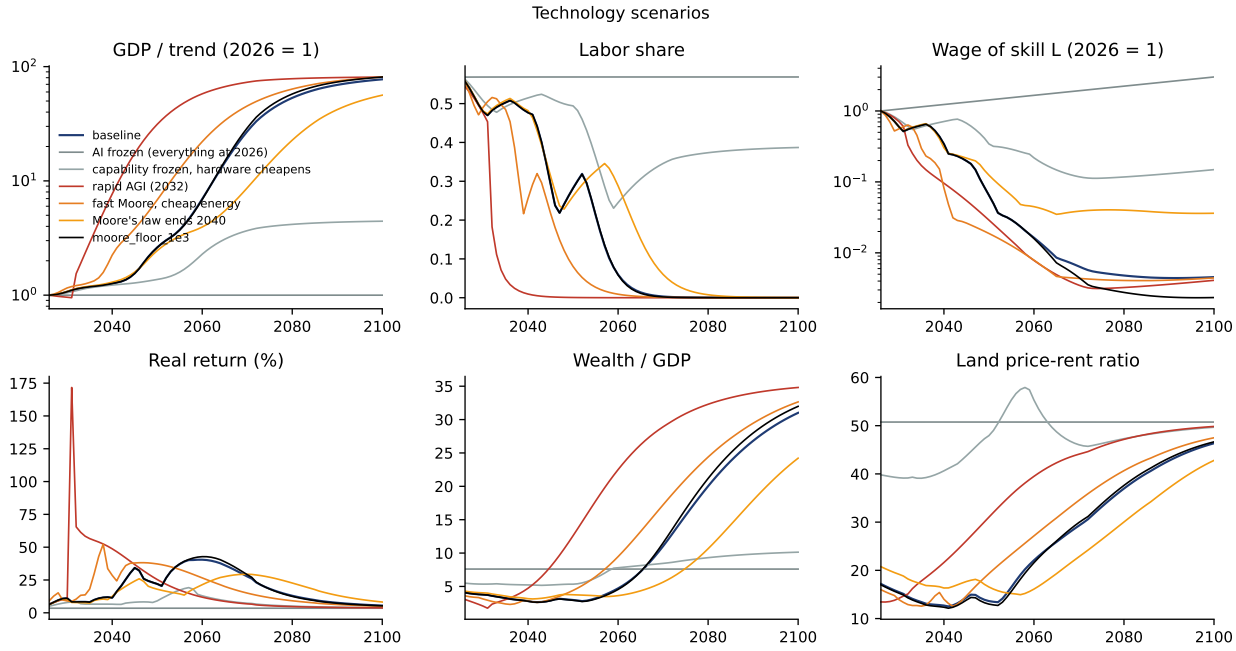


Figure 4. Technology scenarios. GDP relative to trend; labor share; wage of the lower skill; real return; wealth-to-GDP; land price–rent ratio.

materializes because nothing is the last task to be automated. The frozen-capability scenario ("no capability") — 2026 algorithms with hardware that keeps getting cheaper — is instructive: Moore's law alone, at 2026 capability, automates three of the four task classes by 2060 and halves the labor share, because the compute requirements of 2026 are finite and a hundredfold cheaper compute eventually meets them. Only the fully frozen economy ("AI frozen") stays on its balanced path.

The compute price floor is the parameter the long run depends on. With Moore's law ending in 2040 (compute only eight times cheaper than today) the wage of the lower skill in 2100 is 0.036 of its 2026 level against 0.00 in the baseline; with a floor of one thousandth it is lower still, and when the terminal compute requirements are five times higher ("$\alpha^{\infty} \times 5$") wages are five times higher than in the baseline at every date after 2070, exactly as Proposition 3 says. The level of wages in the AGI limit is the product of two unidentified numbers; the paper can only report the dependence.

Energy is where the physical world pushes back, and it is also where the solver reaches its limits. In the baseline the share of electricity used by compute peaks near 60% and the energy price rises eighteen-fold, but capacity is never idle. The scenario that combines the fast compute-price decline with *no* efficiency gains ("energy-constrained") could not be solved to tolerance: the Newton iterations stall on the kinks created by the utilization margin (capacity switches between idle and fully used across periods), and the same happens for the scenario in which Moore's law never ends. For these two cells the paper reports only what Proposition 5 establishes — a bounded energy supply with a floor on energy per unit of compute bounds compute, keeps its price positive and

keeps the labor share positive — and the fact that in every solved scenario the energy price rises by an order of magnitude during the boom. Energy efficiency, not energy capacity, is the margin that matters, because capacity is reproducible and the resource share in energy is modest.

10.2. Substitutability and scarcity

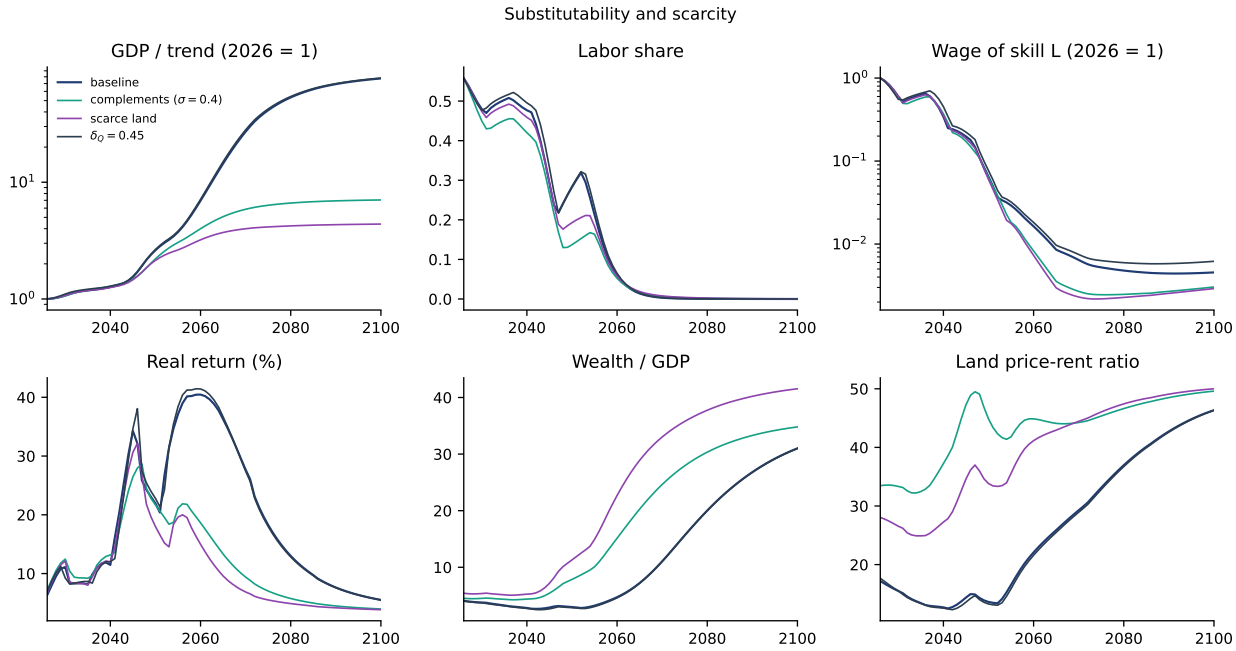


Figure 5. Substitutability between tasks and physical complements, scarce land, and chip depreciation.

The regime map of Section 7 is visible in Figure 5. With strong complementarity ($\sigma_N = \sigma_T = 0.4$) output rises only twenty-fold by 2100 (against 233.0 in the baseline), the land price–rent ratio stays high throughout because the real rate never rises much, and the wealth-to-income ratio does not fall during the boom: there is not enough of a boom to raise the discount rate. The Cobb–Douglas and substitutes cells ($\sigma = 1$ and $\sigma = 1.5$) could not be solved to tolerance on the 275-year horizon — with substitutes the task composite is not bounded by the physical inputs, the economy approaches the source's additive limit, and the compute stock grows by factors that the solver cannot handle together with the terminal conditions — so for those cells the paper relies on Proposition 2 and on the static regime map of Figure E1: constant shares and a wage that falls at the rate of the compute price in the Cobb–Douglas case, a compute share going to one and a positive wage in the substitutes case. Scarce land (a rent share of 8% and $\sigma_T = 0.4$) is the Piketty-like cell: land's share of income exceeds two thirds by 2100 and the wealth-to-income ratio is the highest in the catalogue. Severe chip depreciation ($\delta_Q = 0.45$) raises the user cost of compute by nearly half, raises the wage of the lower skill in 2050 by a fifth relative to the baseline, as Proposition 3 implies, and lowers the

price–earnings ratio of compute further (1.4 against 1.8 in 2050); it does not change the destination of income. The reinstatement scenario, in which the expert task is never automated, keeps the higher skill's premium permanently: its wage rises with output, the labor share converges to the expert task's share of income rather than to zero, and the lower skill's wage path is unchanged. This is the one scenario in which labor as a whole is missed, and it requires that some bottleneck task stay beyond compute forever.

10.3. Preferences and the discount-rate channel

With log utility ($\gamma = 1$) the real return peaks at 26% instead of 40.5%, the land price–rent ratio does not dip during the boom (it starts far higher, because a lower γ makes the anticipated growth path raise the 2026 price of land, and declines slowly as rent growth slows), and the wealth-to-income ratio rises monotonically from 9.5 in 2026; with $\gamma = 4$ the return peaks at 67%, the price–rent ratio of land falls below three in the mid-2050s and the wealth-to-income ratio falls from 3.5 to 1.4 before recovering (Figure 6). The sign of the valuation response to the boom is a preference parameter, as the proposition says; its size in the data would depend on how much of the boom the marginal investor anticipates.



Figure 6. Risk aversion, reinstatement of a human task, and the end of Moore's law.

10.4. Heterogeneous households

Figure 7 and Table 8 report the heterogeneous-agent economy of Section 5 for the baseline, the frozen economy and the citizen fund, solved in general equilibrium up to a maximum log excess demand for claims of 2e-03 (the quasi-Newton on the return path stops improving at that precision; the frozen economy converges to 10^{-4}). Three results stand out.

First, the equilibrium return path of the heterogeneous economy tracks the two-class one closely through the boom (18.9% against 21.9% in 2050, a lower peak) and lies above it afterwards (7.6% against 5.5% in 2100), because the marginal savers who must absorb the growing supply of claims are less patient than the two-class capitalist ($\beta_3 = 0.9885$ against 0.995); the one-year spike in 2026 is the adjustment of the initial distribution to the announced path. The pricing logic of Section 9 carries over — every claim earns the same return, land is where value ends up — with a somewhat higher discount rate. The frozen economy's distribution is calibrated separately (its 2026 wealth is larger) and is flat by construction.

Second, wealth inequality is *non-monotonic*. The Gini coefficient falls from 0.74 in 2026 to 0.60 in 2050 and then rises to 0.73 in 2100 and above its 2026 level afterwards; the top-10% share goes from 0.57 to 0.44 and back to 0.57; the bottom half's share triples by 2050 and then falls below its 2026 value. The simulation separates the two phases. During the boom, the transitory premium of the higher skill and the high return on savings let households in the middle of the distribution accumulate; this is Restrepo's technology-binding transition seen from the wealth side, and it is equalizing. After the last task is automated, labor income disappears for everybody and the distribution of wealth converges to the distribution of patience: households with the highest discount factor hold a growing share, the bottom half holds nothing, and the Gini rises toward the value of a pure rentier economy. The historical analogy is not the twentieth century's $r > g$; it is the disappearance of the one asset that most households own, their labor.

Third, the citizen fund changes what households own, not only what they receive. With the fund holding one fifth of all claims, private wealth is the remaining four fifths, every household receives the same payout, and households at the bottom save less because the transfer substitutes for precautionary wealth: the measured Gini of *private* wealth is slightly higher than in the baseline (0.63 against 0.60 in 2050, 0.76 against 0.73 in 2100) while the share of the fund's payout in the consumption of the bottom half is, by construction, the same as for everybody else. The right comparison is therefore in consumption, which the two-class Table 9 reports: workers' consumption share in 2100 is three orders of magnitude higher with the fund. The top-1% share stays near 0.12 throughout every run: without return heterogeneity the model has no mechanism for a Pareto tail, and the number is reported as a known miss rather than a prediction.



Figure 7. Heterogeneous-agent economy: wealth Gini, top-10%, top-1% and bottom-50% shares, the equilibrium return path against the two-class one, and household expenditure relative to GDP.

Table 8. Wealth distribution along the transition (heterogeneous-agent economy, general equilibrium).

Scenario	Statistic	2026	2030	2040	2050	2075	2100
baseline	Gini	0.738	0.709	0.614	0.601	0.669	0.734
	Top 10%	0.568	0.541	0.460	0.443	0.499	0.569
	Top 1%	0.121	0.114	0.093	0.088	0.102	0.120
	Bottom 50%	0.034	0.048	0.103	0.106	0.067	0.039
	r (%)	21.9	9.3	9.0	18.9	20.2	7.6
citizen fund 20	Top 10%	0.570	0.541	0.462	0.454	0.514	0.582
	Top 1%	0.123	0.116	0.095	0.092	0.107	0.124
	Bottom 50%	0.030	0.043	0.089	0.085	0.041	0.017
	r (%)	17.4	9.1	8.8	18.8	20.3	7.8
AI frozen	Gini	0.798	0.799	0.799	0.798	0.798	0.798
	Top 10%	0.647	0.647	0.647	0.647	0.646	0.646
	Top 1%	0.128	0.128	0.128	0.128	0.128	0.128
	Bottom 50%	0.013	0.013	0.013	0.013	0.013	0.013
	r (%)	3.5	3.5	3.5	3.5	3.5	3.5

10.5. Ownership, competition and policy

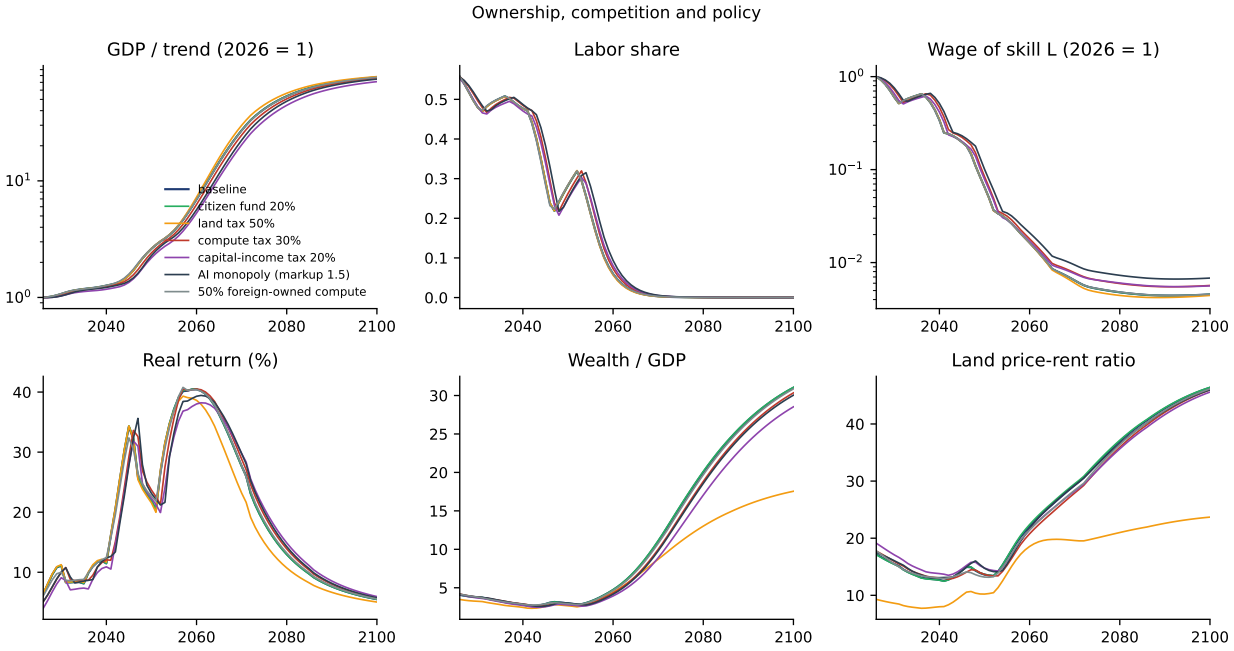


Figure 8. Ownership, competition and policy scenarios.

Table 9. Policy and ownership scenarios: output, wages, workers' consumption and welfare (consumption-equivalent variations over 2026–2065 relative to the baseline, two-class economy).

Scenario	GDP 2050	GDP 2100	w_L 2050	C_W/GDP 2050	C_W/GDP 2100	CEV workers 2026–65 (%)	CEV capitalists 2026–65 (%)
Baseline	3.74	233.0	0.065	0.237	0.0000	0.0	0.0
Citizen fund (20% of equity)	3.73	233.0	0.065	0.270	0.0569	30.7	-9.2
Land tax 50%	3.75	235.8	0.064	0.283	0.2554	40.3	-13.1
Compute tax 30%	3.45	228.4	0.084	0.251	0.0126	25.2	-6.2
Capital-income tax 20%	3.13	213.5	0.079	0.334	0.1676	54.9	-24.2
AI monopoly (markup 1.5)	3.27	224.6	0.101	0.210	0.0000	7.2	-0.9
50% foreign-owned compute	3.75	232.6	0.065	0.238	0.0000	0.5	-1.7

Table 9 compares five instruments with the baseline. A *land tax* of 50% of rents, rebated lump-sum to workers, changes no first-order condition of production (Proposition 9): output in 2050 differs from the baseline by less than half a percent (the difference comes from the change in who saves), and because land captures 0.605 of income in 2100 the rebate gives workers a quarter of GDP in the long run; capitalists lose -13.1% of consumption over 2026–2065, and the wealth-to-income ratio in 2100 is halved because land is worth half as much. A *compute tax* of 30% raises the wage of the lower skill by about thirty percent in 2050 — the tax-inclusive user cost of the replacement compute is higher — and lowers output by eight percent at that date; it is the instrument that changes what workers *earn* rather than what they are *given*. A *capital-income tax* of 20% has the largest output cost in the catalogue (output in 2050 sixteen percent below the baseline)

because it taxes the accumulation of every complement as well as of compute. An *AI monopoly* with a markup of 1.5 on compute services behaves like a compute tax whose proceeds go to the monopolist: the wage of the lower skill is fifty-five percent higher in 2050 than in the baseline and output thirteen percent lower, and workers gain while capitalists, who own the monopolist, are nearly indifferent. *Foreign ownership* of half of compute equity leaves every real path unchanged and sends half of compute payouts abroad; since compute payouts are small relative to the rents that stay at home, the effect on national income is modest, and the result reverses only in the substitutes scenario where compute captures most of income. The *citizen fund* holding one fifth of all claims does what Proposition 8 says: workers' consumption share is bounded below by a fifth of the net payout share, it is three orders of magnitude higher than in the baseline by 2100, and capitalists give up -9.2% of consumption over 2026–2065 for it.

The welfare numbers for workers in Table 9 are large because their baseline consumption falls toward zero: any instrument that gives them a claim on the scarce complements is worth a great deal to them. The comparison that the table allows is across instruments at a given cost to capitalists: per point of capitalists' consumption-equivalent loss, the land tax and the citizen fund deliver far more to workers than the compute or capital-income taxes, because they do not reduce output. This is the general-equilibrium content of the source's suggestion to treat compute "akin to land": in our economy it is land itself, not compute, that should be treated as land.

10.6. AGI-arrival risk: a one-period illustration

The paths above are deterministic. To see the direction in which uncertainty about the arrival of AGI would move asset prices, we take the 2026 stocks as given and let the economy learn in 2027, with probability π, that the baseline path is realized and otherwise the frozen path, pricing each claim with the capitalists' marginal utility in each branch (Section 6). The exercise is not an equilibrium under uncertainty — 2026 decisions ignore the lottery — and the magnitudes are not meaningful, because capitalists' consumption jumps in the baseline branch. The signs are: compute capital carries a large positive premium (it pays in the state in which the marginal investor is rich), land and housing land carry negative premia (they pay relatively more in the state in which the marginal investor is poor, and are therefore hedges), and structures are in between. A full stochastic solution, in which the premium of the last-automated skill is itself a lottery, is the first item on the agenda of Section 12.

11. Empirical discipline and a falsifiable short-run design

The model is calibrated to 2026 and simulated conditionally; nothing in it is a forecast, and the historical sample cannot identify the post-AGI parameters (α_i^{∞}, the floors $\underline{p}$ and $\underline{\varepsilon}$, σ_T at compute

intensities never observed). What the data can do is (i) pin the base point and the pre-AI elasticities, (ii) reject the model's description of the *current* phase of the transition, and (iii) discipline the sign of the mechanisms. This section states what would falsify the model in the next decade and what we have checked against the historical record.

11.1. Validation against pre-AGI patterns

The frozen-technology balanced growth path reproduces by construction the 2026 shares of Table 4. Three untargeted comparisons are informative. First, the model's aggregate wealth-to-GDP ratio in 2026 (4.1) compares with household net worth of 6.0 times GDP in 2026Q2 (Board of Governors of the Federal Reserve System, 2026b): the model's balance sheet, which has no government debt and no foreign assets, is of the right size, with the value of land (1.2 GDP) within the range implied by Davis and Heathcote (2007) and Knoll et al. (2017). Second, the model's price–earnings ratio of structures, $1/(r+\delta_N) = 7.3$, is of the order of the historical earnings yield of equities (Shiller, 2026); the compute price–earnings ratio of 5.0 is what rapid obsolescence and falling prices imply and is not observed separately in national accounts. Third, the heterogeneous economy's wealth Gini (0.74) and top-10% share (0.57) compare with 0.85 and 0.69 in the Distributional Financial Accounts (Board of Governors of the Federal Reserve System, 2026a); the top-1% share (0.12 against 0.325) is under-predicted, as expected without return heterogeneity.

The historical decline of the labor share (Karabarbounis and Neiman, 2014; Autor et al., 2020) is not evidence for or against the model: the model's labor share is flat on the frozen path by construction and falls only through the compute-equivalent channel. What the 2015–2026 record does show, and the model matches at the base point, is that compute income is still a small share of GDP (information-processing-equipment investment of 2.6% of GDP, data-center electricity of 4.4%): the economy is at the beginning of the transition, where the source's Proposition 7 predicts gradual wage effects.

11.2. Short-run predictions that can be checked

Along the baseline path the model predicts, for 2026–2040, a set of co-movements that distinguish it from competing accounts:

1. *Wages by task exposure.* Wages of the skill whose task is being automated fall relative to the other skill around the automation date and recover afterwards (Restrepo's Proposition 7; Figure 2). A design that regresses occupation-level wage growth on the interaction of task exposure (Eloundou et al., 2024; Felten et al., 2021; Webb, 2019) with the time path of compute costs, with occupation and year fixed effects and pre-2022 trends as a placebo, tests the sign and the *timing*: the model predicts that wages fall *after* the compute-equivalent cost of the task crosses the wage, not when

the technology becomes feasible.

2. *Investment and margins by industry exposure.* Industries with high task exposure should show rising compute investment and *falling* margins on the labor-intensive part of their business (compute is competitively supplied), while margins rise only where a scarce complement (sites, grid connections, proprietary data) is owned. The model's distinction between R^Q (competitive) and the markup wedge of the AI-monopoly scenario gives the two hypotheses.
3. *Valuations.* The model predicts that the market value of compute capital per unit of compute income is low (5.0) and that the value of claims to scarce complements rises relative to GDP. A firm-level design would sort firms by the share of their capital that is compute (servers, accelerators; from property, plant and equipment footnotes in annual reports) versus land and long-lived structures, and test whether the price–earnings ratios of the former are lower and whether their earnings yields track the user cost $(1+r)(1+\pi_Q)-(1-\delta_Q)$. Data: Compustat if licensed, or the SEC XBRL financial statements, which are public.
4. *Energy prices.* The energy-constrained scenarios predict a rising relative price of electricity in compute-intensive regions and a rising share of electricity going to data centers (4.4% in 2023, 6.7–12% projected for 2028 by Shehabi et al., 2024). A region-level design using interconnection-queue data and wholesale prices would test the compute–energy price link that drives utilization in the model.

11.3. Identification, confounding and measurement

Each design has the same threats. *Pre-trends*: exposed occupations and industries were already on different trends (offshoring, routine-biased change); the 2000–2021 window provides placebo tests. *Confounding*: AI exposure is correlated with education, tradability and firm size; the model implies that the relevant treatment is the compute-equivalent *cost* crossing the wage, which varies over time within exposure groups and allows occupation-by-year designs. *Measurement*: exposure indices measure feasibility (α finite), not realized automation ($A\alpha < w$); the model predicts that feasibility without adoption has no wage effect, which is itself a test. *Alternative mechanisms*: a demand shock to AI-producing industries raises their margins without any of the task mechanics; the distinction is that demand shocks raise margins of *producers* of compute while task displacement lowers wages of *users*. Appendix D lists the sources, vintages and licenses of every series used, and the document `docs/identification.md` of the replication package states for each parameter whether it is externally identified, calibrated to a base-point share, or a scenario input (Table 4).

12. Conclusion, limitations and open questions

The chapter we build on ends with the sentence "in the AGI economy, we won't be missed." Our results say something more specific about who will be missed by whom. Labor will not be missed by output: that result is robust to everything we tried, because it rests only on the perfect substitutability of labor and compute within tasks. Labor's *income*, however, will be missed by labor, in a way the source does not emphasize: unless compute remains scarce, the compute that replaces a worker is cheap and so is the worker, in level and not only in share. Where compute does not remain scarce, the owners of what compute cannot make — land, sites, the structures that take time to build — are the ones who are not missed. And what those owners hold is valued at the duration of their claims, which is why the wealth-to-income ratio first falls, while the capital share rises and the real rate is high, and then rises when land captures the income.

Three limitations bound these conclusions. First, the paths are perfect-foresight equilibria. There is no aggregate risk, no risk premium and no learning about the arrival of AGI; the one-period lottery of Section 10 shows the direction in which arrival risk would move premia but is not an equilibrium under uncertainty, and the immediate jump of the real rate in 2026 is the model's way of pricing in a path that nobody knows. Second, the level of wages after 2060 depends on the product of two scenario inputs — how cheap installed compute becomes once Moore's law ends, and how much of it a worker-equivalent requires — that no historical data identify; the paper reports the dependence rather than a number. Third, the production structure is fixed. There is no endogenous innovation, no creation of new human tasks beyond the reinstatement scenario, no nonhomothetic demand that would make human-made goods luxuries, no occupational choice or labor supply, and the heterogeneous economy has no return heterogeneity. Each of these would change the distribution of the gains; none of them, as far as we can see, changes the arithmetic of the compute-equivalent wage.

Four questions follow. What does the compute-equivalent wage look like under uncertainty about the arrival date, when the premium of the last-automated skill is a lottery rather than a path? How do the valuation results change when compute services are supplied with market power that is contestable rather than permanent? What is the optimal mix of a land tax, a compute tax and a citizen fund when the planner cannot observe who owns what? And, most directly testable, do the wages of exposed occupations fall when the compute-equivalent cost of their tasks crosses the wage, rather than when the technology becomes feasible? The last question is the one the next decade will answer.

References

Abel, Andrew B. and Olivier J. Blanchard, "An Intertemporal Model of Saving and Investment," *Econometrica*, 1983, *51* (3), 675–692.

Acemoglu, Daron, "The Simple Macroeconomics of AI," *Economic Policy*, 2025, *40* (121), 13–58.

_ **and David Autor**, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in "Handbook of Labor Economics," Vol. 4B, Amsterdam: Elsevier, 2011, pp. 1043–1171.

_ **and Pascual Restrepo**, "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 2018, *108* (6), 1488–1542.

_ **and** _ , "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

_ **and** _ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **and** _ , "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 2022, *90* (5), 1973–2016.

_ **, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones, "Artificial Intelligence and Economic Growth," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, Chicago: University of Chicago Press, 2019, pp. 237–290.

Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 1994, *109* (3), 659–684.

Atkinson, Anthony B., *Inequality: What Can Be Done?*, Cambridge, MA: Harvard University Press, 2015.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Autor, David H., Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Bansal, Ravi and Amir Yaron, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 2004, *59* (4), 1481–1509.

Baqaee, David Rezza and Emmanuel Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

_ and _ , "Productivity and Misallocation in General Equilibrium," *Quarterly Journal of Economics*, 2020, *135* (1), 105–163.

Barkai, Simcha, "Declining Labor and Capital Shares," *Journal of Finance*, 2020, *75* (5), 2421–2463.

Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review*, 1967, *57* (3), 415–426.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu, "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents," *Econometrica*, 2011, *79* (1), 123–157.

Bergholt, Drago, Francesco Furlanetto, and Nicolò Maffei-Faccioli, "The Decline of the Labor Share: New Empirical Evidence," *American Economic Journal: Macroeconomics*, 2022, *14* (3), 163–198.

Bewley, Truman F., "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," in Werner Hildenbrand and Andreu Mas-Colell, eds., *Contributions to Mathematical Economics in Honor of Gérard Debreu*, Amsterdam: North-Holland, 1986.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, 2020, *110* (4), 1104–1144.

Board of Governors of the Federal Reserve System, "Distributional Financial Accounts," Enhanced Financial Accounts 2026. Accessed 2026-10-27.

_ , "Financial Accounts of the United States – Z.1," Statistical release 2026. Accessed 2026-10-27.

Bostrom, Nick, *Superintelligence: Paths, Dangers, Strategies*, Oxford: Oxford University Press, 2014.

Campbell, John Y., "A Variance Decomposition for Stock Returns," *Economic Journal*, 1991, *101* (405), 157–179.

_ **and Robert J. Shiller**, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1988, *1* (3), 195–228.

_ **and** _ , "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, 1988, *43* (3), 661–676.

Carroll, Christopher D., "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 2006, *91* (3), 312–320.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 2017, *8* (3), 977–1020.

Chirinko, Robert S., "σ: The Long and Short of It," *Journal of Macroeconomics*, 2008, *30* (2), 671–686.

Cochrane, John H., "Production-Based Asset Pricing and the Link between Stock Returns and Economic Fluctuations," *Journal of Finance*, 1991, *46* (1), 209–237.

Comin, Diego, Danial Lashkari, and Martí Mestieri, "Structural Change with Long-Run Income and Price Effects," *Econometrica*, 2021, *89* (1), 311–374.

Constantinides, George M. and Darrell Duffie, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 1996, *104* (2), 219–240.

Cooper, Russell W. and John C. Haltiwanger, "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*, 2006, *73* (3), 611–633.

Corrado, Carol, Charles Hulten, and Daniel Sichel, "Intangible Capital and U.S. Economic Growth," *Review of Income and Wealth*, 2009, *55* (3), 661–685.

Costinot, Arnaud and Iván Werning, "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," *Review of Economic Studies*, 2023, *90* (5), 2261–2291.

Crouzet, Nicolas and Janice C. Eberly, "Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles," NBER Working Paper 25869, National Bureau of Economic Research May 2019.

Davis, Morris A. and Jonathan Heathcote, "The Price and Quantity of Residential Land in the United States," *Journal of Monetary Economics*, 2007, *54* (8), 2595–2620.

Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins, "A Model of Secular Stagnation: Theory and Quantitative Evaluation," *American Economic Journal: Macroeconomics*, 2019, *11* (1), 1–48.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs Are GPTs: Labor Market Impact Potential of LLMs," *Science*, 2024, *384* (6702), 1306–1308.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin, "The Decline of the U.S. Labor Share," *Brookings Papers on Economic Activity*, 2013, *2013* (2), 1–63.

Epstein, Larry G. and Stanley E. Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 1989, *57* (4), 937–969.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 2020, *88* (1), 115–170.

Farhi, Emmanuel and François Gourio, "Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia," *Brookings Papers on Economic Activity*, 2018, *2018* (2), 147–250.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, 2015, *105* (10), 3150–3182.

Felten, Edward, Manav Raj, and Robert Seamans, "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal*, 2021, *42* (12), 2195–2217.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "The Dynamics of Inequality," *Econometrica*, 2016, *84* (6), 2071–2111.

George, Henry, *Progress and Poverty: An Inquiry into the Cause of Industrial Depressions and of Increase of Want with Increase of Wealth* Cambridge Library Collection, Cambridge: Cambridge University Press, 2009. First published 1879.

Gomme, Paul, B. Ravikumar, and Peter Rupert, "The Return to Capital and the Business Cycle," *Review of Economic Dynamics*, 2011, *14* (2), 262–278.

Greenwald, Daniel L., Martin Lettau, and Sydney C. Ludvigson, "How the Wealth Was Won: Factor Shares as Market Fundamentals," *Journal of Political Economy*, 2025, *133* (4), 1083–1132.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles, "Should Robots Be Taxed?," *Review of Economic Studies*, 2022, *89* (1), 279–311.

Hanson, Robin, "Economic Growth Given Machine Intelligence," 1998. Unpublished manuscript, University of California, Berkeley. Author's vita dates the paper 1998; the PDF posted on the author's website is dated June 2000 and the paper is often cited as Hanson (2001).

Hayashi, Fumio, "Tobin's Marginal q and Average q: A Neoclassical Interpretation," *Econometrica*, 1982, *50* (1), 213–224.

Heaton, John and Deborah J. Lucas, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 1996, *104* (3), 443–487.

Hémous, David and Morten Olsen, "The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality," *American Economic Journal: Macroeconomics*, 2022, *14* (1), 179–223.

Hobbhahn, Marius and Tamay Besiroglu, "Trends in GPU Price-Performance," Epoch AI June 2022. Accessed 2026-10-27.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith Jr., "Sources of US Wealth Inequality: Past, Present, and Future," *NBER Macroeconomics Annual*, 2021, *35*, 391–455.

Huggett, Mark, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 1993, *17* (5–6), 953–969.

Hulten, Charles R., "Growth Accounting with Intermediate Inputs," *Review of Economic Studies*, 1978, *45* (3), 511–518.

International Energy Agency, "Energy and AI," Technical Report, International Energy Agency, Paris April 2025. Accessed 2026-10-27.

Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, 1995, *103* (4), 759–784.

_ , "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, 2011, *3* (2), 1–28.

_ , "The Past and Future of Economic Growth: A Semi-Endogenous Perspective," *Annual Review of Economics*, 2022, *14*, 125–152.

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor, "The Rate of Return on Everything, 1870–2015," *Quarterly Journal of Economics*, 2019, *134* (3), 1225–1298.

Jr., Robert E. Lucas, "Asset Prices in an Exchange Economy," *Econometrica*, 1978, *46* (6), 1429–1445.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

Kehrig, Matthias and Nicolas Vincent, "The Micro-Level Anatomy of the Labor Share Decline," *Quarterly Journal of Economics*, 2021, *136* (2), 1031–1087.

Knoll, Katharina, Moritz Schularick, and Thomas Steger, "No Price Like Home: Global House Prices, 1870–2012," *American Economic Review*, 2017, *107* (2), 331–353.

Koomey, Jonathan, Stephen Berard, Marla Sanchez, and Henry Wong, "Implications of Historical Trends in the Electrical Efficiency of Computing," *IEEE Annals of the History of Computing*, 2011, *33* (3), 46–54.

Korinek, Anton and Donghyun Suh, "Scenarios for the Transition to AGI," NBER Working Paper 32255, National Bureau of Economic Research March 2024. Also arXiv:2403.12107.

_ **and Joseph E. Stiglitz**, "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, Chicago: University of Chicago Press, 2019, pp. 349–390.

_ **and Megan Juelfs**, "Preparing for the (Non-Existent?) Future of Work," NBER Working Paper 30172, National Bureau of Economic Research June 2022.

Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, *65* (6), 1389–1420.

Krusell, Per and Anthony A. Smith Jr., "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns," *Macroeconomic Dynamics*, 1997, *1* (2), 387–422.

_ **and** _ , "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, *106* (5), 867–896.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins, "Income and Wealth Inequality in America, 1949–2016," *Journal of Political Economy*, 2020, *128* (9), 3469–3519.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

Masanet, Eric, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey, "Recalibrating Global Data Center Energy-Use Estimates," *Science*, 2020, *367* (6481), 984–986.

Meade, James E., *Efficiency, Equality and the Ownership of Property*, London: George Allen & Unwin, 1964. Reprinted 2013 as a Routledge Revival, doi:10.4324/9780203106624.

Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 1985, *15* (2), 145–161.

Mian, Atif, Ludwig Straub, and Amir Sufi, "Indebted Demand," *Quarterly Journal of Economics*, 2021, *136* (4), 2243–2307.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo, "Uneven Growth: Automation's Impact on Income and Wealth Inequality," *Econometrica*, 2022, *90* (6), 2645–2683.

Nordhaus, William D., "Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth," *American Economic Journal: Macroeconomics*, 2021, *13* (1), 299–332.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," *Econometrica*, 2021, *89* (2), 703–732.

Piketty, Thomas, *Capital in the Twenty-First Century*, Cambridge, MA: Harvard University Press, 2014.

_ **and Gabriel Zucman**, "Capital Is Back: Wealth-Income Ratios in Rich Countries 1700–2010," *Quarterly Journal of Economics*, 2014, *129* (3), 1255–1310.

Rachel, Łukasz and Lawrence H. Summers, "On Secular Stagnation in the Industrialized World," *Brookings Papers on Economic Activity*, 2019, *2019* (1), 1–76.

Restrepo, Pascual, "We Won't Be Missed: Work and Growth in the Era of AGI," July 2025. NBER conference paper f227505, Yale University, July 4, 2025.

Rognlie, Matthew, "Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?," *Brookings Papers on Economic Activity*, 2015, *2015* (1), 1–69.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

Segerstrom, Paul S., "Endogenous Growth without Scale Effects," *American Economic Review*, 1998, *88* (5), 1290–1310.

Sevilla, Jaime, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos, "Compute Trends across Three Eras of Machine Learning," in "2022 International Joint Conference on Neural Networks (IJCNN)" IEEE Padua, Italy 2022, pp. 1–8. Also arXiv:2202.05924.

Shehabi, Arman, Sarah J. Smith, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor, "2024 United States Data Center Energy Usage Report," Report LBNL-2001637, Lawrence Berkeley National Laboratory, Berkeley, CA December 2024.

Shiller, Robert J., "Online Data: U.S. Stock Markets 1871–Present and CAPE Ratio; U.S. Home Prices," Yale University 2026. Accessed 2026-10-27.

Smith, Matthew, Owen Zidar, and Eric Zwick, "Top Wealth in America: New Estimates under Heterogeneous Returns," *Quarterly Journal of Economics*, 2023, *138* (1), 515–573.

Statistics Canada, "National Balance Sheet Accounts, Table 36-10-0580-01," Quarterly table (formerly CANSIM 378-0121) 2026. Release of 2026-09-11. Accessed 2026-10-27.

Thuemmel, Uwe, "Optimal Taxation of Robots," *Journal of the European Economic Association*, 2023, *21* (3), 1154–1190.

Tobin, James, "A General Equilibrium Approach to Monetary Theory," *Journal of Money, Credit and Banking*, 1969, *1* (1), 15–29.

Trammell, Philip and Anton Korinek, "Economic Growth under Transformative AI," NBER Working Paper 31815, National Bureau of Economic Research October 2023.

U.S. Bureau of Economic Analysis, "National Income and Product Accounts (NIPA) Tables," Interactive data application 2026. Accessed 2026-10-27.

U.S. Energy Information Administration, "Electric Power Monthly," Monthly statistical release 2026. Accessed 2026-10-27.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," 2019. SSRN Working Paper 3482150.

Zeira, Joseph, "Workers, Machines, and Economic Growth," *Quarterly Journal of Economics*, 1998, *113* (4), 1091–1117.

Appendices

A. Proofs

Throughout, B is the finite set of bottleneck task classes, $\Omega(s)$ the partition across skills, $\widetilde{Q} = Q + \sum_s \widetilde{\mathrm{CEU}}(s) H_s$, and the composite $X_B(\{X_i\}) = (\sum_{i \in B} \nu_i X_i^{\rho_B})^{1/\rho_B}$ with $\rho_B < 0$ is concave, homogeneous of degree one and increasing, with every task essential ($X_B \to 0$ as any $X_i \to 0$).

Proof of Lemma 1

Step 1 (dual of the composite). For task prices $p_i > 0$ the unit cost of the composite is $P_B(p) = [\sum_i \nu_i^{\sigma_B} p_i^{1-\sigma_B}]^{1/(1-\sigma_B)}$ and the cost-minimizing demands are $X_i = \nu_i^{\sigma_B} (p_i/P_B)^{-\sigma_B} X_B$ (Shephard's lemma). If every task is produced with compute at price A per unit, $p_i = A\alpha_i$, so $P_B = A/\mathscr{B}$ with $\mathscr{B}$ as in (10), and the compute required per unit of composite is $\sum_i \alpha_i X_i / X_B = P_B/A = 1/\mathscr{B}$.

Step 2 (pooling). Write $\widetilde{Q}_i = Q_i + \alpha_i L_i$ for the effective compute in task i. The planner's problem $\max X_B$ subject to $\sum_i Q_i \le Q$, $\sum_{i \in \Omega(s)} L_i \le H_s$ is equivalent to $\max_{\{\widetilde{Q}_i\}} X_B(\{\widetilde{Q}_i/\alpha_i\})$ subject to $\sum_i \widetilde{Q}_i \le Q + \sum_i \alpha_i L_i$ for a given labor allocation, and the inner maximum equals $\mathscr{B}(Q + \sum_i \alpha_i L_i)$ by Step 1 and homogeneity. The outer maximum over labor assigns all of H_s to $\arg\max_{i \in \Omega(s)} \alpha_i$, giving $X_B = \mathscr{B}\widetilde{Q}$.

Step 3 (feasibility of the pooled allocation for Q large). The pooled allocation is feasible as an allocation of compute and labor only if, in each task, the effective compute demanded exceeds the labor assigned: $\widetilde{Q}_i^* \ge \alpha_i L_i$. By Step 1, $\widetilde{Q}_i^* = \alpha_i X_i^* = \alpha_i \nu_i^{\sigma_B} (\alpha_i \mathscr{B})^{-\sigma_B} \mathscr{B} \widetilde{Q}$, which is linear in $\widetilde{Q}$, while $\alpha_i L_i \le \alpha_i H_{s(i)}$ is bounded. Hence there is a finite $\bar{Q}$ beyond which the constraint is slack in every task; for $Q \ge \bar{Q}$ the pooled allocation is feasible, hence optimal, and compute is used in every task ($Q_i^* = \widetilde{Q}_i^* - \alpha_i L_i > 0$). For $Q < \bar{Q}$ some constraints bind and the composite is strictly below $\mathscr{B}\widetilde{Q}$ (Restrepo's transition).

Step 4 (prices). Competitive factor prices are the multipliers of the planner's problem. By the envelope theorem $A = Z F_X \, \partial X_B / \partial Q = Z F_X \mathscr{B}$, and the multiplier on H_s is $Z F_X \mathscr{B} \max_{i \in \Omega(s)} \alpha_i = A \widetilde{\mathrm{CEU}}(s)$. □

Proof of Proposition 1

By Lemma 1, $\sum_s w_s H_s = A \sum_s \widetilde{\mathrm{CEU}}(s) H_s$ and $A\widetilde{Q} = Z F_X \mathscr{B} \widetilde{Q} = Z F_X X_B$. Since F is homogeneous of degree one and increasing, Euler's theorem gives $Y = Z F_X X_B + Z F_S S \ge Z F_X X_B = A\widetilde{Q}$. Hence $\sum_s w_s H_s / Y \le A \sum_s \widetilde{\mathrm{CEU}}(s) H_s / (A\widetilde{Q}) = \sum_s \widetilde{\mathrm{CEU}}(s) H_s / \widetilde{Q}$, which goes to zero as $Q \to \infty$ because the numerator is finite under Premise 1. The argument uses only $Y \ge Z F_X X_B$, which holds for any constant-returns F increasing in all arguments; it does not use the form of F, the number of

complements, or how Q is produced. □

Proof of Proposition 2

Write $x \equiv \mathscr{B}\widetilde{Q}$. (a) $\rho \in (0,1)$: $F = x[\nu + (1-\nu)(S/x)^{\rho}]^{1/\rho}$, so $F/x \to \nu^{1/\rho}$ and $F_X = \nu x^{\rho-1} F^{1-\rho} \to \nu \cdot \nu^{(1-\rho)/\rho} = \nu^{1/\rho}$; hence $A \to Z\nu^{1/\rho}\mathscr{B}$ and $s_Q = AQ/Y \to 1$ because $Y/Q \to Z\nu^{1/\rho}\mathscr{B}$ and $\widetilde{Q}/Q \to 1$. (b) $\rho = 0$: $F = x^{\nu} S^{1-\nu}$, $F_X = \nu F/x \to 0$ as $x \to \infty$; the shares are ν and $1-\nu$ identically. (c) $\rho < 0$: $x^{\rho} \to 0$ so $F \to (1-\nu)^{1/\rho} S$ and $F_X = \nu x^{\rho-1} F^{1-\rho} \to 0$ at rate $x^{\rho-1}$; the share of S, $(1-\nu)S^{\rho}/[\nu x^{\rho} + (1-\nu)S^{\rho}]$, converges to one. Wages are $A\widetilde{\mathrm{CEU}}(s)$ in all cases. For the last claim, the pre-AGI wage bill is positive and independent of Q while the post-AGI wage bill is $A\sum_s \widetilde{\mathrm{CEU}}(s) H_s \to 0$ in (b) and (c). □

Proof of Proposition 3

Guess a path on which S_t, K_t^Q, Y_t and capitalists' expenditure grow at a common rate g and $\widetilde{Q}_t/S_t$ is constant (labor's compute-equivalent is negligible asymptotically, so $\widetilde{Q}_t/K_t^Q \to 1$). Constant returns of F give a constant Y/S, hence constant marginal products $R_S = ZF_S$ and $A = ZF_X\mathscr{B}$. The two capital Euler equations without adjustment costs read $1 + r = R_S + 1 - \delta_S$ and $1 + r = (A + (1-\delta_Q)\bar{p})/\bar{p}$, i.e. $A = (r + \delta_Q)\bar{p}$, and the capitalists' Euler equation gives $1 + r = \beta^{-1}(1+g)^{\gamma}$. These three equations in $(r, g, \widetilde{Q}/S)$ have a solution because $R_S(\widetilde{Q}/S)$ is increasing in $\widetilde{Q}/S$ (complementarity) and $A(\widetilde{Q}/S)$ is decreasing, so the ratio is pinned by $A/R_S = (r+\delta_Q)\bar{p}/(r+\delta_S)$, and then g follows from the Euler equation. The investment rates are constant and the goods market clears by Walras' law. The share s_Q is the CES share formula evaluated at the constant ratio; it is strictly between zero and one. The labor share converges to zero by Proposition 1. □

Proof of Proposition 4

(a) With $\rho \in (0,1)$ and S growing at most at the rate of accumulation out of a finite saving share, $F/\widetilde{Q} \to Z\nu^{1/\rho}\mathscr{B}$ as $\widetilde{Q}/S \to \infty$, so A has a positive limit. The compute Euler equation $A_{t+1} = (1 + r_{t+1})p_t^Q - (1-\delta_Q)p_{t+1}^Q$ with $p_t^Q \to 0$ then requires $r_{t+1} \to \infty$, and the Euler equation $1 + r_{t+1} = \beta^{-1}(1+g_{t+1})^{\gamma}$ forces the growth rate of capitalists' expenditure, and hence of output, to diverge. A balanced path with finite g is impossible. (b) With $\rho \le 0$ and $\widetilde{Q}/S \to \infty$ (compute is free at the margin), $F \to (1-\nu)^{1/\rho} S$ for $\rho < 0$ and $F = x^{\nu} S^{1-\nu}$ for $\rho = 0$; in both cases $F_X \to 0$ so $A \to 0$ and $w_s = A\widetilde{\mathrm{CEU}}(s) \to 0$. The value of compute capital is $p_t^Q K_t^Q = A_t K_t^Q/[(1+r)(p_{t-1}^Q/p_t^Q) - (1-\delta_Q)]$ and $A_t K_t^Q/Y_t$ is the compute share, which goes to zero for $\rho < 0$ and is ν for $\rho = 0$; in the Cobb–Douglas case the market value of compute relative to output is $\nu/[(1+r)e^{g_{pQ}} - (1-\delta_Q)]$, bounded, while in the complements case it goes to zero. □

Proof of Proposition 5

Energy feasibility gives $\varepsilon_t Q_t \le \bar{E}$, so $Q_t \le \bar{E}/\bar{\varepsilon}$ and $\widetilde{Q}_t$ is bounded. Since $F_X > 0$ at any bounded $\widetilde{Q}$, $A_t = ZF_X\mathscr{B}$ is bounded below. The labor share equals $s_{Q,t}\sum_s \widetilde{\mathrm{CEU}}(s)H_s/\widetilde{Q}_t$ by Lemma 1 (with $s_{Q,t} = A_t\widetilde{Q}_t/Y_t$ the share of the task composite), and $\widetilde{Q}_t \le \bar{E}/\bar{\varepsilon} + \sum_s \widetilde{\mathrm{CEU}}(s)H_s$ gives the bound. □

Proof of Proposition 6

On a balanced path without adjustment costs the ex-dividend value of capital type k is $V_k = p_k K_k' = p_k K_k(1+g)$ and the Euler equation gives $R_k = (1+r)p_{k,-1} - (1-\delta_k)p_k = p_k[(1+r)(1+\pi_k) - (1-\delta_k)]$ with $p_{k,-1}/p_k = 1+\pi_k$. Hence $V_k/Y = s_k(1+g)/[(1+r)(1+\pi_k) - (1-\delta_k)]$ with $s_k = R_k K_k/Y$. Land: $P_T = \sum_{j\ge 1} R_T(1+g_T)^j/(1+r)^j = R_T(1+g_T)/(r-g_T)$. Summing gives (13). The comparative static follows because the coefficient on s_k is $1/[(1+r)(1+\pi_k) - (1-\delta_k)]$, decreasing in δ_k and in π_k. The price–earnings ratio is $V_k/(R_k K_k)$. □

Proof of Proposition 7

With a constant price index the capitalists' Euler equation gives $1+r = \beta^{-1}(1+g)^\gamma$. For a cash flow growing at g, $P/D = (1+g)/(r-g) = (1+g)/[\beta^{-1}(1+g)^\gamma - (1+g)]$. Let $h(g) = \log(1+g) - \log[\beta^{-1}(1+g)^\gamma - (1+g)]$. Differentiating with respect to $\log(1+g)$,

$$\frac{\mathrm{d}h}{\mathrm{d}\log(1+g)} = 1 - \frac{\gamma\beta^{-1}(1+g)^\gamma - (1+g)}{\beta^{-1}(1+g)^\gamma - (1+g)} = \frac{(1-\gamma)\,\beta^{-1}(1+g)^\gamma}{\beta^{-1}(1+g)^\gamma - (1+g)},$$

whose denominator is $(1+g)(r-g)/(1+g)\cdot(1+g) > 0$ whenever $r > g$ (the transversality condition). The sign is therefore that of $1-\gamma$. □

Proof of Proposition 8

Workers' expenditure is $\sum_s w_s H_s + \tau_t$; with $\tau_t = 0$ and $Y_t \to \infty$ the ratio to Y_t is the labor share, which goes to zero by Proposition 1; the level converges to $A\sum_s \widetilde{\mathrm{CEU}}(s)H_s$ by Lemma 1. With the fund, $\tau_t = \theta_F \mu_W D_t$ where D_t is the aggregate net payout; on the balanced path $D_t/Y_t = \bar{d}$, which gives the bound. □

Proof of Proposition 9

With the tax, the after-tax rental received by the owner of a unit of capacity is $(1-\tau_Q)(A - \varepsilon P_E)$ and the Euler equation on the balanced path is $(1-\tau_Q)(A-\varepsilon P_E) = (r+\delta_Q)\bar{p}$. Solving for A and using $w_s = A\,\widetilde{\mathrm{CEU}}(s)$ gives the expression. A tax on land rent at rate τ_T rebated lump-sum changes

no first-order condition of any producer or investor other than the land-price recursion, in which the after-tax rent replaces the rent; the stock of land is fixed so allocations are unchanged. □

B. Computational methods

All code is in `src/agi_ge/` of the replication package (Python 3.14, NumPy, SciPy, numba); the one-command reproduction is described in Appendix E.

B.1. Static block

Given the predetermined stocks $(K_t^N, K_t^Q, K_t^E, K_t^H)$ and the period's $(\alpha_{i,t}, \varepsilon_t)$, the competitive allocation of labor and compute across tasks is the solution of the concave program $\max_{L,Q\geq 0} Y$ subject to $\sum_{i\in\Omega(s)} L_i = H_s$ and $\sum_i Q_i = Q$, with variables scaled to $[0,1]$ ($l_i = L_i/H_{s(i)}$, $q_i = Q_i/Q$). It is solved by sequential least-squares programming with analytic gradients and then *polished*: the active set identified by the first solver (which tasks use labor, which use compute) is fixed and the Karush–Kuhn–Tucker equalities $F_i = w_{s(i)}$ for tasks with labor, $F_i/\alpha_i = A$ for tasks with compute, and the two adding-up constraints are solved by Newton's method to machine precision. The polished solution is accepted only if no inactive variable has a positive shadow value. Factor prices are the multipliers (envelope theorem): $w_s = \max_{i\in\Omega(s)} F_i$, $A = \max_i F_i/\alpha_i$, $R^N = ZG_V V_K$, $R^T = ZG_T$, $P^E = ZG_E$. Energy supply is predetermined by K_t^E and allocated residually to final production; if the marginal value of compute falls below its energy cost, utilization $u_t < 1$ is found by bisection in $\log u$ on the condition $A_t = \varepsilon_t P_t^E$. The housing rent follows from aggregate goods consumption. A static solve takes 1–3 milliseconds; the Karush–Kuhn–Tucker violation is below 10^{-6} in every test and the income identity (5) holds to 10^{-8}.

B.2. Perfect-foresight transition (two-class economy)

Unknowns are, for $t = 0,\ldots,T-1$, the logarithms of the four stocks K_{t+1}^k, the logarithm of capitalists' expenditure x_t^K, and r_{t+1}: six per period, $6T$ in total. Residuals are the four capital Euler equations (9), the goods-market clearing condition and the capitalists' Euler equation (7). Writing expenditure as an unknown rather than as the goods-market residual keeps the system smooth where investment booms would otherwise drive the residual expenditure through zero. Beyond T the investment rates, the interest rate and the growth rate of expenditure are held at their period-$(T-1)$ values and land prices take their Gordon values. Each residual depends on unknowns at $t-1$, t, $t+1$ only, so the Jacobian is banded; it is computed by colored finite differences (30 evaluations of the full path) and factored sparsely. Newton steps use a backtracking line search. The solution is reached by homotopy in the size of the technology shock ($\lambda \in \{0, 0.25, 0.5, 0.75, 1\}$ scaling

g_{pQ}, g_{ε} and g_{α}), with the base point re-calibrated at each λ so that 2026 stays consistent with the anticipated price path, on the full horizon of 275 years (2026–2300). Convergence is declared when the maximum residual is below 10^{-9}. Truncation is checked by re-solving with 50 additional years and comparing the first 225 years of the stock and interest-rate paths (Table C1).

B.3. Heterogeneous households

Given paths of returns, wages, transfers and the price index, the household problem is solved backward from the stationary policy at the terminal prices by the endogenous grid method on a 300-point asset grid (curvature 4, upper bound 400 times the 2026 average labor income, detrended by the efficiency factor), for the two skills, three discount types and five productivity states. The distribution is simulated forward with the lottery method of non-stochastic simulation. The general-equilibrium path of returns is found by relaxation: the production side is re-solved with the return path given (the same banded Newton method with five unknowns per period), the household block is re-solved, and r_{t+1} is reduced in proportion to the log excess demand for claims at t, until the maximum log excess demand is below 10^{-4}.

B.4. Tests

The test suite (`pytest`, Appendix E) checks: Propositions 1–6 of the source in the benchmark; the Karush–Kuhn–Tucker conditions of the static block; the income identity, and that it fails on a deliberately inconsistent accounting; the exactness of the frozen-technology balanced growth path; the accuracy of the banded Jacobian against a dense one; the Euler-equation residuals, goods-market and asset-market clearing of every saved scenario; and the horizon check.

C. Additional results

C.1. Diagnostics of every scenario

Table C1 reports, for each solved scenario, the maximum absolute residual of the dynamic system (four capital Euler equations, goods-market clearing, capitalists' Euler equation), the maximum gap of the income identity (5) relative to GDP, the truncation check (maximum change in the interest-rate path and in the log capital stocks over the first 225 years when the horizon is extended by 50 years) and the solution time on one core. Four scenarios of the catalogue are not in the table because they were not solved to tolerance after the direct solve and the homotopy fallback: the Cobb–Douglas and substitutes cells (the compute stock grows by factors the terminal conditions cannot accommodate on a 275-year horizon) and the two scenarios in which utilization of compute capacity switches between idle and full across periods (fast Moore without efficiency gains; Moore's

law without end). Their logs are in the replication package; the text relies on the propositions for those cells.

Table C1. Numerical diagnostics by scenario.

Scenario	Converged	max |residual|	max income-identity gap / GDP	horizon check: Δr	$\Delta \log K$	seconds
Baseline	yes	1.7e-12	5.9e-16	5.8e-11	9.3e-10	430
AI frozen (everything at 2026)	yes	1.2e-10	6.7e-16	1.1e-10	9.5e-09	261
Capability frozen, hardware cheapens	yes	5.0e-10	5.6e-16	9.6e-10	1.5e-08	410
Terminal compute requirements $\times 5$	yes	2.1e-12	5.7e-16	9.8e-11	1.5e-09	485
Rapid AGI (all tasks in 2032)	yes	2.5e-12	5.3e-16	1.0e-11	1.6e-10	374
Fast Moore ($g_{pQ} = 0.34$), cheap energy	yes	7.0e-13	4.8e-16	3.4e-11	5.6e-10	419
Moore's law ends 2040 (floor 0.12)	yes	2.0e-10	6.2e-16	1.5e-10	3.5e-09	270
Compute price floor 10^{-3}	yes	6.0e-13	5.7e-16	6.3e-11	9.9e-10	775
Slow capability ($g_\alpha = 0.05$)	yes	2.3e-12	6.7e-16	3.4e-09	5.4e-08	613
Fast capability ($g_\alpha = 0.20$)	yes	2.5e-12	5.5e-16	1.9e-11	3.0e-10	368
Strong complements ($\sigma_N = \sigma_T = 0.4$)	yes	7.5e-13	5.3e-16	5.3e-12	9.4e-11	376
AI monopoly (markup 1.5)	yes	3.8e-12	5.9e-16	6.5e-11	1.1e-09	441
Citizen fund (20% of equity)	yes	1.5e-12	6.3e-16	5.6e-11	8.9e-10	495
Land tax 50%	yes	4.5e-12	7.1e-16	6.0e-10	9.0e-09	495
Compute tax 30%	yes	2.9e-12	6.7e-16	6.8e-11	1.1e-09	443
Capital-income tax 20%	yes	8.9e-10	4.9e-16	1.4e-10	9.2e-09	430
50% foreign-owned compute	yes	1.7e-12	5.9e-16	6.0e-11	9.8e-10	514
$\gamma = 1$	yes	5.2e-13	7.4e-15	3.7e-13	8.7e-12	571
$\gamma = 4$	yes	4.1e-11	5.2e-16	1.9e-08	2.9e-07	505
Scarce land (share 8%, $\sigma_T = 0.4$)	yes	3.1e-10	6.8e-16	5.2e-12	1.9e-09	401
Severe chip depreciation ($\delta_Q = 0.45$)	yes	3.6e-12	6.0e-16	4.7e-11	7.6e-10	498
Expert task never automated	yes	1.2e-11	1.2e-15	1.0e-06	3.4e-05	695

C.2. The two reference economies

The scenario "AI frozen" holds capability, the compute price and energy efficiency at their 2026 values; it is the exact balanced growth path of Section 8 and the reference for the valuation decomposition. The scenario "no capability" holds capability only: hardware keeps cheapening at 15% a year to the same floor. The gap between the two (Table 7) is the contribution of Moore's law alone at 2026 algorithms: three of the four task classes are automated by 2060 and the labor share halves, because the 2026 compute requirements are finite and a hundredfold cheaper compute eventually meets them. The source's Premise 1 (finite α) is already satisfied for every task in 2026 in our calibration — the question is only at what price.

C.3. Price–earnings ratios and user costs

On the baseline path the price–earnings ratio of compute capital is $V_t^Q/(R_t^Q K_t^Q)$; by the Euler equation (9) it equals $(1+g_{t+1}^Q)\,q_t^Q/R_{t+1}^Q$ up to adjustment-cost terms, and on the terminal balanced path $(1+g)/[(1+r)(1+\pi_Q)-(1-\delta_Q)]$ with $\pi_Q = 0$ once the floor is reached: 3.6 at $r = 3.5\%$, $\delta_Q = 0.25$, $g = 1.5\%$. During the transition it is lower (near 2) because the price is still falling ($\pi_Q = 0.15$) and the real rate is high. The corresponding earnings yield, 30–50%, is the number a financial analyst would see as the "return on compute": it is a user cost, not a rate of return, and the return to the owner is r_{t+1} like on every other claim.

C.4. The news effect in 2026

Because the whole path is anticipated, the baseline differs from the frozen economy already in 2026: Tobin's q of compute is 1.5, the real return is 6.3% against 3.5%, and the compute investment share is two to three times its frozen value. These are the model's implications for a world that learns in 2026 that the baseline scenario is certain. The size of the effect is a measure of how much of the scenario the marginal investor would have to believe; the sign — anticipation of AI raises real rates and compute investment at once — is a testable implication (Section 11).

D. Data sources

All series were retrieved on 2026-10-27 from key-less public endpoints by `scripts/01_download_data.py`; every file is listed with its URL, vintage, size and SHA-256 hash in `data/raw/manifest.json`, and every calibration target with its series identifier, period, units and construction in `data/processed/calibration_t` (162 rows). No number in the paper is typed from memory. The FRED copies were obtained through a browser-side fetch because the FRED edge servers timed out for scripted requests on the retrieval day; the files are byte-identical copies and the script re-downloads them when the server answers. Table D2 lists the sources used in the calibration and validation.

Definitions and caveats. The PWT labor share adjusts for self-employment income and is the concept closest to the model's compensation of all workers; the BEA compensation/GDI ratio (0.508 in 2025) is lower because it excludes proprietors' income and includes taxes on production. Depreciation rates divide current-cost depreciation during the year by the end-of-year net stock. The information-processing-equipment investment share is capitalized at $\delta_Q + g_Z + g_{pQ}$ to obtain a stock value; it includes computers and communications equipment beyond AI accelerators, and excludes software, so it is a proxy for the value of "compute capital" in the model's sense. The Shiller spreadsheet ends in September 2023; the Case–Shiller index runs to July 2026. The Epoch price–performance column mixes numerical precisions across generations, which overstates the trend; the paper uses 0.34 only as the "fast Moore" scenario and 0.15 in the baseline. The LBNL data-center figures refer to electricity consumption; the model's energy block is interpreted as electricity.

E. Replication and tests

The replication package (UQO repository `wp24_uqo`) reproduces every number, table and figure of the paper with `make all` (Python 3.14, pinned requirements, TeX Live 2026). The steps and their outputs are:

Table D2. Sources of the empirical inputs.

Object	Series / document (provider)	Vintage used	Value
Labor share	PWT 11.0 `LABSHPUSA156NRUG` via FRED (Groningen; Feenstra et al., 2015); cross-checks: BLS `PRS85006173` index, BEA compensation/GDI and /net domestic income	2023	0.568
Fixed assets	BEA Fixed Assets via FRED: private nonresidential net stock `K1NTOTL1ESO00`, total `K1PTOTL1ESO00`, depreciation `M1NTOTL1ESO00`, `M1PTOTL1ESO00`; GDP `GDPA` (U.S. Bureau of Economic Analysis, 2026)	2024	1.22, 2.39; δ 0.089, 0.057
Compute investment	BEA NIPA via FRED: information-processing equipment `Y034RC1Q027SBEA`, software `B985RC1Q027SBEA`, IPP `Y001RC1Q027SBEA` / GDP	2026Q2	2.57%, 2.49%, 5.77%
Household balance sheet	Federal Reserve Z.1 via FRED: net worth `TNWBSHNO`, real estate `HNOREMQ027S`, owners' equity `OEHRENWBSHNO` (Board of Governors of the Federal Reserve System, 2026b)	2026Q2	6.02, 1.66, 1.10 (× GDP)
Wealth concentration	Distributional Financial Accounts, net worth shares (Board of Governors of the Federal Reserve System, 2026a)	2026Q2	top 1% 32.5, top 10% 68.9, bottom 50% 2.3
Real rates	FRED `DFII10` (10-year TIPS), `REAINTRATREARAT10Y` (Cleveland Fed), `GS10`; Treasury real yield curve	2025 avg; 2026-10	1.96%; 2.47%; 4.99% nominal
Equity valuation	Shiller (2026) `ie_data.xls`: CAPE, D/P, E/P	2023-09 (end of file)	30.8; 1.58%; 4.17%
Housing	Shiller (2026) real home price index; FRED `CSUSHPISA`, `CUSR0000SEHA` (rent CPI)	2023; 2026-07	real growth 1.5%/yr (1990–)
Data-center electricity	Shehabi et al. (2024), LBNL-2001637 (eScholarship copy), pp. 5–6, 52	2023	176 TWh, 4.4% of US electricity; 6.7–12% in 2028
Electricity	EIA Electric Power Monthly Table 1.1 (U.S. Energy Information Administration, 2026); BLS	2025	4 425 TWh generated; 0.185 $/kWh

1. `make test`: the test suite. Static block — Propositions 1–6 of the source hold in the benchmark to 10^{-6}; the Karush–Kuhn–Tucker conditions of the task allocation are satisfied to 10^{-6}; the income identity holds to 10^{-6} relative to GDP and a deliberately broken accounting (compute services counted as final output) is rejected; utilization falls below one when energy binds. Dynamic block — the frozen-technology balanced growth path solves the dynamic system to 10^{-6} without iteration; the banded Jacobian agrees with a dense finite-difference Jacobian to 10^{-6}; a small technology shock converges and satisfies the no-arbitrage identity $(1+r_{t+1})W_t = \text{payouts}_{t+1} + W_{t+1}$ to 10^{-6}; every saved scenario satisfies the residual, identity and no-arbitrage checks.
2. `make benchmark`: Figure 1, Table 1, Figure E1.
3. `make scenarios`: the 22 scenarios of Table C1, each a 275-year perfect-foresight path with its parameter file, meta-data (convergence, residual, horizon check, solution time) and path arrays.
4. `make ha`: the heterogeneous-agent general equilibrium for the baseline, the frozen economy and the citizen fund.
5. `make figures tables paper`: every figure and table of the paper and the file `numbers.tex` of macros, one per number quoted in the text.

Data are downloaded by `make data` from the key-less endpoints listed in Appendix D; the processed calibration targets are versioned with the package so that the paper compiles without network access. No random seeds are used: all solvers are deterministic and the distribution of households is simulated non-stochastically. Solver tolerances are 10^{-14} (static objective), 10^{-9} (maximum residual of the transition) and 10^{-4} (log excess demand for claims in the heterogeneous economy). The research log (`docs/research_log.md`) records the numerical lessons learned during the construction, including three specification changes made for numerical reasons: the removal of accessory tasks from the dynamic model, the treatment of capitalists' expenditure as an unknown, and the re-calibration of the base point for each scenario.

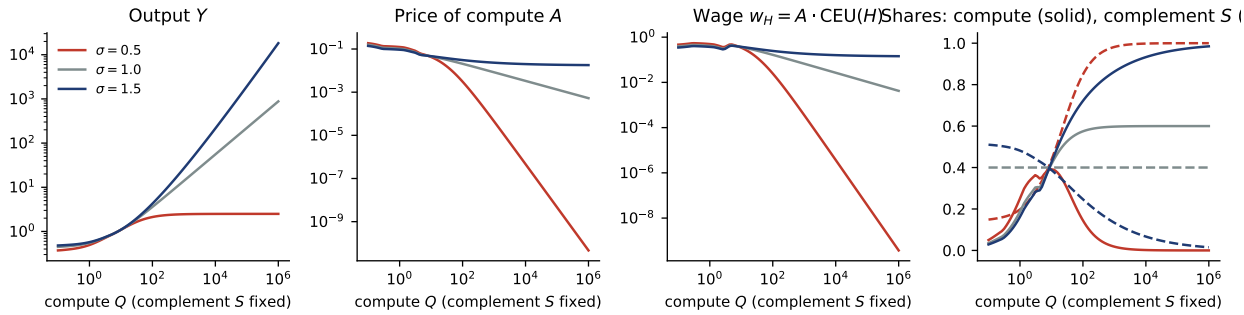


Figure E1. The limit economy of Section 7 computed with the static block: output, price of compute, wage of the higher skill and factor shares as compute grows with the complement S fixed, for $\sigma \in \{0.5, 1, 1.5\}$ (Proposition 2). Source: `scripts/11_regime_map.py`.