

The AI Boom in the US Production Network

Vasco M. Carvalho, Matias Covarrubias, Galo Nuño

The AI boom is an investment boom in data centers and their equipment

- **A macroeconomic shock.** AI investment is 1.75% of US GDP in 2026, and is projected to average 2.43% a year over 2025–32, more than the railroad boom. Van Nieuwerburgh (2026); our construction
- **A larger shock for the industries that supply it.** By 2030 data centers may use 12% of US electricity, up from 2% in 2018, and buy up to 40% of its electrical equipment. LBNL (2024, 2026); Wood Mackenzie (2026)
- **Built at home, filled with imports.** Construction, electrical equipment and utilities build data centers. The servers are mostly imported.
- **Two shocks to the rest of the economy.** A demand shock to a few suppliers travels upstream, a cost shock to everyone downstream. Acemoglu, Akcigit and Kerr (2016)

What must AI deliver to justify this investment, and how does the boom travel?

What we do


1. **Two AI sectors in the US network.** AI facilities and AI compute, built from the engineering recipe of a 2025–26 data center, inside the 2017 production and investment networks of 37 industries.
2. **Demand for AI by task.** Every industry buys AI services for the tasks AI can do, whenever AI does them more cheaply than labour.
3. **News at ChatGPT.** At the end of 2022 agents learn that AI will keep getting more efficient. We size the news to the investment firms plan, and ask what AI must deliver to justify it.
4. **Two accounting identities follow the boom through the network.** We split each industry's change in output by the buyer it serves, and each price change into the costs that the industry and its suppliers pass on.

Results in one slide

1. **Growth.** To justify the investment firms plan, AI must add 2.1 percentage points a year to GDP growth over 2026–40, first through AI productivity, then through the capital that accumulates around it.
2. **The interest rate rises.** Other investment and, for a while, wages make room for AI.
3. **A demand shock to a few suppliers** travels up the supply chain, to metals and minerals. **A cost shock to everyone** travels down it, into prices.
4. **Hours move** from the services AI can do to the industries that build AI.
5. **Scenarios.** Over 2026–2031 the boom adds 2.8 points a year to growth: 3.9 if every announced data center is built, about half a point if AI disappoints. A geopolitical shock to imported servers costs about a third of a point, as much as a regulatory pause, and it lasts.

The model

Two new sectors in the US production and investment network



The 37 sectors of vom Lehn and Winberry in 2017, plus **AI facilities** (shell, power and cooling) and **AI compute** (IT equipment and the labs). We build both from the recipe of a 2025–26 AI data center.

Each task goes to labour or to AI, whichever is cheaper

- In sector j , labour is a bundle of tasks, $\ln T_{jt} = \int_0^1 \ln \tau_{jt}(i) di$. AI can do the share ν_j of them (the **frontier**), with productivity $\psi_t x(i)$, $\ln x(i) \sim N(0, s^2)$.
- The firm uses AI where it is cheaper,

$$\frac{p_{Zt}}{\psi_t x(i)} < W_t \iff \ln x(i) > c_t \equiv \ln \frac{p_{Zt}}{\psi_t W_t}.$$

- AI's share of tasks and the cost of the bundle are

$$a_{jt} = \nu_j \Phi(-c_t/s), \quad \ln P_{T,jt} = \ln W_t - \nu_j E[\max\{0, \ln x - c_t\}].$$

As ψ_t rises, c_t falls. AI takes over more tasks and saves more on those it already does.

► Derivation

- Acemoglu and Restrepo (2022) automate every feasible task, and Acemoglu (2025) fixes the share it pays to automate. Here adoption responds smoothly to the price of AI.

► Relation to Acemoglu–Restrepo

The shock is news that AI will become ever more efficient

- The economy starts in its 2017 steady state. AI can already do the feasible tasks, at a high cost.
- With ChatGPT, at the end of 2022, agents learn the path of AI efficiency,

$$\ln \psi_t = \ln \psi_{\text{pre}} + D w_t,$$

where w_t is logistic. It rises from 0 in 2023 to 1 in the long run, and reaches 0.9 after T_{90} years.

- Three parameters: the level ψ_{pre} , the long-run rise D in log efficiency, and the speed T_{90} .
- Agents have perfect foresight. Investment starts at once, because capacity must be in place before tasks move to AI.

Most of the IT equipment is imported

- AI compute imports 95% of the computers it buys, the import share of US computer investment in 2017.
- The US pays with exports, one basket of US goods in the proportions ξ_j of 2017 consumption for each unit of IT. The terms of trade are fixed.
- Trade is balanced every year, so exports of sector j are $X_{jt} = \xi_j I_t^{\text{imp}}$.
- Imported IT never becomes scarce, and paying for it spreads demand over all industries.
- US households pay for the build-out. A scenario lets the US borrow abroad.



Calibration

Data pin down the network, the recipe and the news

Existing sectors	2017 input-output and capital-flow tables, 37 sectors; parameters of Carvalho, Covarrubias and Nuño (2026).
AI sectors	Inputs, capital and depreciation of a liquid-cooled AI data center of 2025–26.
Frontier ν_j	Share of each industry's tasks that a language model can do (Eloundou et al.), 36% overall, from 18% in food services to 52% in finance.
Dispersion s	An AI–labour elasticity of substitution of 1.5 in 2025.
Preferences	GHH; intertemporal elasticity 0.5, from household evidence (Crump et al.; Havránek).
The news	Level, rise and speed match US AI investment, 0.88% of GDP in 2025, 1.75% in 2026, and 2.43% a year projected for 2025–32.
Check	The revenue of AI providers.

The news matches the investment boom, and revenue follows


— Model ■ Data — Projected average, 2025–32 ● Revenue of AI providers



- Three parameters match AI investment in 2025, in 2026 and its projected average to 2032.
- Revenue is the check. In 2026 the model's AI revenue is 0.36% of GDP, and the revenue of AI providers 0.48%. In both, investment runs years ahead of revenue. [► Adoption](#)

The aggregate boom

Growth comes first from AI, then from capital



- To justify the investment firms plan, AI must add 2.1 points a year to GDP growth over 2026–40.
- To 2030 most of the gain is AI productivity, as firms hand tasks to AI.
- In the 2030s other capital is the largest source, 47% of the gain.
- Capital deepening, as in Imas and Moll. Cheaper tasks are used with more capital.

► Mechanism

► Decomposition


► Long-run arithmetic

Households save for the build-out only at a higher return



- No one borrows. The rate is the return at which households agree to postpone consumption, $r_t \approx \rho + g_t/\text{IES}$, with g_t the growth of the GHH composite.
- With an IES of 0.5, each point of expected growth adds 2 points. The composite grows up to 2.0% a year, and the rate rises 4.3 points to 8.5% in 2029, after a dip to 3.8% in 2025.
- **Real yields.** The curve steepened in data and model (1.1 and 1.5 points, 10y–2y), but the model's level moves too much (3.3 against 1.0, 5y forward). [Details](#)

Other investment and, for a while, wages make room




- With an IES of 0.5, consumption never falls below its 2017 level.
- The investment of existing industries falls 8.4% below its 2017 level in 2026.
- **Workers** make room too. A higher return on capital lowers the wage firms can pay. In 2030 the real wage is 2% below 2017 while GDP is 11.7% higher.
- The labour share falls from 59% to 52% in 2030, and settles at 56%.
- With an IES of 2 consumption pays instead.

► Other elasticities

The boom in other industries

Demand travels up the supply chain




- Imported servers take 66% of AI investment in 2030. The build-out takes more from existing industries through lower investment (2.8% of output in 2026) than it gives through purchases (0.8%).
- For a few industries it is large. Sales to AI add 10.4% to the output of electrical equipment.
- AI buys nothing from primary metals and minerals. They gain by selling to those who build data centers.

► Method

► By industry

► All demand

Other industries invest less where capital lasts longest



- The goods AI buys get cheaper.


► Goods

- Investment falls 1.3 log points more per year of capital life, in 2026, when the one-year rate has barely risen. Two forces: long rates are already up (3.1 points, ten-year), and long-lived capital is the easiest to postpone (House and Shapiro 2008).

- With an IES of 2 the gradient survives, weaker, 3.2 log points per s.d. of capital life against 9.0.

► Regression

The cost of capital travels down the supply chain into prices




- By 2030, real estate +17.3%, utilities +11.8%, petroleum +8.2%; professional services –11.1%, management –12.1%.
- Demand plays almost no role. Electrical equipment sells 31% more, and its price moves 0.0%.
- The cost of capital comes 61% through suppliers. AI's saving pulls the other way. Finance gains 17 points from the first, yet its price falls 7.0%.

Allowed returns and rents pass it on; elsewhere an upper bound (Gormsen and Huber).

► Reading prices

► By industry

Hours move from the services AI can do to the industries that build AI



- Hours rise in electrical equipment and construction, and fall in finance and professional services.
- Correlation with the share of tasks AI can do, -0.66 in 2030.
- Robust to a wealth effect. With separable preferences total hours fall, yet the split by industry correlates 0.99 with the baseline.

► Wealth effect

► Wage and hours

► Cobb–Douglas

Scenarios




Scenarios: how big the boom is, and whether AI delivers

Scenario	Learned	AI investment, % of GDP, 2027–31	GDP growth, % a year, 2026–31	GDP in the long run, % from 2017
Baseline		2.8	2.8	59.4
Transformative AI	2023	3.7	3.9	158.6
Slow diffusion	2027	1.0	0.7	59.4
AI hits a wall	2027	0.5	0.5	6.2
IT import price $\times 1.5$	2027	2.7	2.5 ^a	53.5 ^a
Five-year pause	2030	2.5	2.5	59.4

GDP growth relative to the path without the news. Except for transformative AI, agents follow the baseline until the year learned. Long run: the new steady state. ^aReal income (GDP less the extra exports that pay for the same servers). Online Appendix H.


- The boom adds 2.8 points a year to growth in 2026–2031; 3.9 if every announced data center is built; about half a point if AI disappoints in 2027.
- Geopolitics and regulation cost about a third of a point a year. Bottlenecks move prices, not growth.

Transformative AI doubles the boom



- Every announced US data-center project is built, and the news is recalibrated to the same 2025–26 data. AI must become far more efficient to justify the larger build-out.
- GDP is 68.1% above 2017 by 2040, against 33.3% in the baseline, and 158.6% in the long run. The rate peaks at 11.7%. [▶ Projected build-out](#)

A bust does not reveal its cause



- Agents learn in 2027 that AI disappoints. AI investment falls from 2.5% of GDP to 0.2% with **slow diffusion** and 0.5% with a **wall**. GDP in 2031 is 10.2% or 11.0% below the baseline, and electrical equipment 17.8% or 22.4% below.
- A pause after a security incident in 2030 costs output early, a compliance cost lasts. [In numbers](#)

A geopolitical cost: imported servers half as expensive again

- In 2027 agents learn that imported servers will cost 50% more, for good: export controls or a conflict. The US must export more to buy the same servers.
- AI services are 22.2% more expensive in 2031, firms hand fewer tasks to AI, and AI investment falls from 2.5% to 2.0% of GDP in 2027.
- Growth falls to 2.5 points a year, against 2.8. Real income is 1.4% below the baseline in 2031; GDP, which counts the extra exports, only 0.8%.
- The loss lasts: in the long run real income is 53.5% above 2017, against 59.4% in the baseline. A five-year pause costs as much growth now but nothing in the long run.

If AI disappoints, little capital is stranded

- A wall stops the demand for AI from **growing**. The capacity built is still used, and AI equipment wears out within a few years.
- The value of AI capital falls by at most 0.27% of GDP when the news arrives.
- The economy loses the **future that does not come**. With a wall learned in 2027, GDP ends only 6.2% above 2017 in the long run, against 59.4% in the baseline; with slow diffusion it catches up.

Bottlenecks move rents and prices far more than output

- **Slow data centers.** If they become 10 times slower to build in 2027, the facilities already built are worth 2.23 times their replacement cost. GDP in 2030 is only 0.3% below the baseline.
- **A slow grid** raises the price of electricity by up to 12.3% and leaves GDP almost unchanged.

Conclusion

Conclusion

- So far the AI boom has been an investment boom in data centers and their equipment.
- In the aggregate, growth comes first from AI, then from the capital that accumulates around it, paid for with a higher interest rate. Other investment and, for a while, wages make room.
- Across industries, a demand shock to a few suppliers travels up the supply chain, and a cost shock to everyone travels down it. Hours move to the industries that build AI.
- If AI disappoints, the capacity remains. The loss is the growth that was expected. Geopolitics and regulation cost a third of a point of growth a year.
- The industry patterns hold in every variant we run. The size of the boom depends on the projected build-out, the IES, preferences over hours and the labs' labour. Next, the data of 2023–26 by industry.

Appendix

Model · Mechanisms · Calibration · More results · Other industries
· Scenarios

Model: imports of IT equipment

- In 2017, 95% of computer investment was imported. Imports supply the same share of the computers that AI compute buys.
- The rest of the world sells one unit of IT for one basket of US goods, in the proportions ξ_j of 2017 consumption. The terms of trade are fixed: a unit of imported IT is, by choice of units, what one basket buys.
- Trade is balanced every year, so exports of sector j are

$$X_{jt} = \xi_j I_t^{\text{imp}}.$$

- Tariffs or export controls would make the terms of trade rise over time.
- Imported IT never becomes scarce, and paying for it spreads demand over all industries in proportion to consumption.

Model: the cost of the task bundle

- The cost of one unit of task i is the cheaper of the two:

$$p(i) = \min \left\{ W_t, \frac{p_{Zt}}{\psi_t x(i)} \right\} \Rightarrow \ln p(i) = \ln W_t - \max\{0, \ln x(i) - c_t\}, \quad c_t \equiv \ln \frac{p_{Zt}}{\psi_t W_t}.$$

An automated task saves $\ln x(i) - c_t$ log points; every other task costs W_t .

- With unit elasticity across tasks, the bundle's log price is the average log cost of its tasks:

$$\ln P_{T,jt} = \int_0^1 \ln p(i) di = \ln W_t - \nu_j \underbrace{E[\max\{0, \ln x - c_t\}]}_{\text{expected saving on a feasible task}}.$$

- With $\ln x \sim N(0, s^2)$ the expected saving has a closed form:

$$\ln P_{T,jt} = \ln W_t - \nu_j [s \phi(c_t/s) - c_t \Phi(-c_t/s)].$$

- Every task gets the same share of spending, so AI's share of the bundle's cost is the share of tasks it does:

$$a_{jt} = \frac{p_{Zt} Z_{jt}}{P_{T,jt} T_{jt}} = \nu_j \Phi(-c_t/s) = \frac{\partial \ln P_{T,jt}}{\partial \ln p_{Zt}}.$$

The last equality is Shephard's lemma: the tasks that switch were at the threshold, where AI saved nothing.

Model: the elasticity of substitution between AI and labour

- Definition, at a given AI efficiency: $\sigma_{ZL} = -\frac{d \ln(Z/L)}{d \ln(p_Z/W)}$.
- Spending on AI is $p_Z Z = a P_T T$ and on labour $WL = (1 - a) P_T T$. Dividing, $\ln \frac{Z}{L} = \ln \frac{a}{1 - a} - \ln \frac{p_Z}{W}$.
- A rise in $\ln(p_Z/W)$ raises the threshold c one for one, and $a = \nu \Phi(-c/s)$ falls. Differentiating,

$$\sigma_{ZL} = \underbrace{1}_{\text{intensive margin}} + \underbrace{\frac{\phi(c/s)}{s \Phi(-c/s)} \cdot \frac{1}{1 - a}}_{\text{extensive margin}}.$$

- **Intensive margin.** If no task switches, spending shares are fixed, so quantities move one for one with the price: an elasticity of 1, as in Cobb–Douglas.
- **Extensive margin.** Tasks switch from labour to AI. $\phi(c/s)/s$ is the density of tasks at the threshold, relative to the mass $\Phi(-c/s)$ already automated. The dispersion s sets how many tasks sit near the threshold.
- The elasticity is not a constant. AI does the easiest tasks first, so as c falls fewer tasks sit at the margin relative to those automated, and σ_{ZL} falls: 1.5 in 2025, 1.26 in the long run.

Model: relation to Acemoglu and Restrepo

- In Acemoglu and Restrepo (2018), AI can do the tasks up to a frontier, and does them where it is cheaper than labour. Two regimes:
 - **Technology-constrained.** AI is cheaper in every feasible task: all are automated, many at a large saving. Small price changes do not move adoption.
 - **Cost-constrained.** Firms stop at the task where AI and labour cost the same, which saves nothing (“so-so automation”). The frontier does not bind.
- Moving between the regimes is a kink in adoption, and a transition that crosses it is hard to solve. Later work assumes the first regime: Acemoglu and Restrepo (2022) automate every feasible task, and Acemoglu (2025) fixes the share of exposed tasks that it pays to automate.
- **We add dispersion** in AI's productivity across feasible tasks, $\ln x(i) \sim N(0, s^2)$. Some feasible tasks are always too costly, so the frontier never binds sharply. Both regimes coexist: the frontier caps adoption, the marginal task is so-so, and the inframarginal tasks save a lot. Adoption $\nu_j \Phi(-c_t/s)$ is smooth in prices, so the transition solves with Newton's method, and s sets the AI-labour elasticity. [▶ Elasticity](#)

Mechanism 1: capital deepening

- Each industry combines capital and the task bundle, $Y_{jt} = \text{CES}(K_{jt}, T_{jt})$.
- AI makes the task bundle cheaper per unit of efficiency. Firms use more of it, and each unit of capital becomes more productive: its rental rises above its user cost $\tilde{P}_{jt}^k(r_t + \delta_j)$.
- Industries invest until the rental is back at the user cost. Adjustment costs, and the investment network that has to build the capital, make this take decades: other capital is 20% higher by 2040 and 60% higher in the long run.
- This capital, not AI capital, produces most of the long-run gain: 26.1 of the 59.4 points, against 6.6 for AI capital.
- **Through the network.** With the rate back at its 2017 level, capital is a produced input priced at its rental, so the long-run wage gain is each industry's AI saving times its capital-inclusive Domar weight, $\lambda = (I - \hat{A}^R)^{-1}c$: 24.8 log points, 1.84 times the AI-productivity term (one sector: 1.80). Hours and the fall in the labour share add the rest of GDP.

Mechanism 2: why the wage rises when AI replaces labour

- The price of the task bundle is the wage net of the saving from AI:

$$\ln P_{T,jt} = \ln W_t - \nu_j E[\max\{0, \ln x - c_t\}].$$

- Each price equals unit cost, $P_{jt} = c_j(P_{jt}^M, R_{jt}, P_{T,jt})$. In the long run capital earns its user cost and the interest rate is back at its 2017 level, so R_{jt} is close to its 2017 level relative to the price of output.
- Then $P_{T,jt}$ must be close to it too (the factor-price frontier). The saving from AI goes to labour: the bundle is 27 log points cheaper at a given wage, and the real wage rises by 30%.
- Capital is the key. It accumulates until it earns its usual return, so labour, the factor in fixed supply, gets the gain. The wage rises only as capital builds up: -2% by 2030.
- AI still displaces labour: the labour share of GDP falls from 59% to 56%.
- **Hours.** With GHH preferences hours depend only on the wage, $L \propto W^{0.5}$: 14.2% more in the long run. With a wealth effect on hours (King, Plosser and Rebelo) they would rise much less.

Mechanism 3: why the interest rate rises, and by how much

- No one borrows. The rate is the return at which households agree to postpone consumption, $1 + r_t = \beta^{-1}(u_{t+1}/u_t)^{1/\epsilon}$, so $r_t - \rho \approx g_{u,t}/\epsilon$, with u_t the GHH composite of consumption and leisure and $\epsilon = 0.5$. Each point of expected growth adds 2 points.
- Two forces raise it. **Investment**: the build-out and capital deepening raise the return on capital and need saving. **Wealth**: households expect to be richer and want to consume part of it now. Consumption cannot rise now, so the rate rises until households are willing to wait.
- **How much it rises** is set by ϵ and checked against long real yields:
 - $\epsilon = 0.5$, the household evidence: about 0.5 from planned consumption (Crump et al. 2022); 0.3–0.4 for asset holders in Havránek's (2015) meta-analysis. Asset pricing needs more: 2.18 in Bansal, Kiku and Yaron (2016). Without risk only ϵ matters.
 - Nov 2022 to 2026. The real curve (ten-year minus two-year) steepened 1.1 point in the data and 1.5 in the model. The five-year forward rose 1.0 point; the model gives 3.3 with $\epsilon = 0.5$ and 1.4 with $\epsilon = 2$. Yields also carry policy and term premia.
 - With $\epsilon = 2$ the rate peaks at 6.2% instead of 8.5%.
 - Foreign saving, calibrated to the current account, lowers the peak to 7.5%.

Calibration: parameters

Parameter	Value	Source or target
<i>Existing sectors</i>		
Elasticities $\sigma_q, \sigma_y, \sigma_m, \sigma_l, \sigma_c$	0.5, 0.6, 0.001, 0.5, 0.5	Carvalho et al. (2026)
Discount factor β	0.96	Carvalho et al. (2026)
Frisch elasticity	0.5	Carvalho et al. (2026)
Intertemporal elasticity ϵ	0.5	Household evidence; Carvalho et al.: 0.33
Capital adjustment cost κ	2	Scenario; Carvalho et al.: 4.25
Labour adjustment cost κ_L	1	Scenario
Imports of AI compute's computers	95%	BEA import share, 2017
<i>AI sectors, per \$1 of AI services</i>		
Compute: labour, IT capital services, space	0.151, 0.589, 0.195	Engineering recipe, A3
Facilities: labour, capital services, electricity	0.004, 0.128, 0.046	Engineering recipe, A3
Capital: IT, facilities	1.94, 1.07	Engineering recipe, A3
Depreciation: IT, facilities	0.26, 0.078	Engineering recipe, A3
Adjustment cost κ_{AI}	1	Scenario
<i>Tasks and the news</i>		
Frontier ν_j , share of tasks AI can perform	β_j ; 0.36 overall	Eloundou et al. (2024)
Frontier of the AI sectors	That of Information	NAICS 518 is in 51
Dispersion of AI productivity s	5.62	AI-labour elasticity 1.5 in 2025
AI efficiency before the news, $\ln \psi_{\text{pre}}$	-16.8	AI investment, 2025–26
Long-run rise of AI efficiency, D	12.0	AI investment, 2025–32
Years to 90% of the rise, T_{90}	6.9	AI investment, 2025–26

Calibration: three parameters, three targets

- AI investment = hyperscaler capex $\times$ generative-AI share (82–85%, Morgan Stanley) $\times$ US share (72%) + AI cloud firms (CoreWeave, xAI). Projection: analysts' path to 2028, then a constant share of GDP.
- Each parameter tilts investment differently over time:

Target	Calibrated value	Percent change after raising		
		$\ln \psi_{\text{pre}} + 0.1$	$D + 0.2$	$T_{90} + 0.5$ years
AI investment 2025, % of GDP	0.88	3.4	3.1	−21.7
AI investment 2026, % of GDP	1.75	2.8	4.1	−17.1
AI investment 2025–32, % of GDP	2.43	2.2	4.2	−5.4
Calibrated parameter		−16.78	12.04	6.86

- The level raises early investment most, the size later investment, and the speed shifts it between them.
- Not targeted: investment in 2024, 0.31% of GDP (data 0.47%).

Calibration: the news, year by year


- $T_{90} = 6.9$ years (90% by 2030) and $D = 12.0$ log points.
- Feasible tasks automated: 0.2% before the news, 0.7% in 2025, 16% in 2030, 23% in the long run.

	2023	2024	2025	2026	2027	2028	2030
Rise of AI efficiency ψ_t , $\times$	1.0	2.6	4.5	7.4	9.0	7.0	2.5
Fall of the price per efficiency unit, $\times$	1.0	2.7	4.8	7.7	8.4	6.3	2.4
AI investment, % of GDP	0.10	0.31	0.88	1.75	2.47	2.85	2.93
Data	0.16	0.47	0.88	1.75			
AI revenue, % of GDP	0.04	0.06	0.14	0.36	0.85	1.61	3.06
Direct revenue (check)	0.01	0.02	0.09	0.48			
Korinek–McKelvey	0.10	0.23	0.54				

Epoch AI measure a fall of $13\times$ a year in the cost of a fixed benchmark score over 2023–26. With an elasticity of 1.5, the calibrated path implies that a share $\theta = 0.58$ of it reaches task efficiency.


The news: what it implies for AI efficiency

— Model Epoch: benchmark cost ---- Hardware $\times$ algorithms (Gundlach et al.)



- AI efficiency rises fastest in 2027, $9.0\times$ in a year, and is 90% of the way by 2030. In the long run it is about 169,000 times its level before the news.
- Over 2023–26 it rises $4.4\times$ a year, close to hardware and algorithms together, $3.9\times$ (Gundlach et al.), and below the fall in the cost of a fixed benchmark score, $13\times$ (Epoch AI). After 2025 it is faster.

Adoption: AI's share of the task wage bill



- **Task wage bill:** what firms pay for tasks, to workers or to AI. Every task gets the same spending (unit elasticity), so the share of tasks AI does, $\nu_j \Phi(-c_t/s)$, is its share of the bill.
- In the long run firms automate 23% of the tasks AI can do, so AI does 8.2% of the task wage bill (13.9% with transformative AI).
- Firms automate about the same share of the feasible tasks in every industry, so industries differ through ν_j : from 4.1% in food services to 11.9% in finance.

Calibration: what the price data contain

	× a year
Epoch AI (2026): cost of a fixed benchmark score, 2023–26	13
Korinek and McKelvey (2026): price of inference, quality-adjusted, 2024–25	16
Gundlach et al. (2025): price at fixed performance, 2024–25	5–10
of which hardware price-performance	1.3
of which algorithmic efficiency	3
of which margins and competition	1.3–2.6
Technology: hardware × algorithms	3.9
Model: rise of AI efficiency ψ_t , 2023–26 average	4.4
as a share of the Epoch rate, in logs (θ)	0.58
as a share of the technology rate, in logs	1.09

- Only technology raises AI efficiency in tasks. Lower margins lower the price but are not efficiency.
- The model's AI efficiency, pinned by the take-off of investment with an elasticity of 1.5, rises at about the technology rate over 2023–26 (1.09 of it in logs).
- Other gaps: benchmarks are not tasks; the model's compute has 2017 hardware; users shift to frontier models as old capabilities get cheap.

Calibration: news about efficiency, not about the frontier

	Data	Moving frontier sized to build-out	Moving frontier sized to 2025 investment	Moving frontier free speed	Moving efficiency (baseline)	Both
AI investment 2024, % of GDP	0.47	1.21	0.65	0.91	0.31	0.28
AI investment 2025, % of GDP	0.88	1.63	0.88	1.21	0.88	0.88
AI investment 2026, % of GDP	1.75	2.00	1.09	1.53	1.75	1.75
AI investment 2025–32, % of GDP	2.43	2.43	1.33	2.43	2.43	2.43
AI revenue 2025, % of GDP	0.09	0.56	0.30	0.46	0.14	0.12
AI revenue 2026, % of GDP	0.48	0.97	0.53	0.76	0.36	0.34
Fall of the price per efficiency unit, $\times/\text{yr}$	13	1.0	1.0	1.0	4.6	4.2
Free parameters		1	1	2	3	3
Years to 90% of the news		10.0	10.0	20.0	6.9	5.9
Pass-through θ		–	–	–	0.58	0.55
AI capital, long run, % of GDP		14.2	7.5	29.8	15.0	15.8
GDP 2030, %		2.7	1.4	2.9	11.7	10.1
GDP long run, %		11.0	5.6	25.7	59.4	52.6
from AI productivity		1.4	0.8	3.1	17.3	14.9
Interest rate, peak, %		4.9	4.6	5.3	8.5	7.5


- Bold: targeted; as many free parameters as targets. A moving frontier sized to the projection puts investment too early (1.6% of GDP in 2025, against 0.88%), even with a free speed.
- A moving efficiency fits the take-off and the projection, and its revenue follows the revenue of AI providers. Moving the frontier as well gives similar results.

Calibration: choosing the targets

	Data	Baseline	Trans- formative AI	Van Nieuw. average	Direct revenue	Earlier targets	Required revenue
AI investment 2023	0.16	0.10	0.18	–	0.03	0.41	0.20
AI investment 2024	0.47	0.31	0.40	–	0.34	1.37	0.43
AI investment 2025	0.88	0.88	0.88	–	1.49	2.82	0.88
AI investment 2026	1.75	1.75	1.75	–	2.48	3.79	1.75
AI investment 2025–32		2.43	3.13	–	2.43	3.63	3.50
AI revenue 2023		0.036	0.090	–	0.004	0.100	0.111
AI revenue 2025	0.087	0.139	0.234	–	0.087	0.540	0.272
AI revenue 2026	0.48	0.36	0.47	–	0.48	1.42	0.51
AI revenue 2030		3.1	3.8	–	3.2	4.9	4.3
AI revenue 2032		3.8	5.1	–	3.5	5.2	6.2
Long-run rise D		12.0	12.7	–	14.9	11.5	15.0
T_{90} , years		6.9	8.4	–	5.0	5.2	9.7
GDP 2030, %		11.7	14.1	–	13.8	26.7	15.6
GDP long run, %		59.4	158.6	–	45.3	101.0	566.8
Interest rate, peak, %		8.5	11.7	–	8.6	11.6	17.6

- Bold: targeted. Targeting revenue (Korinek–McKelvey and Van Nieuwerburgh) puts investment at 2.8% of GDP in 2025, against 0.88%: Korinek–McKelvey value the same installed compute at rental prices.
- The projection decides the size: transformative AI (3.13% instead of 2.43%) gives GDP 158.6% in the long run instead of 59.4%.

More results: AI revenue, investment and the price of AI



- AI investment is 0.9% of GDP in 2025 and 2.9% in 2030. It settles at 2.9%, mostly to replace IT that depreciates at 0.26 a year. AI revenue reaches 3.1% of GDP in 2030 and 4.9% in the long run.
- The price of compute first falls, as capacity is built ahead of adoption, then rises by up to 13.8% in 2030. The price per efficiency unit falls by 5.2 orders of magnitude.

More results: sources of GDP growth, the decomposition

Factors earn their marginal products, so nominal GDP is factor income, $\sum_j (R_{jt}K_{jt} + W_{jt}^*L_{jt})$. Year to year,

$$\Delta \ln GDP_t \approx \frac{1}{\overline{GDP}_t^n} \sum_j \left[\underbrace{\overline{P_{T,j}} \overline{T_j} g_{jt}}_{\text{AI productivity}} + \underbrace{\bar{u}_{jt} \Delta K_{jt}}_{\text{capital}} + \underbrace{\bar{W}_t \Delta L_{jt}}_{\text{hours}} + \underbrace{(\bar{R}_{jt} - \bar{u}_{jt}) \Delta K_{jt} + (\bar{W}_{jt}^* - \bar{W}_t) \Delta L_{jt}}_{\text{scarcity wedges}} \right]$$

- **AI productivity** (Hulten): $g_{jt} = (1 - \bar{a}_{jt}) \Delta \ln W_{jt}^* + \bar{a}_{jt} \Delta \ln p_{Zt} - \Delta \ln P_{T,jt}$, the fall in the cost of the task bundle at given prices of labour and AI, weighted by spending on tasks over GDP.
- **Capital** at its user cost $u_{jt} = \tilde{P}_{jt}^k(r_t + \delta_j)$, split into the two AI sectors and the rest. **Hours** at the wage W_t .
- **Scarcity wedges**: rental above user cost, value of an hour above the wage. The sectoral-demand terms of the price decomposition.
- Bars: averages of $t - 1$ and t . Levels, not logs, because AI capital starts from zero. Added up from 2017 and scaled to the % change in GDP.

More results: the price decomposition

- With constant returns and competition, each industry's price moves only when a cost per unit of output moves. For any industry j , dropping the index j :

$$d \ln P = \underbrace{s_M d \ln P^M}_{\text{upstream inputs}} + \underbrace{s_Y s_K d \ln R}_{\text{rental on capital}} + \underbrace{s_Y (1 - s_K) d \ln P_T}_{\text{wage and AI}}.$$

- We split the rental on capital R into three parts:

$$\ln R = \underbrace{\ln \tilde{P}^k}_{\text{upstream inputs}} + \underbrace{\ln(r + \delta)}_{\text{cost of capital}} + \underbrace{\ln \frac{R}{\tilde{P}^k(r + \delta)}}_{\text{sectoral demand}}.$$




- Without adjustment costs sectoral demand is mute: the supply of capital is infinitely elastic at its user cost. The same holds for labour, whose premium over the economy's wage we add to sectoral demand.
- Capital takes a year to install, so demand in the year of the news moves prices whatever the adjustment cost.

More results: why sectoral demand mirrors the cost of capital

- The capital parts of the two terms add up to the change in the rental over the price of capital goods, $\Delta \ln(R_{jt}/\tilde{P}_{jt}^k)$, times capital's share of cost.
- In the short run capital is fixed, so the rental is set by the demand for the industry's capital services, not by the interest rate.
- When r_t moves, the cost of capital moves at once. The rental catches up only as capital adjusts, and the gap shows up as sectoral demand with the opposite sign. Owners of the installed capital take the gap as a capital gain or loss (Tobin's q).
- The rate first stays near its 2017 level (2023–25), then rises to its peak in 2029. So sectoral demand first rises, then falls. When the rate comes down, the rental again lags, and sectoral demand turns positive.
- The gap is largest where capital is long-lived and a large share of cost: real estate and utilities. Their sum, the rental, moves smoothly.

More results: over time, these prices follow the cost of capital


— Price - - - Cost of capital - . - . Sectoral demand — Upstream inputs - - - Wage and AI



- Early on, sectoral demand mirrors the cost of capital: capital is fixed, so its rental lags behind the interest rate. [► Why](#)
- Electrical equipment: tasks are a large part of its cost, and the saving from AI arrives decades before the wage catches up.

More results: the suppliers of AI capital sell much more, at almost the same price

— Output, % from 2017 - - - Price, % from 2017



- Once their capital and labour adjust, supply is elastic: the extra demand raises output, not price.

More results: downstream, price rises travel with fuel

- The intermediate bundle whose price rises most is that of Petroleum products, by 9.1% in 2037, when crude spikes with the news. Next come Education and Utilities.
- Oil and gas extraction peaks at 12.3% in 2037, 0.9 points of it from sectoral demand: the news moves spending from consumption to investment at once.
- The investment bundle whose price rises most is that of Oil and gas extraction, by 8.3%.
- The suppliers of AI capital are small shares of other industries' inputs, and their prices barely move.

More results: adjustment costs

	AI invest- ment, % of projection	GDP long run	Electrical equipment			Utilities		
			Peak price	Sectoral demand	Long-run price	Peak price	Sectoral demand	Long-run price
Baseline, $\kappa = 2$, $\kappa_L = 1$	100	59.4	0.4	-0.1	-0.7	14.3	-2.9	1.3
<i>Capital</i>								
Fast, $\kappa = 1$	101	59.4	0.1	0.0	-0.7	14.1	-2.0	1.3
Slow, $\kappa = 4$	99	59.4	1.1	-0.1	-0.7	14.7	-4.5	1.3
<i>Labour</i>								
Free, $\kappa_L = 0$	100	59.4	0.0	0.1	-0.7	14.2	-2.4	1.3
Fixed by sector, $\kappa_L = \infty$	94	40.6	12.3	6.6	4.0	23.4	-4.4	7.8
<i>Grid</i>								
Slow, $\kappa_{Util} = 20$	100	59.4	0.9	0.3	-0.7	19.2	-5.5	1.3

- Same news in every row. Sectoral demand is its average part of the price, 2023–30.
- Capital. Halving or doubling κ barely changes aggregates: GDP in 2030 is 11.5 or 12.0 percent above 2017, against 11.7.
- Labour. Free labour absorbs demand at one wage. Fixed hours are a permanent friction and change long-run prices.
- Grid. A grid ten times slower to build, $\kappa_{Util} = 20$, raises the price of utilities by up to 19.2% in 2043.


More results: other intertemporal elasticities

Intertemporal elasticity ϵ	Interest rate peak rise, points	Consump- tion, lowest	GDP		Utilities		Real estate	
			90% of 2030	gain in	Peak price	Cost of capital	Peak price	Cost of capital
0.33	5.02	2.0	9.0	2078	15.8	19.1	24.3	27.1
0.5 (baseline)	4.30	0.5	11.7	2067	14.3	16.3	21.4	23.7
1	3.10	-2.2	16.0	2055	12.1	11.7	17.0	17.0
2	2.05	-4.7	19.6	2048	10.4	8.3	13.5	11.2

- Same news. The household's Euler equation sets the rate, $1 + r_t = \beta^{-1}(u_{t+1}/u_t)^{1/\epsilon}$, with u_t consumption net of the disutility of work. Baseline $\epsilon = 0.5$; Carvalho, Covarrubias and Nuño (2026) has $\epsilon = 0.33$.
- With $\epsilon = 0.33$ the rate rises by 5.02 points instead of 4.30, and utilities and real estate peak at 15.8% and 24.3%. GDP reaches nine-tenths of its gain in 2078 instead of 2067.
- With $\epsilon = 2$ (asset pricing) the rate rises by 2.05 points, the boom comes sooner, and consumption pays for the build-out.

More results: a higher intertemporal elasticity, and transformative AI

— Baseline - - - Intertemporal elasticity 1 - - - Transformative AI



- Elasticity 1, same news: the rate peaks at 7.3% instead of 8.5%, and the prices of long-lived capital rise less.
- Transformative AI: more than twice the long-run gain in GDP, and a rate that peaks at 11.7%.

► How much the rate rises

More results: is the boom reasonable?

	Baseline	IES 0.33	IES 1	IES 2	Both	Moving frontier
GDP growth 2026–2031, % a year	2.8	2.5	3.4	3.8	2.4	0.4
GDP growth 2026–2036, % a year	2.4	2.1	2.8	3.0	2.1	0.4
GDP growth 2026–2056, % a year	1.3	1.2	1.4	1.4	1.2	0.3
Consumption growth 2026–2031, % a year	1.2	0.9	1.8	2.3	1.0	0.0
Consumption growth 2026–2036, % a year	1.4	1.1	2.0	2.5	1.2	0.2
Consumption growth 2026–2056, % a year	1.1	1.0	1.3	1.5	1.0	0.2
Fastest GDP growth, % (year)	3.6 (2030)	3.2 (2030)	4.2 (2029)	4.6 (2029)	2.9 (2029)	0.5 (2030)
Interest rate, peak, %	8.5	9.2	7.3	6.2	7.5	4.9
AI share of task cost, 2030 / long run, %	6 / 8	5 / 8	6 / 8	6 / 8	5 / 9	5 / 8
Feasible tasks automated, 2030 / long run, %	16 / 23	15 / 23	16 / 23	16 / 23	12 / 15	91 / 97
Saving per automated task, 2030 / long run, %	77 / 80	77 / 80	77 / 80	78 / 80	75 / 76	23 / 29
Fall in task-bundle cost, 2030 / long run, log pts	16 / 27	16 / 27	17 / 27	17 / 27	14 / 23	1 / 3
Real wage, long run, %	30	30	30	30	26	3
Hours, long run, %	14	14	14	14	12	1
Labour share of GDP, 2017 / long run, %	59 / 56	59 / 56	59 / 56	59 / 56	59 / 55	59 / 56

- Growth is above trend (the model has none): fast in the next five years, plausible over thirty.
- AI does few tasks, at a fraction of the cost: the task bundle becomes 27 log points cheaper. The real wage rises 30%, and GHH hours follow it.

Other industries: two identities split output and prices

- Every year, in nominal flows, market clearing $G_i = \sum_j F_{ij}^M + \sum_j F_{ij}^I + C_i + X_i$ and zero profits $G_j = \sum_i F_{ij}^M + z_j + W_j^* L_j + R_j K_j$ hold **whatever the elasticities**.
- **Output by the buyer it serves.** $G_t = (I - A_t)^{-1} \sum_k f_{kt}$, with A_t that year's coefficients. The change since 2017 splits exactly into AI direct, AI through suppliers, exports, consumption, other investment and the change in coefficients.
- **Prices by the costs passed on.** $d \ln P = \bar{S}' d \ln P + e$, solved through the network into the industry's own costs and those its suppliers pass on, each split into the cost of capital, the wage and AI's saving.
- We read the solution of the full model, so every number refers to the same 2017 economy. AI is about a millionth of GDP in 2017, so moving AI within the network, or making one industry slow to build, leaves 2017 unchanged.

Other industries: what each industry's output serves

Year	Industry	AI			Consumption	Other investment	Change in		Real output
		Direct	Indirect	Exports			coefficients	Total	
2026	Electrical equipment	10.4	1.8	0.7	0.2	-5.7	0.5	7.9	8.2
	Construction	3.2	0.0	0.2	0.2	-14.3	-0.4	-11.2	-8.0
	Utilities	0.5	0.4	1.2	0.7	-1.2	0.7	2.3	1.5
	Primary metals	0.0	2.1	0.7	0.0	-5.0	-0.2	-2.5	-1.1
	Machinery	2.3	0.5	0.4	0.0	-6.2	-0.3	-3.3	-1.3
	Professional services	0.4	0.4	0.7	0.1	-4.0	-0.7	-3.1	-1.3
	Finance	0.0	0.2	1.2	0.5	-0.8	0.0	1.1	1.4
	All existing	0.4	0.4	1.0	0.4	-2.8	0.0	-0.7	
2030	Electrical equipment	14.3	2.7	1.4	2.3	9.6	0.9	31.3	31.3
	Construction	4.7	0.1	0.3	1.1	13.6	-0.5	19.3	23.0
	Utilities	4.7	0.9	2.9	7.4	1.7	8.2	25.9	12.0
	Primary metals	0.0	3.1	1.5	2.5	7.6	3.7	18.4	16.1
	Machinery	3.1	0.8	0.8	1.2	13.2	-0.3	18.9	21.3
	Professional services	0.8	0.7	1.2	1.5	5.5	-5.5	4.3	16.5
	Finance	0.0	0.4	2.4	2.6	1.2	-3.4	3.4	10.8
	All existing	0.8	0.7	2.1	3.4	3.8	0.1	10.8	
Long run	Electrical equipment	12.5	2.7	2.3	26.4	31.1	-0.4	74.6	75.8
	Construction	4.8	0.2	0.6	6.3	55.8	1.0	68.6	61.7
	Utilities	10.8	1.4	4.2	47.3	6.8	1.7	72.1	69.9
	Primary metals	0.0	3.1	2.5	27.8	29.0	3.3	65.7	63.4
	Machinery	2.8	0.9	1.3	15.2	43.9	0.1	64.2	64.6
	Professional services	1.4	1.1	2.1	25.3	28.8	-4.2	54.5	64.3
	Finance	0.1	0.7	4.0	45.6	5.0	-4.1	51.3	60.9
	All existing	1.2	0.9	3.4	38.4	16.1	-0.3	59.7	
Long run, closed	Electrical equipment	12.5	4.9	4.6	52.4		0.2	74.6	75.8
	Construction	4.8	2.1	4.8	52.9		4.1	68.6	61.7
	Utilities	10.8	1.7	4.7	53.1		1.8	72.1	69.9
	Primary metals	0.0	4.5	4.7	52.7		3.8	65.7	63.4
	Machinery	2.8	3.3	4.6	52.6		0.8	64.2	64.6
	Professional services	1.4	3.0	4.4	51.6		-5.9	54.5	64.3
	Finance	0.1	0.9	4.3	50.1		-4.1	51.3	60.9
	All existing	1.2	1.7	4.6	52.5		-0.3	59.7	

Change in gross output since 2017, % of 2017 output. Exports pay for imported IT under balanced trade; they fall instead when the US borrows abroad. Long run, closed: other investment as replacement.

Other industries: what explains the fall in other investment

Characteristic	2026	2030
Capital share	6.0	3.4
Life of capital	-9.0	-3.7
Overlap with AI's basket	-0.1	0.4
AI network exposure	6.0	6.6
AI frontier ν_j	-4.4	-5.2
Mean	-10.1	17.6
R^2	0.87	0.67

OLS across existing industries of the change in real investment since 2017 (log points $\times 100$) on standardized characteristics fixed before the news. Overlap: the industry's 2017 investment bundle weighted by the share of each good's sales the AI sectors take in 2026.

Other industries: where AI lands, in full

AI frontier	GDP 2030 vs base	GDP long run, %	Multiple of AI prod.	AI prod. term	Rate peak, %	Lowest other invest., %	AI invest. 2025–32	Hours long run	Welfare, % of cons.	Network multiple
Observed		59.4	3.44	17.3	8.47	−8.4	2.43	14.2		1.84
Uniform	0.14	59.7	3.43	17.4	8.45	−8.5	2.46	14.3	0.12	1.85
Most amplified	0.12	63.4	3.60	17.6	8.13	−9.1	2.48	15.2	0.21	1.98
Least amplified	0.14	55.2	3.26	16.9	8.92	−7.9	2.40	13.3	−0.11	1.73
Consumer-facing	0.03	47.5	2.96	16.0	9.64	−7.3	2.25	11.7	−0.74	1.58

The same news. Network multiple: capital-inclusive over ordinary weights applied to the AI saving.
Welfare: permanent change in consumption from the news on.

Other industries: what survives across cases

Case	AI demand, 2026		Other investment, 2026				Prices 2030, R^2		Capital multiple	
	Multiplier	% of output	Mean	Life	Overlap	Max. ΔP^k	Own	Network	Network	One sector
Baseline	1.90	0.8	-10.1	-9.0	-0.1	-0.7	0.62	0.87	1.84	1.80
No imported IT equipment	1.57	1.5	-12.2	-11.6	0.8	-0.8	0.61	0.84	1.83	1.80
Foreign saving	1.89	0.8	-4.2	-7.2	-0.8	0.6	0.61	0.86	1.82	1.80
IES 2	1.89	1.0	15.5	-3.2	-3.5	4.3	0.41	0.71	1.78	1.80
Wealth effect on hours	1.89	0.8	-4.2	-6.2	-0.6	0.2	0.58	0.84	1.81	1.80
Hours fixed by sector	1.89	0.8	-8.4	-6.5	-0.1	1.3	0.29	0.54	1.85	1.81
Slow grid	1.89	0.8	-10.2	-9.6	3.4	-0.7	0.59	0.85	1.84	1.80
Transformative AI	1.89	0.8	-16.4	-13.1	0.4	-0.8	0.66	0.89	1.99	1.90

AI demand: output per \$ AI buys, and AI's direct and indirect demand in % of output (2026). Other investment: mean change, coefficients on life and overlap, largest bundle price change. Prices: R^2 on own and network exposure. Capital multiple: network and one sector.

Other industries: the boom industry by industry

Industry	Output, 2026				2030	
	AI sectors		Other invest.	Total	Hours	Rate via suppliers
	Direct	Indirect				
Electrical eq.	10.4	1.8	−5.7	7.9	16.4	12.4
Construction	3.2	0.0	−14.3	−11.2	11.2	11.3
Utilities	0.5	0.4	−1.2	2.3	9.1	15.4
Primary metals	0.0	2.1	−5.0	−2.5	9.4	17.4
Machinery	2.3	0.5	−6.2	−3.3	7.8	12.6
Prof. services	0.4	0.4	−4.0	−3.1	−7.0	8.6
Finance	0.0	0.2	−0.8	1.1	−10.2	12.8
All existing	0.4	0.4	−2.8	−0.7		

Output: change since 2017 in % of 2017 output, split exactly by the demand it serves (AI: 2026 coefficients; other investment: 2017 coefficients). Hours: log points $\times 100$. Rate via suppliers: the part of the 2030 price change due to the cost of capital of the industry's suppliers. Online Appendix F.

- Each \$1 the AI sectors spend on existing industries raises their gross output by \$1.90, because their suppliers buy from others in turn. Other investment takes 2.8% of output away in 2026.
- AI sits [upstream](#) (2.87; existing industries 1.02–3.30).

Other industries: the goods AI buys get cheaper

Good	Share of AI's basket, %	2024			2026			2030		
		To AI	To others	Price	To AI	To others	Price	To AI	To others	Price
Construction	37	0.7	-8.1	-1.9	3.8	-13.3	-3.5	5.2	19.0	-3.0
Electrical equipment	20	13.6	-2.3	-2.1	71.2	-9.2	-0.3	97.5	27.4	0.0
Machinery	13	0.8	-4.6	-1.6	4.3	-6.2	-2.1	5.9	21.6	-2.0
Wholesale trade	10	0.7	-5.2	-0.7	4.2	-8.6	-1.2	8.8	17.0	-3.3
Professional services	10	0.2	-3.8	-0.8	1.1	-5.1	-1.9	2.0	22.7	-10.5
Computers	7	0.6	-4.1	-1.6	3.9	-7.2	-2.5	8.2	15.6	-4.9
Transportation and warehousing	2	0.9	-5.8	-0.3	6.0	-9.3	-0.5	12.7	15.0	0.1
Fabricated metals	1	1.4	-4.6	-1.2	7.5	-7.4	-1.9	10.2	19.4	-1.0

Real sales as investment goods, change since 2017 in % of the good's 2017 investment sales, to the AI sectors and to everyone else; price in % from 2017.

- Other buyers buy less of every good AI buys, and these goods get **cheaper**. In 2026 every industry's investment bundle costs 0.7–3.1% less. With constant returns, supply is elastic.
- Competition for goods appears when everyone invests at once. With an IES of 2 other investment rises and investment bundles cost up to 4.3% more. The pattern by life of capital holds in every case we run.



Other industries: prices split into own costs and suppliers' costs

Industry	Price	Cost of capital		Sectoral demand		Wage		AI saving	
		Own	Suppliers	Own	Suppliers	Own	Suppliers	Own	Suppliers
Electrical equipment	0.0	5.2	12.4	1.5	-0.9	-0.5	-1.2	-4.4	-12.1
Construction	-3.0	1.7	11.3	0.6	-1.3	-0.7	-1.0	-4.3	-9.3
Utilities	11.8	18.7	15.4	-3.4	-2.0	-0.3	-1.3	-2.8	-12.5
Primary metals	2.0	3.8	17.4	0.0	-2.2	-0.2	-1.4	-1.5	-13.8
Machinery	-2.0	3.5	12.6	0.8	-0.9	-0.4	-1.2	-3.7	-12.5
Professional services	-11.1	2.5	8.6	0.2	-1.0	-0.8	-0.8	-11.0	-8.7
Finance	-7.0	4.3	12.8	-1.1	-1.7	-0.5	-1.2	-7.0	-12.7
Real estate	17.3	27.3	15.7	-7.4	-1.5	-0.1	-1.6	-0.8	-14.4
Information	0.6	8.4	13.8	0.0	-1.1	-0.4	-1.3	-4.7	-14.1
Oil and gas extraction	7.1	12.8	22.4	-5.5	-5.4	-0.3	-1.4	-1.8	-13.7

2030, log points $\times 100$ from 2017. Each term split into the industry's own and what its suppliers pass on.
Online Appendix F.

- Long run: sectoral demand drops out of every price. AI demand moves **quantities**. [► Derivation](#)
- In 2030 the cost of capital adds 27.3 points to the price of real estate through its own buildings and 15.7 through its suppliers. Weighted by consumption, 61% of its effect comes through suppliers. [► Peaks](#) [◀ Back](#)

Other industries: hours with and without a wealth effect



- Hours by industry relative to all hours, log points $\times 100$ from 2017, with GHH preferences and with separable ones. The correlation is 0.99 in 2030 and 1.00 in the long run.
- Total hours depend on preferences, 14.2% with GHH and -7.6% with separable preferences. The wage rises 26.5 log points in both (Caselli and Manning). [► Wage and hours](#)
- With Cobb–Douglas nests, construction loses hours in 2030 as other investment falls further. The long-run pattern holds (correlation 0.88). [◀ Back](#)

Other industries: where AI sits

AI frontier	Corr. with amplif.	Long-run GDP, %	Multiple of AI prod.	
			Model	Network
Observed	-0.19	59.4	3.44	1.84
Uniform	0.00	59.7	3.43	1.85
Most amplified	0.96	63.4	3.60	1.98
Least amplified	-0.86	55.2	3.26	1.73
Consumer-facing	-0.73	47.5	2.96	1.58

Same news; the frontier ν_j moved across industries with its 2017 wage-bill average fixed, so 2017 and the build-out (2.25–2.48% of GDP) stay the same. Amplification: capital-inclusive over ordinary Domar weight.

- AI's buyers resemble the economy. Weighted by AI purchases, upstreamness is 1.85 and the frontier 0.38, against 1.84 and 0.36 by value added.
- Amplification ranges from 1.35 (Health) to 2.63 (Construction). The observed frontier is uncorrelated with it (-0.19).
- Long-run GDP goes from 47.5% (AI in consumer-facing industries) to 63.4% (AI where amplification is highest).

► In full

◀ Back

More results: how we read prices, each price equals unit cost

- We split the change in each industry's price into four parts:

$$\Delta \ln P_j = \text{upstream inputs} + \text{cost of capital} + \text{sectoral demand} + \text{wage and AI}.$$

- **Cost of capital**: the user cost $r_t + \delta_j$. It moves with the interest rate, most for long-lived capital.
- **Sectoral demand**: the rental on capital above its user cost, and the premium of an hour over the wage. It is zero in the long run, when supply is infinitely elastic.

► Derivation

► Why sectoral demand mirrors the cost of capital

◀ Back

More results: the cost of capital, not demand, moves the prices that rise most

Industry	Price $\Delta\%$	Decomposition of the price change, $\Delta\%$					Price, long run $\Delta\%$
		Upstream inputs	Cost of capital	Sectoral demand	Wage and AI	Output $\Delta\%$	
Real estate	21.4	-2.0	23.7	0.8	-1.0	8.4	1.6
Utilities	14.3	-0.5	16.3	2.2	-3.7	18.6	1.3
Oil and gas extraction	12.3	3.7	9.5	0.9	-1.8	26.5	0.9
Petroleum products	10.6	4.5	5.1	1.5	-0.5	15.4	0.2
Electrical equipment	0.4	-1.3	5.0	-0.6	-2.6	17.4	-0.7
Construction	-0.5	-0.5	0.8	0.1	-0.8	48.0	4.2
Machinery	-1.6	-0.7	-0.2	-0.4	-0.2	-2.4	-0.2
Computers	-1.7	-0.5	-0.3	-0.4	-0.4	-1.3	-4.8

Percent change from 2017, relative to the consumption basket, at each industry's highest price after the news. Across the 37 industries, the peak price rises with the share of the interest rate in unit cost (correlation 0.90).

Conclusion. The prices that rise most are those of industries with long-lived capital. They rise because capital costs more, not because of demand.

[► Over time](#)[► Suppliers of AI capital](#)[► Adjustment costs](#)[◀ Back](#)


More results: AI lowers the prices of the industries it can do most for

Industry	Price $\Delta\%$	Decomposition of the price change, $\Delta\%$				Output $\Delta\%$	Price, long run $\Delta\%$
		Upstream inputs	Cost of capital	Sectoral demand	Wage and AI		
Management	-15.0	-1.0	1.6	0.1	-15.6	18.6	-6.5
Professional services	-14.6	-2.7	2.1	0.6	-14.6	26.0	-6.1
Computers	-10.4	-5.9	4.4	1.0	-9.9	33.0	-4.8
Administrative services	-9.8	-2.7	1.7	0.3	-9.1	18.8	0.6
Education	-9.6	1.3	3.3	-0.1	-14.1	17.1	-3.1
Finance	-9.0	-3.6	3.5	0.5	-9.4	18.5	-6.1
Health	-8.6	-1.0	2.1	-0.2	-9.5	14.9	1.4
Apparel	-8.3	-3.2	1.5	0.2	-6.9	16.7	2.2

Percent change from 2017, relative to the consumption basket, at each industry's lowest price after the news.

Conclusion. The saving from AI arrives before the wage catches up, so these prices fall most in the 2030s. Where AI can do many tasks, they stay lower for good.

More results: the suppliers of the build-out grow, the services AI can do shrink



The suppliers of the build-out grow most; finance, management and professional services shrink.

Domar weight: gross output over GDP. AI services, not shown, reach 4.9% of GDP.

More results: who grows, by characteristics fixed before the news

Characteristic	Real output		Price	
	2030	Long run	2030	Long run
AI network exposure	3.1	2.1	0.1	0.1
AI frontier ν_j	0.0	0.5	-3.2	-2.6
Sold as investment	1.1	-0.3	-0.4	0.3
Capital share	-1.5	0.1	4.9	0.1
Life of capital	-0.2	-0.2	0.5	0.6
Sold to consumers	-0.1	-0.4	-0.5	0.4
R^2	0.84	0.96	0.86	0.91

- OLS across 37 industries on characteristics fixed before the news; standardized coefficients. AI network exposure: the output needed, directly and through the 2017 network, for the AI sectors' purchases.
- Output follows AI network exposure (2.1 per s.d. in the long run). Prices do not: they follow the capital share in 2030 (4.9, the cost of capital) and the AI frontier in the long run (-2.6).




More scenarios: known from 2022

Scenario	AI, % of GDP		GDP, $\Delta\%$		Interest rate, peak	Peak price, $\Delta\%$				
	Investment, 2025–32	Revenue, 2030	2030	Long run		AI services	Computers	Electrical equipment	Utilities	Real estate
Baseline	2.43	3.1	11.7	59.4	8.5%	13.8	−1.7	0.4	14.3	21.4
Transformative AI	3.13	3.8	14.1	158.6	11.7%	15.4	−2.8	−0.7	24.5	38.0
A smaller build-out	1.75	2.3	8.8	27.2	7.2%	18.1	−0.9	1.1	8.5	12.1
A power bottleneck	2.43	3.1	11.6	59.4	8.5%	13.8	−1.6	0.9	19.2	21.6
No imported IT equipment	2.45	3.1	12.0	60.1	8.2%	16.5	4.4	0.9	14.2	20.9
Foreign saving	2.43	3.0	12.5	57.6	7.5%	12.1	−1.2	1.9	13.6	20.4

Transformative AI and the smaller build-out recalibrate the news to the same 2025–26 data and another projection for 2025–32. The next two keep the baseline news and change the economy. Foreign saving recalibrates it with the current account. Peak prices in % from 2017, relative to the consumption basket.


- The projection sets the size of the boom, more than proportionally.
- Frictions in supply move the prices of the scarce input, not GDP.
- Foreign saving finances 32% of AI investment, yet the rate still peaks at 7.5%.

More scenarios: transformative AI, every announced data center built



- Same data for 2025–26. The projection is every announced US data-center project (Van Nieuwerburgh) times the generative-AI share: 3.13% of GDP a year over 2025–32, against 2.43% in the baseline.
- AI efficiency keeps rising for longer: nine-tenths of the rise after 8.4 years, against 6.9.

More scenarios: GDP growth in the baseline and with transformative AI



- Over 2023–40 GDP grows 1.7% a year on average in the baseline and 3.1% with transformative AI.
- In the long run GDP is 59.4% above 2017 in the baseline and 158.6% with transformative AI.
- The interest rate peaks at 11.7%, against 8.5%.

More scenarios: a power bottleneck, and no imported IT equipment


A power bottleneck

- The grid is ten times slower to build: utility capital grows at most 2.5% a year.
- The price of electricity peaks at 19.2% in 2043, against 14.3% in 2033. Sectoral demand, not the cost of capital, now drives it.
- GDP barely moves: electricity is 0.046 of each dollar of AI services.

No imported IT equipment

- All the computers of the AI sectors are made at home, an extreme form of trade restrictions.
- Output of computers rises 74.0% by 2030, against 17.1%.
- Computers peak at 4.4%, against -1.7%, and AI services at 16.5%, against 13.8%. GDP barely moves.

More results: capital deepening, an exact identity



- Long run, log points: $\text{GDP } 46.4 = \text{wage } 26.5 + \text{hours } 13.2 + \text{fall in the labour share } 6.6$ ($59.4\% \rightarrow 55.6\%$: AI capital earns 3.6% of GDP).
- The [network](#) predicts the wage: with the rate back at its 2017 level, capital is a produced input, and each industry's AI saving counts with its capital-inclusive Domar weight. 24.8 log points, 1.84 times the AI-productivity term. [► Mechanism](#)
- One sector gives 1.80: the network adds little *here*.

More scenarios: the projected build-out sets the size of the boom

Projected AI investment, 2025–32	AI efficiency 90% by	AI revenue, long run	GDP, $\Delta\%$		Interest rate, peak	Peak price, $\Delta\%$	
			2030	long run		Utilities	Real estate
1.75%	2028	2.9%	8.8	27.2	7.2%	8.5	12.1
2.00%	2029	3.5%	10.0	36.2	7.5%	10.3	15.1
2.25%	2029	4.3%	11.0	48.1	7.9%	12.5	18.5
2.43% (baseline)	2030	4.9%	11.7	59.4	8.5%	14.3	21.4
2.75%	2031	6.3%	12.8	88.3	9.5%	18.1	27.6
3.13% (transformative)	2031	8.5%	14.1	158.6	11.7%	24.5	38.0

Each row recalibrates the news to a projected average of AI investment over 2025–32, in % of GDP, with the same data for 2025–26.

- From 1.75% to 3.13% of GDP, the long-run gain in GDP goes from 27.2% to 158.6%. The interest rate and prices respond much less.
- The data for 2025–26 fix the take-off. A larger projection for the same take-off means that AI keeps becoming cheaper after 2032; a smaller one builds less capacity, so compute becomes scarcer.

Robustness: Cobb–Douglas nests

Elasticities	AI invest. 2025–32	GDP long run	Rate peak	Other invest. 2026	C 2026	Invest. per s.d. of life, 2026	Prices 2030, R^2 own / network	Corr. of hours 2030 / long run
Baseline elasticities	2.43	59.4	8.5	−8.4	0.5	−9.0	0.62 / 0.87	
Cobb–Douglas, same news	2.67	59.9	8.3	−18.1	3.1	−10.7	0.57 / 0.82	0.22 / 0.88
Cobb–Douglas, news recalibrated	2.43	50.6	7.8	−16.4	2.8	−9.6	0.56 / 0.81	0.29 / 0.88
consumption only	2.49	59.5	8.5	−9.8	1.0	−7.1	0.60 / 0.86	0.96 / 1.00
intermediates only	2.51	59.3	8.5	−11.0	1.3	−7.8	0.61 / 0.85	0.73 / 0.98
value added vs. intermediates only	2.47	59.6	8.5	−9.8	1.0	−7.4	0.61 / 0.86	0.97 / 0.99
capital vs. task bundle only	2.55	59.6	8.2	−13.8	1.8	−14.9	0.62 / 0.86	0.89 / 0.99
investment bundles only	2.44	59.5	8.5	−8.5	0.6	−9.0	0.63 / 0.87	1.00 / 1.00

Every CES nest at 0.999, or one at a time. Same 2017 economy: all prices are one, so the weights are the 2017 shares. The AI–labour elasticity stays at 1.5.

- Same news: the long run hardly moves (GDP 59.9%, wage 26.7 log points). Recalibrated to the same data, long-run GDP rises 50.6%.
- Other investment falls 18.1% in 2026, mostly through capital against tasks. Prices, network exposure and the life of capital hold; hours in 2030 do not (correlation 0.22).

More results: the long-run arithmetic

Log points $\times 100$	GHH	Separable
AI productivity (Hulten)	13.5	13.8
$\times$ capital-inclusive weights	24.8	25.1
Wage	26.5	26.5
$+$ fall in the labour share	6.6	6.6
Output per hour	33.1	33.1
$+$ hours	13.2	-7.9
GDP	46.4	25.3

- The Hulten term, 13.5 log points, times a [capital multiplier](#) of 1.84 (one sector, 1.80) predicts the real wage, 24.8 against 26.5 in the model.
- Hours and the shift of income to AI capital do the rest of the 46.4 log points.
[► Identity](#)
- The network amplifies AI as an average industry would. The frontier correlates -0.19 with capital amplification.
- [Where AI lands](#) could move long-run GDP from 47.5% to 63.4%. [► Where AI sits](#)

More scenarios: scenarios revealed later, in numbers

Agents follow the baseline until 2027 or later, learn another future, and re-optimize from the capital already built. Investment cannot be negative.

		GDP, % from baseline			Welfare,	In the year learned: scenario / baseline			Electrical
Scenario	Learned	2030	2040	Long run	% of consumption	AI investment, % of GDP	q of AI capital	Interest rate, %	equipment, 3 years on
<i>Disappointing revenue</i>									
Slow diffusion	2027	-8.4	-10.1	0.0	-3.94	0.2 / 2.5	0.98 / 1.09	3.9 / 6.8	-16.7
Models hit a wall	2027	-8.4	-21.9	-33.4	-13.91	0.5 / 2.5	1.00 / 1.09	3.8 / 6.8	-19.7
Models hit a wall	2030	0.1	-7.5	-13.5	-4.71	2.3 / 2.9	1.04 / 1.07	6.2 / 8.3	-7.8
the same, known early	2027	0.3	-7.4	-13.5	-4.71	2.4 / 2.5	1.09 / 1.09	6.7 / 6.8	-1.5
Models hit a wall	2035	0.0	-0.2	-0.6	-0.17	2.6 / 2.7	1.02 / 1.03	7.2 / 7.3	-0.4
<i>A security incident</i>									
Five-year pause	2030	0.1	-2.4	0.0	-0.95	2.3 / 2.9	1.04 / 1.07	6.3 / 8.3	-7.9
Compliance cost	2030	-0.5	-1.3	-2.3	-0.82	2.8 / 2.9	1.06 / 1.07	8.0 / 8.3	-1.5
<i>A slow build-out</i>									
Data centers	2027	-0.3	-0.5	0.0	-0.09	2.1 / 2.5	1.48 / 1.09	6.8 / 6.8	-2.8
Grid	2027	-0.1	-0.3	0.0	-0.04	2.4 / 2.5	1.09 / 1.09	6.8 / 6.8	-0.6
<i>A geopolitical cost</i>									
Imported IT 50% more expensive	2027	-0.6	-1.7	-3.0	-1.14	2.0 / 2.5	1.05 / 1.09	6.0 / 6.8	-2.6

Slow diffusion: the rest of AI's rise takes 3 times as long (Internet traffic grew about a third as fast as planned in the late 1990s). Compliance cost: 17% per unit of AI services (upper estimate for high-risk AI under the EU AI Act). Slow build-out: adjustment cost $\times 10$. Welfare: permanent % change in consumption from 2027. Online Appendix H.