

THE AI BOOM IN THE US PRODUCTION NETWORK

By VASCO M. CARVALHO, MATIAS COVARRUBIAS, AND GALO NUÑO*

Preliminary draft, October 2026. Comments welcome.

So far the AI boom has been an investment boom in data centers and their equipment. We add two AI sectors, built from the engineering recipe of an AI data center, to the US production and investment networks, and let every industry buy AI services for the tasks that AI performs more cheaply than labour. We then ask what AI must deliver to justify the investment that firms plan. The answer is 2.1 percentage points of extra GDP growth a year from 2026 to 2040, first from AI itself and then from the capital that accumulates around it. The build-out raises the interest rate, and other industries make room by investing less, most of all where capital lasts longest. To the rest of the economy the boom is a demand shock to a few suppliers and a cost shock to everyone. Its purchases travel up the supply chain, the higher cost of capital travels down it into prices, and hours move from the services AI can do to the industries that build AI. Finally, we run an array of scenarios: a larger build-out, AI that disappoints, a geopolitical shock to the price of imported servers, and a regulatory pause.

So far the AI boom has been an investment boom. US investment in AI data centers and their equipment, as we measure it, rose from 0.16 percent of GDP in 2023 to 0.88 percent in 2025 and 1.75 percent in 2026, far faster than the revenue that AI services earn (Korinek and McKelvey, 2026; Carpinelli et al., 2026). An AI data center is a building with a power plant and a cooling plant, filled with computers. Construction firms put up the shell, makers of electrical equipment supply the transformers and the switchgear, and utilities connect it to the grid. The servers, the largest item, are mostly imported.

To the rest of the US economy, the build-out is therefore a demand shock to a few suppliers. It is also a cost shock to everyone, because it absorbs saving years before AI pays off and so raises the cost of capital for every industry. In the language of Acemoglu et al. (2016), the first shock travels upstream, from buyers to their suppliers, and the second downstream, from suppliers to their customers. We ask three questions. How much growth does the investment that firms plan imply, and where does that growth come from? Who makes room for a build-out that needs resources before it pays off? And how does the boom reach industries that sell nothing to a data center?

To answer them we start from the dynamic multi-sector model of Carvalho et al. (2026), calibrated to the input-output and investment networks of the United States in 2017 (vom Lehn and Winberry, 2022), and add two sectors that did not exist in those data. One owns data-center facilities; the other owns the computers and sells AI services. Their inputs follow the engineering recipe of an AI data center built in 2025–26. Their customers are the other industries, each

of which buys AI services for the tasks that AI can do (Eloundou et al., 2024), whenever AI does them more cheaply than labour (Acemoglu and Restrepo, 2018; Acemoglu, 2025). At the release of ChatGPT in late 2022, agents learn that AI will keep getting more efficient at these tasks. We size this news so that the model reproduces US AI investment in 2025 and 2026 and its projected average to 2032. The model then tracks the revenue of AI providers, which the calibration leaves free (Figure 1).

The investment that firms plan implies 2.1 percentage points of extra GDP growth a year from 2026 to 2040. Until 2030 that growth comes mostly from AI itself, as firms hand tasks to AI as it becomes cheaper. In the 2030s it comes mostly from capital: with tasks cheaper, every industry accumulates capital to work with them. Households that expect to be richer save for the build-out only at a higher return, so the real interest rate rises by 4.3 points. With the intertemporal elasticity that household evidence supports, consumption never falls below its 2017 level. The investment of other industries makes room instead, and so, for a while, do wages.

To follow the boom across industries we use two accounting facts. Each industry's sales add up to what its buyers purchase, and each price adds up to the costs it covers. Applied to each year's flows in the model, they split every industry's change in output by the buyer it serves, and every price change by the industry whose costs it passes on. Four patterns emerge. First, the AI sectors' purchases travel up the supply chain: they are large for electrical equipment and construction, and they reach primary metals and minerals, which sell nothing to an AI firm, through their customers. Second, other industries invest less, most of all where capital lasts longest, which is where a higher interest rate bites hardest. In the aggregate, the build-out takes more from existing industries through

*Carvalho: University of Cambridge and CEPR. Covarrubias: Bank of Spain. Nuño: Bank of Spain, CEMFI and CEPR. The paper does not necessarily reflect the views of the Bank of Spain or the Eurosystem. An online appendix states the full model, the data and further results.

their lower investment than it gives them through its purchases. Third, the cost of capital travels down the supply chain and tilts relative prices: real estate, utilities and energy become more expensive, while the services whose tasks AI can do become cheaper. Fourth, hours move from those services to the industries that build and power the boom.

Last, we ask what alternative futures imply for growth over the next five years. If firms build every data center announced so far, the boom adds 3.9 points a year instead of 2.8. If agents learn in 2027 that AI disappoints, it adds about half a point, whether progress slows or stops, and the bust lands on the industries that build. Little capital is stranded, because the capacity built is still used and AI equipment wears out within a few years. A geopolitical shock that makes imported servers half as expensive again costs about a third of a point a year, as much as a regulatory pause, and bottlenecks in the grid or in data centers move prices rather than growth.

The paper adds the supply side of AI to work that measures AI in the national accounts (Highfill and Samuels, 2026; Korinek and McKelvey, 2026; Makridis and Brynjolfsson, 2026) and to task-based models of AI (Acemoglu, 2025; Korinek et al., 2026; Imas and Moll, 2026; Blumenfeld et al., 2026). Aldasoro et al. (2024) give each sector its own exposure to AI, Wachter and Wachter (2026) size an AI boom by the investment it calls for, and Caballero (2026) studies the interest rate while AI is installed. We place AI sectors, built from their recipe, inside the input-output and investment networks (Foerster et al., 2011; Baqaee and Farhi, 2019; vom Lehn and Winberry, 2022), so that the boom reaches other industries both through what the AI sectors buy and through the cost of capital that every industry shares.

I. MODEL AND CALIBRATION

A. Two AI Sectors in the Network

The economy has 37 sectors linked by input-output and investment networks. Each sector accumulates its own capital, built from the investment goods of other sectors, and the economy is calibrated to 2017 (vom Lehn and Winberry, 2022; Carvalho et al., 2026). We add two sectors. AI FACILITIES owns the shell and the power and cooling plant of data centers and rents out powered space. AI COMPUTE rents that space, owns the IT equipment, employs the labs and sells AI services to the other industries. We build both recipes per MW of IT load for a liquid-cooled AI data center of 2025–26, from cost-consultant, engineering and company data, and scale them to \$1 of AI services (online Appendices A and B). The two sectors hold 3.01 dollars of capital per dollar of AI services sold in

a year. Construction makes up 49 percent of the investment of AI FACILITIES and electrical equipment 26 percent; computers make up 90 percent of the investment of AI COMPUTE. As in 2017, 95 percent of these computers are imported, so imported IT equipment is 66 percent of AI investment in 2030. The economy pays for it with exports in the proportions of its 2017 consumption.

B. Demand for AI

In every sector, including the AI sectors, labour input is a Cobb–Douglas bundle of tasks $i \in [0, 1]$. Labour performs any task with productivity one. AI can perform a share v_j of the tasks, the frontier, with productivity $\psi_t x(i)$. AI efficiency ψ_t is common to all tasks, and $\ln x(i)$ is normal across feasible tasks with mean zero and standard deviation s . A firm automates a feasible task when AI does it more cheaply, $p_{Zt}/(\psi_t x(i)) < W_t$, where p_{Zt} is the price of AI services and W_t the wage. Let $c_t = \ln(p_{Zt}/(\psi_t W_t))$ be the log cost of AI relative to labour in the median feasible task. The unit cost of the bundle and AI’s share of it are

$$\begin{aligned} \ln P_{T,jt} &= \ln W_t - v_j [s \phi(c_t/s) - c_t \Phi(-c_t/s)], \\ a_{jt} &= v_j \Phi(-c_t/s), \end{aligned} \quad (1)$$

where ϕ and Φ are the standard normal density and distribution. As the price of AI per efficiency unit, p_{Zt}/ψ_t , falls relative to the wage, firms hand AI a smoothly rising share of the tasks it can do, $\Phi(-c_t/s)$, and the cost of their task bundle falls. The frontier v_j is the share of j ’s wage bill in tasks that a large language model can do at least twice as fast at equal quality, as rated by Eloundou et al. (2024), with occupations weighted by their 2025 wage bill. It is 36 percent for the economy, ranging from 18 percent in FOOD SERVICES to 52 percent in FINANCE.

C. The News and Its Calibration

At the release of ChatGPT at the end of 2022, agents learn the path of AI efficiency, and they act on it from 2023:

$$\ln \psi_t = \ln \psi_{\text{pre}} + D w_t, \quad (2)$$

where w_t is zero until 2023 and rises along a logistic curve to one, nine-tenths of the way after T_{90} years. We choose the three parameters, ψ_{pre} , D and T_{90} , to match three facts about US AI investment, the spending of the two AI sectors. We measure it as the capital spending of the five largest cloud providers (Van Nieuwerburgh, 2026), times the share of data-center spending that serves generative AI (Morgan Stanley Research, 2025) and the US share of their long-lived assets, plus the capital spending of CoreWeave and xAI (online Appendix C). It was 0.88 percent of GDP in 2025 and 1.75 percent in 2026. The third target is its

projected average over 2025–32, 2.43 percent of GDP, which follows analysts’ path of capital spending to 2028 and keeps AI investment at its share of GDP after that. We set s so that the elasticity of substitution between AI services and labour is 1.5 at the adoption of 2025 (Blumenfeld et al., 2026).

Figure 1 shows the fit. The news reproduces the take-off of AI investment, which peaks at 3.0 percent of GDP in 2029 and settles a little lower once the stock of AI capital has caught up with demand. AI efficiency grows fastest in 2027, by 9.0 percent that year, and the share of feasible tasks automated rises from 0.7 percent in 2025 to 16 percent in 2030 and 23 percent in the long run. AI revenue is an outcome of the model. It is 0.14 and 0.36 percent of GDP in 2025 and 2026, close to the revenue of AI providers, 0.087 and 0.48 percent. The model thus reproduces a boom in which investment runs years ahead of revenue.

Capital pays a quadratic adjustment cost, $\kappa = 2$, which makes fast growth of a sector’s capital costly, and moving hours across sectors pays a cost too, $\kappa_L = 1$. Preferences are GHH, with a Frisch elasticity of 0.5 and an intertemporal elasticity of substitution (IES) of 0.5, the value best supported by evidence on households (Crump et al., 2022; Havránek, 2015). Firms are competitive and price at unit cost. We solve for the perfect-foresight path from the 2017 steady state.

II. THE AGGREGATE BOOM

A. Growth and Its Sources

Figure 2 splits the growth that the news adds to GDP into its sources. GDP is flat until 2026. AI is still too expensive to take over many tasks, and hours dip slightly with the wage. Growth then accelerates: GDP grows 2.8 percentage points a year faster from 2026 to 2030 and 1.8 points faster in the 2030s, and by 2040 it is 33.3 percent above its 2017 level.

The sources of that growth change over time. To 2030 the boom is mostly AI productivity. Firms hand tasks to AI as it becomes cheaper, and the fall in the cost of their task bundles accounts for 80 percent of the gain in GDP from 2026 to 2030. In the 2030s capital takes over: other capital accounts for 47 percent of the gain from 2030 to 2040, AI productivity for 25 percent and hours for 18 percent. As AI takes over some of labour’s tasks, every industry’s bundle of tasks becomes cheaper, and industries use more of it together with more capital. This is the induced capital accumulation of Imas and Moll (2026), and it takes time because capital is built slowly and saving is scarce while the interest rate is high.

The long run shows the arithmetic most clearly, because the interest rate is then back at its 2017 level.

AI’s direct gain, the Hulten term, is 13.5 log points. The capital that accumulates in response multiplies it by 1.84, close to the one-sector multiplier $1/(1 - \alpha)$ of 1.80, and this prediction matches the long-run rise of the real wage almost exactly, 24.8 log points against 26.5 in the model. Hours, which rise with the wage under GHH preferences, and the shift of income towards AI capital account for the rest of the long-run gain of 46.4 log points (online Appendix F). The network adds little to the multiplier, because AI’s tasks fall on industries of average capital amplification: across industries, the frontier has a correlation of -0.19 with the capital an industry’s output embodies. Where AI lands could matter. With the same average frontier, long-run GDP would rise by 63.4 percent if AI’s tasks were concentrated where capital amplification is highest, and by 47.5 percent if they were concentrated in the industries that sell to consumers, against 59.4 percent in the baseline.

The growth of hours depends on preferences. GHH preferences have no wealth effect on labour supply, so hours rise with the wage, by 14.2 percent in the long run. With separable preferences, under which richer households work less, hours fall by 7.6 percent and long-run GDP rises by 29.1 percent instead. The wage is the same in both cases, because once the interest rate is back at its 2017 level it depends only on technology.

B. The Interest Rate and Who Makes Room

The model has no debt, so the real interest rate is the return at which households agree to postpone consumption, and it equals the return on capital. The Euler equation ties it to expected growth: approximately, the rate equals the rate of time preference plus the growth of the GHH composite of consumption and leisure divided by the IES. Households who expect to be richer would like to consume part of that wealth already, and only a higher return keeps them saving for the build-out. With an IES of 0.5, each point of expected growth adds 2 points to the rate. The composite grows by up to 2.0 percent a year, and the one-year real rate rises by 4.3 points, from 4.2 percent in 2017 to 8.5 percent in 2029. Before that it dips to 3.8 percent in 2025. At the news consumption jumps, then edges down while the build-out ramps up, and an expected fall in consumption lowers the rate. The rise comes in every case we run, by 2.1 to 10.2 points: more when the boom is larger and households less willing to wait.

How willing households are to wait decides who makes room. With an IES of 0.5, consumption never falls below its 2017 level. The investment of existing industries falls instead, to 8.4 percent below its 2017 level in 2026, and recovers it only in 2029. Workers




FIGURE 1. THE INVESTMENT BOOM AND ITS CALIBRATION

Notes: Baseline. Panel A: investment of the two AI sectors, including imported IT equipment. The squares are our measure of US AI investment in 2023–26; the targets are 2025, 2026 and the projected average over 2025–32 (the horizontal line). Panel B: sales of AI services. The circles are the revenue of AI providers, (OpenAI + Anthropic)/0.67, which the calibration does not use (online Appendices C and D). The vertical line marks the news.




FIGURE 2. THE SOURCES OF GDP GROWTH

Notes: Baseline, relative to the 2017 steady state, which is also the path without the news. Contributions to the annual growth of GDP at 2017 prices, percentage points. AI productivity is each industry's fall in the cost of its task bundle at given prices of labour and AI services, weighted by its spending on tasks. Capital is valued at its user cost and hours at the wage; hours include the small scarcity wedges. The vertical line marks the news.

make room too. A higher return on capital lowers the wage that firms can pay at given technology, so in 2030 the real wage is 2 percent below its 2017 level although GDP is 11.7 percent higher, and the labour share falls from 59 to 52 percent. In the long run the wage is 30 percent higher and the labour share settles at 56 percent. With an IES of 1 or 2, consumption pays instead: it falls 2.2 or 4.7 percent below its 2017 level, and other investment never falls.

Real yields offer a test. The real yield curve steepened after November 2022, in the data and in the model: by 2026 the gap between ten-year and two-year real yields had widened by 1.1 point in the data and 1.5 in the model. The model's level moves too much, however. The real forward rate five years ahead rose by 1.0 point in the data and 3.3 in the model. An IES of 2 brings the model to 1.4, but then consumption, rather than investment, makes room for the build-out. Yields also move with monetary policy and term premia, which the model leaves out (online Appendix E).

III. THE BOOM IN OTHER INDUSTRIES

To follow the boom industry by industry, we decompose the model's solution with two accounting facts. Market clearing says that an industry's sales are the sum of what its buyers purchase; tracing purchases back through each year's input-output coefficients, we split each industry's change in output by the final demand it ultimately serves. Zero profits say that a price equals unit cost; solving the cost equations through the network, we split each price change into the costs of the industry itself and those its suppliers pass on (online Appendix F). The first decomposition follows demand up the supply chain and the second follows costs down it. Figure 3 shows what they reveal.

A. Demand Travels Up the Supply Chain

For existing industries as a whole, the AI sectors' purchases are modest, because most of their spending goes to imported servers. In 2026 these purchases, direct and through suppliers, add 0.8 percent to the output of existing industries, while lower investment




FIGURE 3. THE BOOM IN OTHER INDUSTRIES

Notes: Baseline. Panel A: change in gross output from 2017 to 2026, percent of 2017 output, by the demand it serves: the AI sectors' purchases from the industry (intermediate inputs and investment goods), the intermediate inputs that the AI sectors' suppliers buy from it, directly and indirectly, with the input-output coefficients of 2026, and the investment of existing industries, which includes the investment of the AI sectors' suppliers; the dot is the total change, which also includes consumption and exports. Panel B: investment of each of the 37 existing industries in 2026, log points times 100 from 2017, against the average life of its capital, one over its depreciation rate; the dashed line is the OLS fit. Panel C: the change in each price from 2017 to 2030, relative to the consumption basket, log points times 100 (the dot), split into the costs of the industry itself and those its suppliers pass on, directly and indirectly (including its purchases from itself). Panel D: hours of each existing industry relative to all hours, log points times 100 from 2017 to 2030, against the share of its tasks that AI can do; the dashed line is the OLS fit.

elsewhere takes 2.8 percent away. The build-out thus takes more from existing industries through their lower investment than it gives them through its purchases. Paying for the imported servers with exports adds a further 1.0 percent in 2026, which reflects our assumption of balanced trade and disappears when foreigners lend to the United States (online Appendix F).

For a few industries, however, the boom is large

(Figure 3A). Sales to the AI sectors add 10.4 percent to the output of ELECTRICAL EQUIPMENT in 2026, and 3.2 percent to that of CONSTRUCTION. From there the demand travels up the supply chain. The AI sectors buy nothing directly from primary metals or minerals, yet these industries gain 2.1 and 1.8 percent of their output by selling to the industries that build data centers. Each dollar the AI sectors spend on existing industries raises their gross output by 1.90 dollars.

Exposure to the boom is a property of an industry's position in the network: an industry can be part of the AI supply chain without selling anything to an AI firm.

B. Other Industries Invest Less Where Capital Lasts Longest

The same panel shows who makes room. In 2026 CONSTRUCTION sells 3.2 percent of its 2017 output more to the AI sectors and 14.3 percent less to other investors, so its output falls. Scarce construction capacity does not explain this. The price of every industry's investment bundle falls that year, by 0.7 to 3.1 percent, and the industries that buy the same goods as the AI sectors cut investment no more than the rest. The adjustment happens in the price of waiting rather than in the price of the goods AI buys.

What predicts the cut is how long an industry's capital lasts (Figure 3B). Investment falls by 1.3 log points more for each additional year in the life of capital. Two forces produce this gradient in 2026, when the one-year rate has barely risen. Long-lived capital is priced off long rates, and those have already risen: the model's ten-year real rate is up 3.1 points by 2026. Long-lived capital is also the easiest to postpone. A firm whose buildings last decades loses little by waiting while the rate is high, because its stock barely shrinks in the meantime (House and Shapiro, 2008). EDUCATION, ACCOMMODATION and REAL ESTATE are among the industries that cut most, and REAL ESTATE invests 18 percent less than in 2017. ELECTRICAL EQUIPMENT, whose sales to the AI sectors grow fastest, invests more. With an IES of 2 the gradient survives, but it is weaker, 3.2 log points per standard deviation of capital life against 9.0 in the baseline, because everyone invests at once and goods become scarce: investment bundles then become up to 4.3 percent more expensive (online Appendix F).

C. The Cost of Capital Travels Down the Supply Chain

By 2030 the boom has tilted relative prices (Figure 3C). Against the consumption basket, real estate is 17.3 percent more expensive, utilities 11.8 percent and petroleum products 8.2 percent, while professional services and management are 11.1 and 12.1 percent cheaper. Demand plays almost no role in this tilt. With constant returns and competition, an industry's price equals its unit cost, which depends on the prices of its inputs, the wage and the cost of its capital, and not on how much the industry sells: in the long run, with the interest rate back at its 2017 level, the AI sectors' demand leaves every price unchanged, as in Samuelson's non-substitution theorem. Even during the boom, ELECTRICAL EQUIPMENT sells 31 percent more in 2030 than in 2017, and its relative

price changes by 0.0 percent.

Two forces move prices instead. The first is the cost of capital, and it reaches prices mostly through suppliers. Every input embodies the capital of the industry that made it, so a higher interest rate raises the cost of steel, electricity and freight, and each buyer passes these costs on. In 2030 the cost of capital adds 27.3 points to the price of real estate through its own buildings and 15.7 points through its suppliers; for utilities the two parts are 18.7 and 15.4 points, and for petroleum products, whose own capital is modest, 6.0 and 26.6. Weighted by consumption, 61 percent of the effect of the interest rate on prices comes through suppliers. The second force is AI's saving, which lowers costs most in the industries whose tasks AI can do and, through them, in their customers. FINANCE gains 17 points from the cost of capital, yet its price falls by 7.0 percent. The industries that become more expensive are those whose costs, their own and their suppliers', are mostly capital; those that become cheaper are those whose tasks AI can do.

How far the cost of capital reaches actual prices depends on institutions. Regulators let utilities charge an allowed return on their rate base, so their prices follow the cost of capital closely, and rents adjust toward the user cost of buildings over time. Elsewhere the model gives an upper bound. It passes the market rate fully and at once into the return that firms require, whereas firms move their discount rates by about a quarter of a change in their cost of capital (Gormsen and Huber, 2025).

D. Hours Move to the Industries That Build AI

Labour follows the same two forces, the build-out's demand and AI's taking over of tasks (Figure 3D). Hours move into the industries that build and power the boom and out of those whose tasks AI can do. In 2030, relative to all hours, hours rise by 18 log points in ELECTRICAL EQUIPMENT and 13 in CONSTRUCTION, and fall by 9 log points in FINANCE and 5 in PROFESSIONAL SERVICES. Across industries, the change in relative hours has a correlation of -0.66 with the share of tasks AI can do in 2030, and of -0.93 in the long run, once the build-out is over. In the long run every industry produces more, but those whose tasks AI can do need fewer hours for each unit of output.

This reallocation does not depend on the wealth effect on labour supply. With separable preferences, hours by industry relative to all hours have a correlation of 0.99 in 2030 and 1.00 in the long run with those of the baseline, although total hours fall instead of rising. Where hours go is a robust prediction of the model; how many there are is a statement about preferences.

TABLE 1—SCENARIOS

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

Notes: GDP growth: average annual growth of GDP at 2017 prices from 2026 to 2031, relative to the path without the news. Except for transformative AI, agents follow the baseline until the year learned and then re-optimize from the capital built (online Appendix H). Transformative AI: the news recalibrated to the same data for 2025–26 and to a build-out equal to every announced US data-center project, 86 percent of it for generative AI. Slow diffusion: the rest of AI's rise in efficiency takes 3 times as long. Wall: AI efficiency stops rising. IT import price $\times 1.5$: imported servers cost half as much again as the exports that pay for them. Five-year pause: no new models for 5 years after a security incident in 2030. GDP in the long run: the new steady state. ^aReal income: GDP less the extra exports that pay for the same servers; GDP grows 2.7 percent a year and is 54.7 percent above 2017 in the long run.

IV. SCENARIOS

The baseline is one future among many. For the next few years, the question is how each alternative changes growth. Table 1 reports, for the main scenarios, AI investment and GDP growth from 2026 to 2031, and how much higher GDP ends up in the long run. In the baseline the boom adds 2.8 percentage points a year to growth over these five years. The scenarios answer three questions: how large the boom will be, whether AI delivers, and what it costs to build.

A larger boom. Analysts' plans may understate the build-out. If every US data-center project announced so far is built (Van Nieuwerburgh, 2026), 86 percent of it for generative AI (Morgan Stanley Research, 2025), AI investment averages 3.13 percent of GDP over 2025–32, against 2.43 in the baseline. Recalibrated to the same data for 2025–26, the news then implies a stronger technology, and the boom adds 3.9 points a year to growth over the next five years. The industries that build feel it first: electrical equipment produces 11.5 percent more than in the baseline in 2031. The gap widens later: in the long run GDP is 158.6 percent above its 2017 level, against 59.4 in the baseline, because the faster the build-out that agents expect, the more efficient AI must become to justify it.

AI disappoints. Suppose that in 2027 agents learn that the rest of AI's rise in efficiency will take 3 times as long, as Internet traffic grew about a third as fast as planned in the late 1990s (Coffman and Odlyzko, 1998), or that AI stops improving altogether. Over the next five years these two futures look almost the same. AI investment collapses, the boom adds only 0.7 or 0.5 points a year to growth, and GDP in 2031 is 10.2 or 11.0 percent below the baseline. The bust lands on the supply chain of the build-out: in 2031 the output of electrical equipment is 17.8 or 22.4 percent below the baseline, and that of construction 21.3 or 19.3 percent.

The first years of a bust do not reveal whether the technology has stalled. The two futures separate only later. With slow diffusion GDP catches up with the baseline and ends 59.4 percent above 2017; with a wall it ends only 6.2 percent above. Little capital is stranded in either. The capacity built is still used, AI equipment wears out in a few years, and the value of AI capital falls by at most 0.27 percent of GDP when the news arrives. What the economy loses is the growth that does not come.

A geopolitical cost. Export controls or a conflict could make the servers that the United States imports more expensive. If agents learn in 2027 that imported servers will cost half as much again, the price of AI services in 2031 is 22.2 percent above the baseline, firms hand fewer tasks to AI, and AI investment falls from 2.5 to 2.0 percent of GDP in 2027 although each server costs more. The boom adds 2.5 points a year to growth instead of 2.8. Part of the cost is a transfer abroad, because the United States exports more to buy the same servers. GDP at constant prices counts those exports as output and is only 0.8 percent below the baseline in 2031; real income, which nets them out, is 1.4 percent below. The loss lasts: in the long run real income is 53.5 percent above 2017, against 59.4 in the baseline. A five-year pause on new models after a security incident in 2030 costs as much growth over these five years, 2.5 points a year, but nothing in the long run, while a permanent compliance cost on AI services (Renda et al., 2021), like the geopolitical cost, lasts.

Bottlenecks. Making data centers or the grid ten times slower to build from 2027 barely moves growth over the next five years: GDP in 2031 is within 0.4 percent of the baseline. Bottlenecks act on rents and prices instead. The data centers already built become worth 2.23 times their replacement cost, and a slow

grid raises the price of electricity by up to 12.3 percent over the baseline, in 2048 (online Appendix H).

Over the next five years, then, the growth the boom delivers depends on two things: how large firms expect it to be and whether AI delivers. Geopolitics and regulation cost about a third of a point of growth a year, and bottlenecks move prices rather than output.

V. CONCLUSION

So far the AI boom has been an investment boom in data centers and their equipment. Placing the sectors that build and run AI inside the US production and investment networks shows what the planned investment implies and how it travels. In the aggregate, it implies growth that comes first from AI and then from the capital that accumulates around it, paid for with a higher interest rate; with households reluctant to postpone consumption, other investment and, for a while, wages make room. Across industries, the boom is a demand shock to a few suppliers that travels up the supply chain and a cost shock to everyone that travels down it, and hours move toward the industries that build AI. If AI disappoints, the capacity remains and the loss is the growth that was expected. Geopolitics and regulation cost about a third of a point of growth a year, and bottlenecks move prices rather than growth.

The industry patterns hold in every variant we run (online Appendix F). The size of the boom depends on the projected build-out, the IES, the absence of a wealth effect on hours and the labour of the labs in the recipe of AI services. The next step is to confront the model with the data of 2023–26 industry by industry: investment by the life of capital, prices in utilities and real estate, and hours in the industries that build AI.

REFERENCES

- Acemoglu, Daron, “The Simple Macroeconomics of AI,” *Economic Policy*, 2025, 40 (121), 13–58.
- 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.
- , Ufuk Akcigit, and William Kerr, “Networks and the Macroeconomy: An Empirical Exploration,” *NBER Macroeconomics Annual*, 2016, 30 (1), 273–335.
- Aldasoro, Iñaki, Sebastian Doerr, Leonardo Gambacorta, and Daniel Rees, “The Impact of Artificial Intelligence on Output and Inflation,” BIS Working Paper 1179, Bank for International Settlements 2024.
- Baqae, David Rezza and Emmanuel Farhi, “The macroeconomic impact of microeconomic shocks: Beyond Hulten’s theorem,” *Econometrica*, 2019, 87 (4), 1155–1203.
- Blumenfeld, Alex, Jonathon Hazell, Chen Lian, and Andreas Schaab, “The Macroeconomic Effect of AI: Sizing the Software Engineering Channel,” Working Paper July 2026.
- Caballero, Ricardo J., “Speculative Growth and the AI “Bubble”,” Working Paper 34722, National Bureau of Economic Research 2026.
- Carpinelli, Luisa, Filippo Natoli, and Marco Taboga, “Artificial Intelligence and the US Economy: An Accounting Perspective on Investment and Production,” Working Paper, Bank of Italy January 2026. arXiv:2601.11196.
- Carvalho, Vasco M., Matias Covarrubias, and Galo Nuño, “Planning Against Disasters in Dynamic Production Networks,” 2026. Working paper.
- Coffman, K. G. and Andrew M. Odlyzko, “The Size and Growth Rate of the Internet,” *First Monday*, 1998, 3 (10).
- Crump, Richard K., Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa, “Subjective Intertemporal Substitution,” *Journal of Monetary Economics*, 2022, 126, 118–133.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, “GPTs are GPTs: Labor Market Impact Potential of LLMs,” *Science*, 2024, 384 (6702), 1306–1308. Working paper version: arXiv 2303.10130.
- Foerster, Andrew T, Pierre-Daniel G Sarte, and Mark W Watson, “Sectoral versus aggregate shocks: A structural factor analysis of industrial production,” *Journal of Political Economy*, 2011, 119 (1), 1–38.
- Gormsen, Niels Joachim and Kilian Huber, “Corporate Discount Rates,” *American Economic Review*, 2025, 115 (6), 2001–2049.
- Havráněk, Tomáš, “Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting,” *Journal of the European Economic Association*, 2015, 13 (6), 1180–1204.
- Highfill, Tina and Jon D. Samuels, “Early Estimates of the Impact of AI Within BEA’s Industry Economic Accounts,” BEA Working Paper WP2026-3, U.S. Bureau of Economic Analysis February 2026.
- House, Christopher L. and Matthew D. Shapiro, “Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation,” *American Economic Review*, 2008, 98 (3), 737–768.
- Imas, Alex and Benjamin Moll, “A Standard Task-Based Solow Model Delivers AI-Driven Double-Digit Growth,” Technical Report, Working paper 2026.
- Korinek, Anton and Patrick McKelvey, “Measuring the AI Economy,” PIIIE Working Paper 26-9, Peterson Institute for International Economics May 2026.
- , Charles I. Jones, Szymon Sacher, Tess Cotter, and Peter McCrory, “Economic Scenarios for Transformative AI,” Working Paper 2026-02, The Anthropic Institute 2026.
- Makridakis, Christos A. and Erik Brynjolfsson, “Counting AI: A Blueprint to Integrate AI Investment and Use Data into US National Statistics,” Brookings Research, Brookings Institution January 2026.
- Morgan Stanley Research, “Bridging a \$1.5tr Data Center Financing Gap,” Research report, Morgan Stanley July 2025.
- Renda, Andrea et al., “Study to Support an Impact Assessment of Regulatory Requirements for Artificial Intelligence in Europe,” Technical Report, European Commission 2021.
- Van Nieuwerburgh, Stijn, “Financing the AI Buildout,” 2026.

Brookings Papers on Economic Activity, draft.
vom Lehn, Christian and Thomas Winberry, "The investment network, sectoral comovement, and the changing US business cycle," *The Quarterly Journal of Economics*, 2022, 137 (1), 387–433.
Wachter, Jessica and Jonathan Wachter, "What Investment Data Implies about the AI Transition," Working Paper 35290, National Bureau of Economic Research 2026.