

AI leadership and Europe

What US and Chinese AI strategies mean for Europe

“Whoever wins AI, wins”

An inquiry into productivity, who receives AI-created value, and control of the systems we use.

Full research report

Brussels | Evidence to 24 September 2026

AI USE AND VERIFICATION

This report was produced largely using generative AI for research, drafting and synthesis. Unchained PA has not independently validated the report as a whole. Factual claims, quotations, figures and interpretations should not be treated as verified by Unchained PA. Readers should check consequential points against the cited primary sources.

The question behind the AI race

The US and China are building different positions in AI. Europe can benefit from both as a user and supplier, while facing questions about productivity, income and dependence.

What does “winning” mean?

US strategy links a large AI ecosystem to standards, economic returns and security. Chinese plans emphasise industrial diffusion, domestic components and the reach of its models. Neither approach guarantees that one country captures all the gains.

What could change for Europe?

Shallow adoption could widen an existing productivity gap. Deep adoption could raise European output while creating income for foreign suppliers. Both outcomes can be true at once; their size is not yet measured.

Where does control matter?

Models, chips, cloud, data and applications create different kinds of dependence. European firms also supply important parts of the global stack. The practical question is which capabilities can be replaced and which interruptions would be costly.

How to read this report

The report compares evidence and competing interpretations. It describes trade-offs and measurable outcomes without recommending a policy programme or treating either the US or China as a single, uniform actor.

How certain is the argument

The mechanisms are plausible and some dependencies are visible today. The size of AI's productivity effect, the durability of supplier profits and the possibility of self-reinforcing frontier leadership remain uncertain. Evidence labels distinguish observed facts, reported estimates, model outputs, scenarios and the report's own inferences.

Contents

The question behind the AI race	2
Executive summary	4
Part 0 The phrase US strategy and China's different approach.....	7
Part I Why productivity matters the foundations.....	11
Part II The productivity challenge in Europe.....	15
Part III Why AI could be economically different	22
Part IV How big could the productivity effect be	26
Part V Productivity is not the same as value capture.....	31
Part VI AI capital and the distribution of economic power.....	36
Part VII The flywheel of technological leadership.....	40
Part VIII US China and Europe where is power actually located	43
Part IX China in the AI landscape	51
Part X From economic power to geopolitical power	54
Part XI Standards dependence and freedom of action.....	57
Part XII Does Europe need to lead the frontier to benefit	61
Part XIII Industrial AI An alternative path for Europe	64
Part XIV Sweden as a case study	66
Part XV European AI policy as an object of study	72
Part XVI Regulation what the evidence shows and does not show	76
Part XVII Competing interpretations	79
Part XVIII The answer to the title question	81
Annex A Chart list	83
Annex B Principal sources with links	84
Annex C How to read the evidence	86

Executive summary

The question. What does “Whoever wins AI, wins” mean when the United States and China pursue AI leadership in different ways? This report follows a personal question through the evidence: how AI could change productivity, where the resulting income may flow, and what control over key systems means for Europe and Sweden. It examines possibilities and uncertainties; it does not set out a policy programme.

The central argument. AI could change how much economies produce, who receives the resulting income, and who controls technologies others depend on. These are related but distinct questions. Europe can gain as a user of American or Chinese AI even if a share of supplier income accrues abroad. It can also fall behind in productivity if adoption is shallow. The size of either effect is uncertain, and ownership alone does not settle who benefits.

Twelve conclusions

1. Two major powers frame AI leadership differently. The US Action Plan connects the largest AI ecosystem to standards and economic and military benefits; a 2025 executive order promotes exports of the full American technology stack to allies. Chinese plans emphasise diffusion into industry and society, domestic components and the global reach of Chinese models [OFFICIAL CLAIM]. These are strategic ambitions, not evidence that either country captures all gains. Europe is already a customer, partner and supplier in this changing landscape.
2. **Productivity differences compound into large differences in resources.** A sustained one-point difference in annual growth makes one economy about 22% larger than another after 20 years and 34% after 30 [ILLUSTRATION]. That gap affects wages, pensions, tax revenue and the room for defence. Institutions decide how much of it becomes capability.
3. **Europe's productivity growth has trailed the US over most of three decades.** Real output per hour grew about 1.7% a year in the US over 1995–2025, against 1.1% in the EU27 and 0.9% in the euro area [ESTIMATE] (AMECO). The ECB finds US growth of about 50% over 1995–2019 against 28% in the euro area [FACT]. The large western euro economies were close to US levels in the mid-1990s; in 2024 the EU27 produced about 80% of US output per hour [ESTIMATE]. The gap opened in 1995–2005 and again after 2019. Research points to both weaker digital production and less effective deployment; their relative weights are contested.
4. **AI is a plausible general-purpose technology, but the size and timing of its macroeconomic payoff remain uncertain.** Published estimates must be read in their own units. Added TFP growth is mostly put at about 0.1–0.7 points a year over a decade (one model's range reaches 1.2); added labour-productivity growth at about 0.4–1.3 points in US and UK scenarios, or 1.5 at full adoption in one bank estimate. Europe-specific estimates sit at the lower end [MODEL OUTPUT]. Larger effects appear in scenarios that assume capabilities not yet demonstrated [SCENARIO]. Task-level gains do not translate directly into economy-wide growth: task shares, adoption, costs, bottlenecks and the redeployment of saved time affect the aggregate result.

5. **Using AI, receiving its income, owning it, taxing it and controlling it are different things.** A Swedish firm using American AI can create income for Swedish workers, customers, shareholders and the state while paying foreign suppliers; both countries can gain. How the gains divide is not measured for business use. Consumer-surplus studies show large benefits to US consumers, not the national split **[INFERENCE]**.
6. **Some suppliers are very profitable today; that does not prove permanent rents.** Nvidia earned \$63.7bn of operating income in the quarter to July 2026, an operating margin of about 66% **[FACT]**; the leading model labs appear still to be loss-making **[ESTIMATE]**. Current margins reflect scarcity, quality and an investment boom as well as scale economies. Europe's strongest current position is as a supplier, above all ASML.
7. **If AI shifts income towards capital, ownership and tax design matter more — but a lower labour share is not automatically lower labour income.** Every transformative scenario reviewed shows labour's share falling (in one US model, from 60% to 45% in the extreme case) **[SCENARIO]**. Because output also rises, total labour income stays close to the no-AI path on the author's approximation from that model's outputs **[ILLUSTRATION]**, while its distribution between occupations shifts sharply. Europeans own about 9% of US equities, which gives income but not control **[ESTIMATE]**.
8. **Concentration today is strongest in chips, fabrication, lithography, cloud and distribution; the model layer is more contestable.** Model leads between firms shift within two years, and open-weight and Chinese models trail the frontier by months **[ESTIMATE]**. That does not mean a frontier lead is unimportant or cheap to match. If frontier AI begins to accelerate research, including AI research, a lead could compound **[SCENARIO]**.
9. **Power in AI is split by layer.** In the best-documented sample the US hosts about 75% of AI-supercomputer performance (the sample covers an estimated 10–20% of global capacity), and five US firms own about 71% of global AI compute **[ESTIMATE]**. The US leads in chips, private capital and frontier models; China leads in robots, electricity, manufacturing and open-weight releases **[FACT]/[ESTIMATE]**. Europe leads in no layer outright but holds the EUV lithography chokepoint and strong positions in power equipment.
10. **Productivity becomes power only through institutions, and dependence can create leverage.** Higher productivity enlarges potential resources; political decisions, procurement, industrial capacity and alliances convert them. The EU is not a single state, so conversion depends on coordination among governments. US rules governing advanced-chip access have changed, while Chinese rare-earth and Nexperia controls have disrupted production **[FACT]**. These episodes differ in cause and consequence, but each shows why the location of control matters. A sanctions-related email disconnection at the International Criminal Court remains disputed as to its initiator **[CONTESTED]**.
11. **Europe is active; the question is delivery.** AI Factories are being built, a gigafactory call is open and the Data Act applies. But much of the public money sits in a budget not yet agreed, the first gigafactories are unlikely to operate before about 2028, and adoption depth, energy and scale-up finance receive less money and measurement than compute **[INFERENCE]**. On regulation, the evidence documents real costs from

specific rules, fragmentation and uncertainty, and real reasons for some safeguards; it does not support a general verdict either way.

12. **Sweden has strong assets and real constraints;** conversion remains a hypothesis. Sweden leads the EU on R&D intensity and ICT specialists, ranks third on enterprise AI use and exports 33 TWh of electricity a year [FACT]. Grid capacity, scale-up capital and slowing productivity growth are constraints [ESTIMATE]. Whether these assets turn into Swedish productivity, resident income and dependable access is measurable but not yet established.

The paths under examination. Part XII compares near-frontier capability, deep adoption, industrial AI, assured operation of critical functions, and hosting foreign capacity. Each has different costs, benefits and conditions. The report does not rank them or select a preferred combination.

The answer in one paragraph. The slogan is useful because it draws attention to productivity, supplier income and strategic leverage. It is misleading if taken to mean that AI has one race, one winner or a fixed division of gains. The US leads in frontier models and much of the compute stack; China combines a strong manufacturing base with a growing model ecosystem and a particular emphasis on industrial diffusion. Europe is a user, supplier, regulator and potential builder within both ecosystems. The questions are how these roles interact, where European value is created, and which dependencies matter in practice.

Indicators that would change this assessment (replacing version 1's numerical probabilities):

- Measured productivity growth in AI-intensive sectors in national accounts, in Europe and the US.
- The share of firms using AI in several business processes, by country, including Sweden.
- Evidence that frontier AI shortens research cycles outside narrow domains.
- The market share of non-Nvidia accelerators and the capability gap between open-weight and closed models.
- Operating European compute capacity, not announced capacity; and grid connection times.
- Any use of access to US or Chinese AI, cloud or chips as explicit leverage against an EU government or institution.

Part 0 The phrase US strategy and China's different approach

What was said when and in what context

The written version was posted on Truth Social on the morning of Monday 14 September 2026 (09:58 EDT according to an archive copy) [FACT] (timestamp from the [archive](#), wording corroborated by a [contemporaneous news report](#)). The key line reads: "WHOEVER WINS AI, WINS! We are leading China, and all others, and will continue to do so." The post opened: "The only control or 'guardrails' that AI needs is a STRONG AND SMART (High IQ!) PRESIDENT." It stated that the administration has "tremendous CRIMINAL and REGULATORY power over these companies" and described opposition to AI and data centres as a "SICK conspiracy" of which "the only one that is happy about it is China" [FACT].

The phrase is also documented in spoken remarks. During a live call with Nvidia's Jensen Huang on 14 September, the President said: "I have an expression. It's whoever wins AI, wins. That's how big it is. It's bigger than the internet." [FACT] ([Roll Call Factbase transcript](#), 00:03:54). A similar remark had been made to reporters in Ireland the day before [FACT] (as reported).

The immediate trigger was Dario Amodè's essay "[We Must Pace the Frontier](#)", published on 12 September 2026. It called for slowing frontier-capability gains, embedded and external evaluation, coordination among democracies and eventual global coordination [FACT]. Sam Altman and Elon Musk were reported to have endorsed the call [ESTIMATE]. The post also referred to a running dispute: the Pentagon had designated Anthropic a "supply-chain risk" in March 2026, a decision a federal judge ruled to be unlawful retaliation in late August 2026 [FACT].

The statement is useful because it makes an implicit theory explicit: AI leadership may bring economic and geopolitical advantage. Published US policy gives that theory concrete form. China is pursuing its own mix of model development, industrial deployment and infrastructure. The report compares these approaches and tests what each might mean for Europeans, without assuming that either country's gains are Europe's losses.

Rhetoric published doctrine and what economics establishes

Three layers need to be kept apart.

1. Rhetoric. The post and the spoken remark present AI as a single-winner contest and recast guardrails as presidential discretion [INFERENCE]. That is a political framing.

2. The administration's published ambitions [OFFICIAL CLAIM]. These are more detailed and more useful for European policymakers.

- [America's AI Action Plan](#) (July 2025) opens: "The United States is in a race to achieve global dominance in artificial intelligence." It argues that "whoever has the largest AI ecosystem will set global AI standards and reap broad economic and military benefits", and describes AI as "an industrial revolution, an information revolution, and a renaissance — all at once" (p. 1). Its three pillars are innovation,

infrastructure, and international diplomacy and security. It is not only about frontier models. It states that "the bottleneck to harnessing AI's full potential is not necessarily the availability of models, tools, or applications. Rather, it is the limited and slow adoption of AI" (p. 5). It calls for "exporting its full AI technology stack — hardware, models, software, applications, and standards — to all countries willing to join America's AI alliance" (p. 20).

- [Executive Order 14320](#) (23 July 2025) makes it US policy "to preserve and extend American leadership in AI and decrease international dependence on AI technologies developed by our adversaries by supporting the global deployment of United States-origin AI technologies". Its "full-stack" export packages cover chips, servers and accelerators, data-centre storage, cloud services and networking; data pipelines; models; security; and applications.
- [Executive Order 14363, the Genesis Mission](#) (24 November 2025) builds an integrated AI platform on federal scientific datasets to accelerate discovery, in an effort described as "comparable in urgency and ambition to the Manhattan Project".

3. What economics establishes. The published doctrine rests on claims that are partly supported and partly open:

- *Supported:* productivity differences compound; some layers of the AI stack are highly concentrated; effective adoption is necessary to realise economic gains, while capability, reliability, cost and access can also constrain particular uses (Parts I, IV and VII).
- *Open:* whether the "largest ecosystem" captures durable rents, whether a model lead compounds through faster research, and how gains are divided between countries (Parts V, VII and XII).
- *Not implied:* that US companies' success automatically benefits all Americans — the scenario evidence suggests AI may shift income from knowledge workers to capital owners (Part VI) — or that it harms all Europeans, who are also users, customers, shareholders and suppliers (Part V).

Why it matters for Europe. American firms provide tools that can raise European productivity, and the US remains an important European partner. The US plan also seeks worldwide adoption of an integrated American technology stack. China's industrial AI strategy, open-weight models and manufacturing scale offer another source of innovation and competition, alongside distinct security, governance and supply-chain concerns. European outcomes depend on adoption, bargaining power, domestic capabilities and access across both ecosystems; they cannot be inferred from the nationality of a supplier alone.

The China comparison. Chinese plans put particular emphasis on applying AI across manufacturing and public life, building domestic components, and extending the reach of Chinese models. These aims are not the same as the US full-stack export strategy, although both countries seek economic and strategic influence. Part IX examines what each approach could offer or constrain for European users. [OFFICIAL CLAIM]/[INFERENCE]

What could winning AI mean

"Winning AI" can refer to at least sixteen different things. They are not the same, and a country can lead in one while lagging in another.

Dimension	What "winning" would mean	Who leads today	Does leadership translate into national gain?
Frontier models	Building the most capable models	US; China ~7 months behind	If the lead persists and is monetised, or accelerates research
Advanced chips	Designing AI accelerators	US (Nvidia >60% of AI compute)	Currently: high margins, export leverage
Chip fabrication	Making leading-edge chips	Taiwan (TSMC)	Currently: high margins, but geopolitical exposure
Chip equipment	Lithography and tools	Netherlands (ASML), US, Japan	Chokepoint and high margins
Compute infrastructure	Owning AI data centres	US (~75% of documented AI-supercomputer performance; five US firms own ~71% of AI compute)	Where owned; less so where only hosted
Energy	Cheap, abundant power	China (scale); Nordics (price)	Only if converted into compute and industry
Cloud/hyperscalers	Selling compute as a service	US (~70% of the European market)	Margins, lock-in, legal reach
AI companies and IP	Owning the firms and patents	US (value), China (patent counts)	Profits, control; tax depends on where activity is located
Talent	Attracting researchers	US (~59% of top researchers work there)	Yes, with a lag
Finance	Funding AI at scale	US (83% of private AI investment)	Enables the rest
Economy-wide adoption	Using AI deeply across firms	Mixed; Nordics, Singapore, UAE high	The main productivity channel
Industrial application	AI in manufacturing and engineering	China (volume), Germany/Nordics (density)	If firms keep the differentiating layer
Scientific discovery	AI that accelerates R&D	US/UK labs (DeepMind), China	Potentially large, highly uncertain
Robotics / physical AI	AI in machines	China (54% of robot installations)	For manufacturing power
Military and intelligence	AI in ISR, targeting, cyber	US; China investing; Ukraine as test bed	Security and deterrence
Standards and platforms	Setting technical and governance rules	Contested; US stack exports vs China's open-weight and UN track	Long-term lock-in

Sources: Epoch AI (Apr 2025; Jan 2026; Apr 2026); Stanford AI Index 2026; Synergy Research (Jul 2025); IFR World Robotics 2025; MacroPolo Global AI Talent Tracker 3.0. Detailed data in Part VIII.

So what? "Winning AI" is not one race. It is a bundle of distinct contests with different economics. Some layers (chips, cloud, distribution) are highly concentrated today. Others (models) are contested. Others (adoption, industrial application) are open to any country that organises itself well. A national strategy that treats them as one thing will spend in the wrong places.

The central question winner-takes-most diffusion or research acceleration

Three models of the world compete.

- **The winner-takes-most model.** AI capability compounds: better models attract users, revenue, capital, compute and talent, which produce better models. The leader pulls away, captures most of the profits, and converts them into economic and military advantage. Followers become dependent customers.
- **The diffusion model.** AI is a general-purpose technology like electricity. Much of its economic value is realised by users across the whole economy, not by inventors. Models commoditise quickly. Countries that adopt broadly and deeply capture much of the productivity gain, whoever owns the frontier.
- **The research-acceleration model.** If frontier AI speeds up science and AI research itself, a lead of months could compound into a lasting advantage. This is the bet behind the Genesis Mission. It is a serious possibility, not an established fact.

Jeffrey Ding's historical study of the three industrial revolutions (*Technology and the Rise of Great Powers*, Princeton, 2024) found that power shifted to the states best at diffusing general-purpose technologies across their economies, not to those that pioneered the leading innovations [CONTESTED] (an influential scholarly argument). The evidence in this report suggests each model is partly right, **in different layers of the stack and under different conditions**. That layered answer is the report's main analytical claim, and Parts V–VIII test it.

Part I Why productivity matters the foundations

Why are some societies much richer than others

Over the long run, a country's income per person is mainly determined by how much each hour of work produces. Natural resources, trade surpluses and financial engineering matter at the margin and for particular countries. But across decades and across most economies, the difference between a rich and a poor society is largely the difference in **output per hour of work**: labour productivity.

Output per hour depends on three things:

- **Capital per worker.** A farmer with a tractor produces more than one with a hoe. An engineer with simulation software produces more than one with a slide rule. Economists call the growth of capital per worker *capital deepening*.
- **Skills.** Education, training and experience let people use capital well.
- **How well capital and skills are combined.** This covers technology, management, organisation, competition, institutions and the quality of the legal system. Economists call this residual *total factor productivity* (TFP). It is the part of growth that does not come from simply adding more inputs.

Labour productivity and TFP are different measures. Labour productivity can rise because workers have more capital, even if TFP does not change. Estimates of AI's effect are published in both units, and sometimes as effects on the level of GDP. The report keeps these apart (Part IV).

A general-purpose technology such as electricity or AI can raise all three at once. It is new capital, it changes which skills are valuable, and it changes how firms can be organised.

From productivity to resources the chain

The transmission chain is straightforward but rarely spelled out. Each link is a tendency, not a guarantee.

Link	Mechanism
Productivity → output per worker	Each hour produces more goods and services
Output per worker → real wages	Over time, wages tend to follow productivity, with lags and large differences in distribution
Real wages → living standards	Higher real incomes buy more, or buy the same with fewer hours
Output → profits	Firms earn more on the same capital, attracting investment
Wages + profits + consumption → tax base	Revenue can rise with the economy at unchanged tax rates, if the composition of income does not shift towards lightly taxed forms

Link	Mechanism
Tax base → public services	Health, education, pensions and care become easier to fund, although demand and public-sector wages also rise
Profits + savings → investment	Capital deepening and R&D feed the next round of productivity
All of the above → potential resources for power	More room for defence, R&D, aid and external investment. Institutions, procurement and alliances decide what those resources achieve (Part X)

A country whose productivity stagnates faces harder choices between living standards, taxes, debt and public services. Faster productivity growth makes those choices easier. It does not remove them. This is why productivity is not a technical subject: it shapes the budget constraint of democratic politics.

The power of compounding


Small differences in annual growth become very large differences in levels. The arithmetic is simple: a level grows by a factor of $(1 + g)$ each year. The results are not intuitive.

Annual growth	After 10 years	After 20 years	After 30 years	Doubling time
1.0%	110	122	135	70 years
1.5%	116	135	156	47 years
2.0%	122	149	181	35 years
2.5%	128	164	210	28 years
3.0%	134	181	243	23 years
4.0%	148	219	324	18 years

[ILLUSTRATION] Index, year 0 = 100. Pure arithmetic, not a forecast.

Small differences in growth become large differences in wealth

Left: level after 10, 20 and 30 years at different annual productivity growth rates. Right: how much larger the faster economy becomes relative to a 2%/yr baseline. Arithmetic illustrations, not forecasts.



Source: author's arithmetic, level = $100 \times (1+g)^t$. ILLUSTRATION, NOT A FORECAST.

Chart 1

Chart 1. Compounding: small differences in growth become large differences in output. Author's arithmetic, not a forecast.

Chart 1 is unchanged from version 1. It shows output levels; output per worker, GDP per person and wealth are related but different measures.

The right-hand panel of Chart 1 answers the question that matters for policy: **how much larger does an economy become if it grows one percentage point faster than another?** Against a 2% baseline, a sustained one-point advantage gives a level about 10% higher after ten years, 22% after twenty and 34% after thirty (3% against 2%: +21.55% after 20 years, +34.00% after 30). Even a quarter-point advantage adds up to about 8% after thirty years [ILLUSTRATION].

To make that concrete [ILLUSTRATION]:

- EU27 GDP in 2025 is about €18.8 trillion (Commission estimate). One per cent of it is roughly €190bn a year, of the same order as the EU's entire annual budget.
- A 22% level difference after twenty years, applied to today's EU economy, is about €4 trillion a year. That is close to the size of the whole German economy (approximate).
- For Sweden, one per cent of 2025 GDP is roughly SEK 66bn a year. At the 2025 tax ratio of 40.5%, that would be about SEK 27bn in public revenue if the extra output were taxed at the average rate [ILLUSTRATION] (Statistics Sweden; 2025 figures are preliminary).

So what? A persistent one-point productivity gap is not a rounding error. Over a political generation it separates an economy that can fund pensions, healthcare and a higher defence budget with less strain from one that must cut, borrow or raise taxes further. It also enlarges the resources a state can mobilise abroad, although turning resources into influence is a separate step (Part X).

Three cautions

Growth rates are not destiny. Advanced economies rarely sustain growth above 2% a year in output per hour. The US managed about 3% (non-farm business sector) only during 1995–2004 [ESTIMATE] (standard BLS-based literature value). The author's judgement is that a realistic range for Europe over the next decade, without a technology shock, is roughly 0.5–1.5% a year [JUDGEMENT].

Productivity and welfare are not identical. Europeans work fewer hours than Americans, partly by choice. Some gains from technology come as free services and better quality that GDP does not capture (consumer surplus). Distribution matters too: average productivity can rise while median wages stagnate.

Levels and rates are different things. Much confusion in the AI debate comes from mixing up two claims. The first is a one-off increase in the level of output ("GDP 7% higher after ten years"). The second is a permanent increase in the growth rate ("growth 1.5 points higher every year"). The first is worth about 0.7 points a year for a decade and then stops. The second compounds indefinitely. The distinction is used throughout this report.

Part II The productivity challenge in Europe

Before asking what AI will do, the baseline must be clear. Europe enters the AI era with a productivity gap that has widened, in growth terms, over most of the past three decades.

The facts

Real growth: the robust part. Real GDP per hour worked grew by about 1.68% a year in the US over 1995–2025. It grew by 1.14% in the EU27 and 0.86% in the euro area (EA20, a fixed historical grouping of 20 countries; Bulgaria became the 21st member on 1 January 2026) **[ESTIMATE]** (author's calculation from the Commission's [AMECO database](#), Spring 2026 vintage). The ECB, using a different dataset, finds that US output per hour grew by about 50% between 1995 and 2019, against 28% in the euro area **[FACT]** ([ECB Economic Bulletin 6/2024, Box 1](#)).

US output per hour grew faster in 1995–2005 and again after 2019

Left: real growth, the right measure for change over time. Over 1995–2025 US output per hour grew about 1.7% a year, EU27 1.1% and the euro area 0.9%. Right: a single-year level comparison. The 1995 starting level is uncertain and depends on the price method, so the chart does not show a 1995–2025 level path.




Chart 2. US output per hour grew faster in 1995–2005 and again after 2019. Left: real growth index. Right: a single-year level benchmark.

Sources: AMECO Spring 2026 (OVGD, NLHT), author's calculations; [Conference Board Total Economy Database, Sept 2025](#). Values and definitions in the source-table workbook (sheet C2).

The gap came in two waves **[ESTIMATE]**:

Period	US	EU27	Euro area (EA20)	US lead over EU27 (pp/yr)
1995–2005	2.47%	1.82%	1.31%	+0.65
2010–2019	0.63%	1.00%	0.83%	–0.37
2019–2025	1.79%	0.51%	0.29%	+1.28

Real GDP per hour, compound annual growth. AMECO Spring 2026; author's calculations. The 2010s result is data-sensitive: the Conference Board shows the US slightly ahead (0.8% vs 0.7% for the euro area over 2011–19). EU27 growth includes the catch-up of Central and Eastern European members.

The first wave coincided with the US ICT boom. In the 2010s Europe roughly kept pace in this dataset. Since 2019 a second wave has opened. Between late 2019 and mid-2024, output per hour rose by 6.7% in the US and 0.9% in the euro area **[FACT]** (ECB). The ECB attributes part of the recent gap to cyclical factors, including the energy shock, and part to structural ones.

Where did Europe start? It depends on the geography and the price method. There is no single, well-defined answer to "were Europe and the US equal in the 1990s?"

- The ECB's Isabel Schnabel reported that labour productivity in the four largest euro-area economies rose from 25% of the US level in 1945 to about 100% in 1995 **[ESTIMATE]** ([speech of 16 Feb 2024](#), long-run historical dataset). Bergeaud (ECB Forum, 2024) similarly finds euro-area countries close to the US per hour in 1995.
- The Draghi report states that EU productivity "converged from 22% of the US level in 1945 to 95% in 1995" **[ESTIMATE]** ([Draghi Part A, p. 23](#), using the same long-run database at 2010 prices).
- Converting today's AMECO data back to 1995 at 2020 purchasing-power parities puts the EU27 at about 92% of the US level and the euro area above it. That figure is sensitive to projecting 2020 prices 25 years backwards, so this report no longer uses it as a headline.
- The EU27 of today included economies that were much poorer in 1995. Its starting level was therefore lower than that of the large western euro economies.

What is robust is narrower and still important. **The core western European economies were close to US output per hour in the mid-1990s and are clearly below it today.** In 2024 the EU27 produced about 80% of US output per hour and the euro area about 86%, measured at current purchasing-power parities for that year **[ESTIMATE]** (AMECO). Germany was at about 95%, France 89%, Italy 76% and Spain 74% **[ESTIMATE]** (Conference Board, 2024 PPP).

Per person versus per hour. Draghi states that "around 70% of the gap in per capita GDP with US at PPP is explained by lower productivity in the EU" **[ESTIMATE]**. The author's decomposition of AMECO data gives 64% for 2019 and about 76% for 2024–25 for the EU27. The remainder comes from shorter hours and a lower employment rate. In 1995 the per-capita gap was mostly an *employment* gap. Today it is mostly a *productivity* gap, because European employment rates have caught up **[ESTIMATE]**. In the euro area, about a quarter of the gap reflects shorter working hours, which is partly a choice rather than a failure.

Investment, R&D and firms.

- Real private fixed investment grew 4% in the EU between Q1 2020 and Q2 2025, against 18% in the US. EU business investment remains about 1.5 points of GDP below the US **[ESTIMATE]** ([EIB Investment Report 2025/26](#)).
- EU27 R&D spending was 2.24% of GDP in 2024, against 3.44% in the US (2023). Almost the whole gap is business R&D: 1.49% against 2.69% **[FACT]** (Eurostat).
- Of the world's 2,000 largest corporate R&D spenders, EU firms account for 16% of R&D and US firms for 47%. EU firms hold 2% of R&D among the top 50 software firms; US firms hold 88% **[FACT]** (EU Industrial R&D Investment Scoreboard 2025).

- About a fifth of European employees work in firms with ten or fewer people, roughly double the US share. European firms aged 25 or more employ about 10 people on average, against about 70 in the US [ESTIMATE] (IMF).

Measurement warnings


Three measurement issues must accompany any EU–US comparison.

1. **Growth and levels need different tools.** Use real (constant-price) growth to measure change over time. Use purchasing-power parities for one year to compare levels across countries. Eurostat states that current-price PPS series "should be used for cross-country comparisons in a specific year but do not for comparisons over time since they do not constitute time series" [FACT] ([Eurostat metadata](#)). Version 1 presented a current-PPP series as an alternative growth narrative. That was wrong and has been removed.
2. **Market exchange rates exaggerate the gap.** At market exchange rates, the EU27 economy went from about 11% larger than the US in 2008 to about 31% smaller in 2025. In purchasing-power terms the two economies are almost the same size (EU27 about 98% of the US) [ESTIMATE]. Most of the swing reflects the exchange rate and US price levels. But market exchange rates are the relevant measure for buying things priced in dollars: chips, compute, defence equipment and US software. That is precisely where much AI spending goes.
3. **Ireland.** Multinational activity inflates Irish GDP and therefore EU aggregates. Irish modified gross national income (GNI*) was €334bn in 2025 against GDP of €602bn [FACT] ([CSO](#)). Excluding Ireland lowers EU27 output per hour in 2025 from about 78% to about 77% of the US level on the AMECO constant-2020-PPP basis [ESTIMATE]. (The 80% figure above is a 2024 current-PPP benchmark; the two bases differ.) Part V explains what GNI* does and does not measure.

Why did Europe fall behind Accounting candidate causes and contested explanations

Where the gap shows up: an accounting view, not a causal verdict

The information & communications sector accounts for roughly a third of the 1995–2019 per-hour growth gap in the ECB's decomposition; a broader ICT definition (Draghi, ten EU countries, 2000–19) accounts for most of it. Decompositions show where growth was recorded, not why. ICT use, management, market size and finance are candidate causes.



Sources: left – AMECO Spring 2026, author's calculations (CAGR). Right – ECB Economic Bulletin 6/2024, Box 1 (EU KLEMS): ICT sector 20% (US) / 12% (EA) of per-hour growth 1995–2019; cumulative growth 50% (US) / 28% (EA); contributions derived by author. Samples differ: ECB = euro area, NACE J; Draghi = EU10, C26+J.

Chart 3

Chart 3. Where the gap shows up: an accounting view, not a causal verdict.

Sources: AMECO Spring 2026; ECB Economic Bulletin 6/2024 Box 1 (EU KLEMS); Draghi Part A (2024). Contributions derived by the author.

What the accounting shows [ESTIMATE]. These are decompositions. They show where growth was recorded, not why it happened.

- The US total factor productivity advantage opened in 1995–2005 (US 1.7% a year, euro area 0.8%). It narrowed afterwards (0.6% against 0.4% over 2005–22) ([Bergeaud, ECB Forum 2024](#)).
- The ECB finds that the information and communication sector contributed about 20% of US per-hour growth over 1995–2019, against 12% in the euro area. That is roughly 0.2 points of a 0.7-point annual gap, or about a third.
- Draghi reports that, excluding the main ICT sectors, "EU productivity has been broadly at par with the US in the period 2000–2019": 0.8% a year in the US against 0.6% in a ten-country EU sample.

The two decompositions use different samples and sector definitions, so their shares cannot be reconciled precisely.

Candidate causes: producing digital technology and deploying it well [ESTIMATE] (/CONTESTED). The research points to both.

- **Production.** Europe had smaller ICT-producing sectors and fewer large digital firms (van Ark, O'Mahony and Timmer, 2008).
- **Deployment.** US productivity accelerated "particularly in sectors that intensively use information technologies". US multinationals operating in Europe also got more out of IT than European firms did, which Bloom, Sadun and Van Reenen (2012) attribute to management practices. Schnabel (2024) describes European firms' "failure to reap

the efficiency gains brought about by information and communication technologies" as "one of the root causes". She notes that new technologies spread "to all firms *within* an economy ever more slowly".

The evidence supports the view that **both producing and effectively deploying digital technology mattered**. It does not establish their relative weights. Version 1 said that most of the gap came from depth of use. That goes beyond what the accounting shows, and v2 withdraws it.

Paying for foreign technology is not what caused the gap. Productivity measures what is produced in Europe per hour. Buying software or cloud services from US firms is buying an input; it does not by itself lower European output per hour. The gap arises from how much European firms produced with the inputs they had, and how many high-productivity firms and sectors Europe grew. Payments abroad matter for other questions — income, ownership and control (Part V) — not as an explanation of measured productivity.

Contested explanations (the research dossier gives full sources):

Explanation	Evidence for	Evidence against	Status
Fragmented single market	IMF: internal barriers equivalent to a 44% tariff on manufactured goods and 110% on services; ECB (Jan 2026): 67% and 95%; 34 mobile-operator groups in the EU	Head & Mayer (CEPR, 2025): the same data give 15–32%, near zero with controls; gravity estimates overstate policy barriers	Frictions real; magnitude contested
Capital markets and scale-up finance	EU raises ~5% of global VC vs 52% for the US; ~30% of EU unicorns (2008–21) relocated abroad	Relocation may reflect US market depth; causal share unquantified	Facts solid; causality contested
Energy prices	EU industrial electricity ~1.9× US, gas ~3.5–4.3× (2024–25)	The divergence began in 1995 in low-energy sectors	Weak historically; relevant for AI compute and energy-intensive industry now
Regulation	~100 tech laws and 270 regulators (Draghi); GDPR costs documented	High-regulation Nordic economies are near US per-hour levels; aggregate effects poorly identified	Contested (Part XVI)
Management practices	Bloom et al.: US IT advantage mainly management	Mostly cross-sectional; German and Swedish scores near US	Strong evidence, partial explanation
Firm size and dynamism	More micro-firms, fewer scale-ups	Firm size is itself a result of market size and finance	Symptom and cause

Explanation	Evidence for	Evidence against	Status
Sector composition ("middle-technology trap")	85% of US business R&D is high-tech; ~50% of EU business R&D is mid-tech (autos); 60% of the gap is composition (Fuest et al. 2024)	Specialisation reflects comparative advantage	Accepted as description
Demographics	Workforce shrinking ~2m a year by 2040 (Draghi)	Affects per-capita growth, not the per-hour gap	Future relevance

What falling behind in the ICT cycle meant

"Losing" should not be used as a single number. There is no measured bill for the digital revolution. There are several indicators, each measuring something different.

1. **Growth.** The observed per-hour growth gap of 0.54 points a year over 1995–2025, compounded over thirty years, corresponds to a level difference of about 17%. The ICT-sector part alone (about 0.2 points) corresponds to about 6%. This is arithmetic, not a counterfactual estimate of what Europe would have achieved with different policies [\[ILLUSTRATION\]](#).
2. **Firms and market value.** On 24 September 2026, 3 of the world's 50 largest technology firms by market value were EU-headquartered (ASML, SAP, Schneider Electric), together about 2.4% of the group's value; 32 were American, with about 80% [\[ESTIMATE\]](#) (market data on that date; classification by a third-party ranking).
3. **Income flows for intellectual property.** The EU27's deficit with the US in charges for the use of intellectual property rose from €16bn in 2019 to about €178bn in 2025 (provisional) [\[FACT\]](#) ([Eurostat bop_its6_det](#), EU27 with partner US). Three cautions apply. The jump after 2019 partly reflects multinationals relocating IP and booking flows through Irish subsidiaries. Royalties pay for valuable inputs, so they are not a welfare loss in themselves. The deficit shows *where* valuable IP is owned; it does not measure what Europe lost.
4. **Market structure.** Three US providers hold about 70% of Europe's cloud infrastructure market. European providers' share fell from 29% in 2017 to about 15%, even though their own revenues more than tripled as the market grew about six-fold [\[ESTIMATE\]](#) ([Synergy Research, Jul 2025](#)).
5. **Global revenue share.** Between 2013 and 2023, the EU's share of global technology revenues fell from 22% to 18% while the US share rose from 30% to 38% (Draghi) [\[ESTIMATE\]](#).

What Europe did not lose. Europeans use the same smartphones, search engines and cloud services as Americans, often at similar prices. Much of the welfare gain from *using* ICT accrued to European consumers and firms regardless of where producers were headquartered. What Europe forwent were producer profits, royalty income, some corporate tax base, agglomeration effects and strategic control over infrastructure it now depends on. The size of each is contested.

So what? The ICT precedent is the best available guide to the AI era, with limits. It suggests three lessons. First, productivity growth depended on both producing digital

technology and deploying it well across the economy; the relative weights are debated, but deployment is the part Europe controls most directly. Second, profits and control went to the owners of scarce platforms, which Europe mostly did not build. Third, those ownership effects are harder to measure in GDP than in strategic terms. AI may raise the stakes on all three, but that is a possibility to test, not a settled outcome.

Part III Why AI could be economically different

AI as a general-purpose technology

Economists reserve the term *general-purpose technology* (GPT) for a small class of innovations with three properties (Bresnahan and Trajtenberg, 1995). They are **pervasive**, used across most sectors. They **keep improving**. And they **spawn complementary innovation**: users must redesign products and processes to exploit them. Steam, electricity, the internal combustion engine, the computer and the internet qualify. The question is whether AI does.

The case is strong on the first two criteria. Eloundou and co-authors (2023, published in *Science* 2024) estimate that about 80% of US workers have at least 10% of their tasks exposed to large language models. They also estimate that 47–56% of all tasks could be done significantly faster at the same quality once software is built around the models [MODEL OUTPUT]. Capability continues to improve rapidly. The training compute used for frontier models has grown about fivefold a year since 2020. The price of reaching a given capability level falls by factors between 9 and 900 a year, depending on the task [ESTIMATE] (Epoch AI).

The third criterion, complementary innovation, is where the economic payoff is decided. AI also has a feature earlier GPTs lacked: it may speed up invention itself, including research on AI. If that happens at scale, gains would compound faster than in past technology waves, and an early lead would matter more. This is a serious possibility in the literature, not an established fact [SCENARIO]. Part VII returns to it.

What history teaches and where the analogies break

Technology	Lag from invention to measurable productivity	What had to change
Steam	Negligible contribution to British productivity before about 1830, some 70–100 years after Watt's engine (Crafts 2004)	Railways, high-pressure engines, new industrial locations
Electricity	~40 years from first power stations (1880s) to the US manufacturing productivity surge of the 1920s (David 1990)	Factories redesigned around unit drive rather than a central shaft
Computers	Microprocessor 1971; US productivity boom 1995–2004	ICT capital deepening, reorganised retail, logistics and finance
Internet	Commercial from mid-1990s; gains concentrated in ICT-using services	New business models, platforms, e-commerce

Sources: Crafts, *Economic Journal* 2004; David, *AER P&P* 1990; Oliner & Sichel 2000; Jorgenson, Ho & Stiroh 2008. Historical values are standard literature estimates [ESTIMATE].

Three lessons follow, and each has a limit.

1. **Payoffs lag, sometimes by decades.** Complementary investments in organisation, skills, software and data take time. Brynjolfsson, Rock and Syverson (2021) call this the *productivity J-curve*: measured productivity is *understated* early, because firms are investing in intangibles the statistics do not capture. Adjusting for computer-related intangibles raises US total factor productivity by about 16% relative to official figures by 2017 **[ESTIMATE]**. *Limit*: AI spreads through existing digital infrastructure, so adoption can be much faster than for electricity. Worker-level use of AI in the euro area went from 26% to 52% in two years **[ESTIMATE]** (ECB survey).
2. **Access converges; intensity diverges.** Comin and Hobijn (2010) found that countries adopted technologies on average 45 years after invention, and that differences in adoption lags explain at least a quarter of cross-country income differences. Comin and Mestieri (2018) found that by the twentieth century the lags had converged, but the **intensity of use** diverged, and that this divergence explains much of modern income inequality between countries **[ESTIMATE]**. *Implication*: for AI, most countries will have access; much of the contest will be over depth of use.
3. **Much of the social return accrues to users, not inventors.** William Nordhaus estimated that innovators captured only a few per cent of the social returns to innovation in the US between 1948 and 2001 **[MODEL OUTPUT]**. *Limits*: this is an estimate for one economy, not for the split between countries. And it held when complementary assets were cheap to replicate. In AI, some complementary assets — advanced chips and hyperscale compute — are scarce and concentrated. Part V therefore distinguishes the layers.

The channels through which AI could raise productivity

Channel	Mechanism	Current evidence
1. Augmentation	People do existing tasks faster or better	Strongest evidence: +14% in customer support, +26% completed tasks for developers in field trials
2. Automation	Software and compute substitute for human input in some tasks	Growing: Anthropic data show "directive" (delegating) use of its chatbot rising from 27% to 39% of conversations in 2025; enterprise API traffic about 75% automation
3. Innovation	AI accelerates R&D and discovery (protein structure, materials)	Real in narrow domains; macro effect unquantified. If it extends to AI research itself, growth could compound faster [SCENARIO]
4. Organisational change	Firms redesign workflows around AI	Early; the EIB finds 81% of US AI-using firms apply it in two or more processes, against 55% in the EU
5. New products and industries	Activities that were not possible before	Visible in software (AI coding platforms), media, customer service

Channel	Mechanism	Current evidence
6. Capital deepening	Compute, software and robots add productive capital	Large investment: AI-related investment contributed about 39% of US growth in the first nine months of 2025
7. Knowledge diffusion	Expert-level help becomes cheap and widely available	Largest gains for novices (+34% vs ~0 for experts in customer support)

Sources: Brynjolfsson, Li & Raymond (QJE 2025); Cui et al. (2025); [Anthropic Economic Index](#) (Sept 2025, Jan 2026); [EIB Investment Survey 2025](#); St Louis Fed (Jan 2026). **[ESTIMATE]** throughout.

Micro gains versus macro gains why a 30% saving on a task is not 30% more output

This distinction causes more confusion in political debate than any other. Experimental studies routinely report task-level gains of 15–55%. Yet the macroeconomic estimates are measured in tenths of a percentage point per year. Both can be true, for eight reasons.

1. **Task share.** A task that becomes 30% cheaper matters in proportion to its share of total costs. Economists formalise this as *Hulten's theorem*: to a first approximation, the aggregate productivity gain equals the cost share of the affected tasks multiplied by the cost saving on them. A 30% saving on tasks worth 5% of GDP adds 1.5%, not 30%.
2. **Adoption share.** Only some workers and firms use the tool. In the EU, 20% of firms with ten or more employees used any AI in 2025; only 7% of euro-area firms report *significant* use **[ESTIMATE]** (Eurostat; ECB survey).
3. **Time saved is not output produced.** Saved time must be redeployed into something valuable: more output, better quality or shorter hours. Danish evidence (Humlum and Vestergaard, 2025) found that most users reallocated saved time to other tasks, and ruled out effects on earnings or hours larger than 2% two years after ChatGPT **[ESTIMATE]**.
4. **Bottlenecks (the Baumol effect).** When some tasks become much cheaper and others do not, the economy becomes limited by the slow tasks, and spending shifts towards them. Chad Jones (2026) shows that with strong complementarity between tasks, automating half of the economy's tasks raises output by only about 19% in total **[MODEL OUTPUT]**.
5. **Quality and errors.** Outside the "jagged frontier" of what AI does well, consultants using GPT-4 were 19 percentage points less likely to reach the right answer (Dell'Acqua et al., 2023). Experienced open-source developers were 19% *slower* with AI tools in a 2025 trial (METR) **[ESTIMATE]**. Checking AI output takes time.
6. **Reallocation takes time.** Displaced tasks and workers must find new uses. Search frictions and rigid wages create temporary unemployment.
7. **Complementary investment.** Data infrastructure, redesign and training are costly, and they depress measured productivity early (the J-curve). Tool subscriptions and compute are costs too.

8. **Measurement.** Free chatbot services create consumer value that GDP largely misses. AI investment shows up first as demand for construction and equipment, not as productivity.

A simple worked example [ILLUSTRATION]. Suppose AI cuts by 30% the time needed for tasks that make up 25% of working time, for 40% of workers. Total working time needed for the same output falls by about 3% ($0.25 \times 0.40 \times 0.30$). If the saved time is redeployed into equally valuable work, output rises by roughly 3%. Spread over ten years, that is about 0.3 percentage points a year — close to where the moderate macro estimates sit.

Two cautions. First, "30% faster" is not the same as "30% less time". If workers produce 30% more per hour on those tasks, the time needed falls by about 23% ($1 - 1/1.3$), and the example gives about 2.3%. Second, this is a gross, first-order calculation. The net gain is smaller once the costs of tools, verification and reorganisation are deducted, and it depends on whether saved time becomes more output, better quality or shorter hours.

So what? When a minister hears that "AI makes workers 40% more productive", the right follow-up questions are: on what share of the job, for what share of workers, with what error rate, at what cost, and redeployed into what? The answers turn dramatic micro numbers into modest macro numbers. Compounded over a decade or two, modest numbers are still significant.

Part IV How big could the productivity effect be

The serious estimates

The literature does not converge on one number. It produces a range, driven by five assumptions: **capabilities** (what AI can do), **adoption** (how widely and deeply it is used), **autonomy** (how much it automates rather than assists), **productivity per task**, and **labour-market adjustment** (whether displaced workers move to equally productive work).

The estimates are also published in **three different units**, which must not be ranked against each other:

- **TFP growth** — efficiency after accounting for measured inputs.
- **Labour-productivity growth** — output per hour or per worker, which also rises when capital per worker rises.
- **GDP level** — how much higher output is than on a no-AI path at a given date. A level effect spread over several years is not the same as an annual productivity gain, especially if employment changes.

Source (date)	Measure, geography	Estimate	Key assumptions
Acemoglu, <i>Economic Policy</i> (2025)	TFP and GDP level, US	[MODEL OUTPUT] TFP +0.66% over 10 years (~0.07 pp/yr); GDP +0.9–1.2%	~20% of tasks exposed, 23% of those profitably automated, 27% cost saving; no R&D effect
Penn Wharton Budget Model (Sept 2025)	GDP level; TFP, US	[MODEL OUTPUT] GDP +1.5% by 2035, +3.7% by 2075; peak TFP contribution 0.2 pp (2032)	40% of labour income exposed; adoption like PCs/internet
IMF, Misch et al. (Apr 2025)	TFP, Europe	[MODEL OUTPUT] ~1% cumulative over 5 years	Acemoglu framework, European exposure. A scenario with 50% lower AI exposure in regulated tasks and sectors cuts the gain by >30%
ECB (Lane) (Mar 2026)	TFP, euro area, 10 years	[SCENARIO] ~0.2 pp/yr (slow adoption); 0.3–0.4 pp/yr (fast)	40–60% of GDP exposed; ~40% gains on exposed tasks
OECD, Filippucci et al. (Nov 2024)	TFP and labour productivity, US, 10 years	[MODEL OUTPUT] TFP +0.25–0.6 pp/yr; labour productivity +0.4–0.9 pp/yr	Micro gains × exposure × adoption; Baumol drag
OECD, Schief et al. (Jun 2025)	Labour productivity, G7, 10 years	[SCENARIO] +0.4–1.3 pp/yr in the US and UK; "up to 50% smaller" in several other G7 economies	Sector mix and adoption speed

Source (date)	Measure, geography	Estimate	Key assumptions
OECD, Yildirim, "AI meets trade" (Mar 2026)	Real income per head, many countries	[MODEL OUTPUT] +0.1–0.95 pp/yr, varying widely by country	Multi-country trade model; lagging adopters gain from cheaper imports but need domestic adoption to stay competitive in AI-exposed sectors
Aghion & Bunel (2024)	TFP, 10 years	[MODEL OUTPUT] median +0.68 pp/yr (range 0.07–1.24)	Larger exposure and cost savings than Acemoglu
Anthropic Economic Index (Nov 2025 → Jan 2026)	Labour productivity, US	[MODEL OUTPUT] 1.8 pp/yr, revised to 0.7–0.9 pp after adjusting for task success and complementarity	Universal adoption within 10 years; current capabilities; published by an AI developer
Goldman Sachs (2023)	Labour productivity, US; GDP level, world	[CONSULTANCY] US +1.5 pp/yr after full adoption; world GDP +7% over 10 years; time-phased +0.3–0.4 pp by 2034	Two-thirds of occupations exposed; full reallocation
McKinsey (2023; 2026 Europe update)	Labour productivity	[CONSULTANCY] Gen-AI +0.1–0.6 pp/yr to 2040; 58% of European work hours technically automatable	Technical potential × adoption scenarios
Anthropic Institute, Korinek, Jones et al. (Sept 2026)	GDP level, US	[SCENARIO] +1.6% / +8.3% / +32% by 2030 vs no-AI path (modest / substantial / extreme)	Task-based model; excludes policy responses; published by an AI developer
Epoch AI "GATE" model (2025)	GDP growth rate	[SCENARIO] Explosive growth (~30% a year by ~2035) in default runs	Full automation reachable with enough compute


Two recent reality checks matter. In March 2026 Goldman Sachs analysts wrote that they "still do not find a meaningful relationship between productivity and AI adoption at the economy-wide level" **[CONSULTANCY]**. And the US labour-market evidence remains mixed. There is no aggregate disruption so far (Yale Budget Lab, 2025). But employment of 22–25-year-olds in the most AI-exposed occupations is about 19% below the trend of less-exposed peers ([Brynjolfsson, Chandar and Chen](#), revised August 2026) **[ESTIMATE]**. That is a *relative* shortfall for one group, not a 19% fall in total employment.

A scenario range not an average

These estimates use different units, horizons and geographies. Averaging them would be meaningless. Chart 11 therefore shows each estimate in its own unit.

Published AI productivity estimates, kept in their own units

TFP, labour productivity and output levels are different quantities and cannot be ranked on one axis. Mainstream estimates are modest over a decade; large level effects come from scenarios conditional on capabilities not yet demonstrated.



Sources: Acemoglu, Economic Policy (2025); PWB (8 Sept 2025); ECB (Lane, 23 Mar 2026); OECD AI Papers (Filippucci et al., Nov 2024; Schief et al., Jun 2025); Aghion & Bunel (2024); McKinsey (2023); Anthropic Economic Index (Jan 2026); Goldman Sachs (Mar 2023); IMF WP/25/067 (*~1% cumulative TFP gain over five years, shown as ~0.2 a year); Anthropic Institute WP 2026-02 (US, 2026-30).

Chart 11

Chart 11. Published AI productivity estimates, kept in their own units: TFP growth, labour-productivity growth and output levels.

Only ranges published by the source are drawn as ranges; single estimates are shown as points. Version 1 drew narrow intervals around some point estimates (for example 1.45–1.55 around Goldman's 1.5); those intervals were not documented and have been removed. The IMF's ~1% cumulative TFP gain over five years is shown as about 0.2 points a year. Source table: workbook sheet C11.

The scenarios below are defined by their assumptions. The published estimates that fall within each are listed in their own units.

Scenario	Capabilities	Adoption	Autonomy	Labour market	Published estimates of this kind
Low impact — "useful software"	Roughly today's	Slow, limited by profitability; ~5% of tasks affected	Mostly assistance	Little displacement; little gain	Acemoglu TFP ~0.07 pp/yr; Penn Wharton peak TFP 0.2 pp; ECB slow ~0.2 pp TFP
Moderate — "internet-like"	Steady improvement	30–60% of firms by mid-2030s	Assistance and some task automation	New tasks offset part of displacement	ECB fast 0.3–0.4 pp TFP; OECD US 0.4–0.9 pp labour productivity

Scenario	Capabilities	Adoption	Autonomy	Labour market	Published estimates of this kind
High — "electricity-like, fast diffusion"	AI can do about half of knowledge work by ~2030	Rapid, agent-based	High automation of cognitive tasks	Weaker offset; some unemployment	OECD US/UK up to 1.3 pp labour productivity; Goldman 1.5 pp at full adoption; Anthropic Institute "substantial" GDP +8.3% by 2030 (US)
Transformative — near-full automation of cognitive work	AI better than humans at most cognitive tasks, possibly accelerating research	Very rapid	Near-full	Labour share falls sharply; outcome depends on policy	Anthropic Institute "extreme" GDP +32% by 2030 (US); Epoch GATE explosive growth

Mapping by the author **[INFERENCE]**. The transformative scenario is included because serious economists model it, not because current evidence makes it likely.

How to read this without false precision. Version 1 attached numerical probabilities to these outcomes (for example "~55%" for at least 0.3 points a year in the EU). They had no documented method or calibration record and have been removed. The evidence supports three qualitative statements. Mainstream estimates for Europe over the next decade sit in the low-to-moderate range. The high and transformative scenarios depend on capabilities and adoption speeds not yet observed. The single largest uncertainty for Europe is the speed and depth of organisational change in firms and public services, not the technology alone.

Indicators that would change this assessment:

- Measured productivity growth in the sectors with the highest AI use, compared with less-exposed sectors, in national accounts rather than surveys.
- The share of firms using AI in several business processes (the EIB's depth measure), tracked over time and by country.
- Evidence that AI is shortening research cycles outside narrow domains, including in AI research itself.
- Employment and wages for exposed occupations, beyond the early-career effects seen so far.


What being early or late could mean

The OECD's G7 work is the clearest quantitative anchor. With identical technology, differences in sector mix and adoption speed alone make the gains in several G7 economies "up to 50% smaller" than in the US and UK **[SCENARIO]** (OECD 2025). Lane's two ECB scenarios imply a difference of roughly 1–2 percentage points in the euro-area TFP level after ten years between slow and fast adoption **[SCENARIO]** (author's arithmetic from 0.2 vs 0.3–0.4 pp a year).

What that means for national income can be shown with simple arithmetic.

Early or late: an arithmetic illustration of AI diffusion

Same technology, different speed and depth of use. Baseline growth 1% a year. A late start leaves a permanent level gap even when the growth rate catches up; a smaller increment keeps widening the gap. Not a forecast for any country.



Source: author's arithmetic [ILLUSTRATION]. Increment sizes informed by the range in OECD (Nov 2024; Jun 2025) and ECB (Mar 2026). Output per worker; ignores changes in hours, capital shares, terms of trade and where income is received.

Chart 12

Chart 12. Early or late: an arithmetic illustration of AI diffusion. Not a forecast.

Assume a baseline of 1% annual productivity growth and three economies with the same technology. The increments are constant, start in the stated year and are chosen from within the published range [ILLUSTRATION]:

- **A — early and deep adopter:** AI adds 0.8 points a year from 2027.
- **B — shallow adopter:** AI adds 0.4 points a year from 2027 (half of A, echoing the OECD's "up to 50% smaller").
- **C — late starter:** AI adds 0.8 points a year, but from 2032.

By 2035, A's output per worker is 3.6% higher than B's and 4.0% higher than C's. By 2045 (after 19 years of AI effects), A is 7.8% above B, while the gap to C stays at 4.0%. A late start creates a permanent level gap even when the growth rate catches up; a smaller increment keeps widening the gap [ILLUSTRATION].

Fiscal translation, with its conditions. For Sweden, a 7.8% lower output level, taxed at the 2025 tax ratio of 40.5%, means gross public revenue lower by about 3.2% of the lower-output economy's GDP ($0.078 \times 0.405 \approx 0.032$) [ILLUSTRATION]. That is of the same order as planned annual defence spending. Three conditions matter. The calculation assumes the extra output is taxed at the same average rate, although Part VI shows AI may shift income towards forms taxed differently. It does not raise the tax-to-GDP ratio; it changes the size of the base. And gross revenue is not money available for new programmes, because public-sector wages, pensions and demand for services also rise with the economy. For the EU27, a 7.8% level difference is roughly €1.5 trillion of annual output at today's size.

So what? Under mainstream assumptions, the difference between an early, deep adopter and a shallow or late one is not a transformation of living standards within five years. It is a slowly widening gap that, within twenty years, becomes as large as the budget of a major national priority. In the illustration, how deeply AI is used matters more over time than when use begins. That result follows from the assumed increments; it is a reason to measure depth, not proof of it.

Part V Productivity is not the same as value capture

This is the part of the argument most often skipped in political debate, and one of the most important for Europe. Version 2 separates five things that version 1 sometimes merged: **value creation, income capture, ownership, taxation and control**.

Different things gaining from AI can mean

Concept	What it measures	Where it is recorded
Consumer surplus	Value to users above what they pay	Not in GDP; accrues where users live
Productivity	Output per hour	Where the work is done
GDP	Value added produced on a territory	Where production is located, whoever owns it
GNI / national income	Income received by residents, including from abroad	Where the recipients are resident
Profits	Returns to firms	Where firms book them (often not where they operate)
Intellectual property	Rights to royalties and licences	Where IP is legally held
Capital income and gains	Dividends, interest, retained earnings, capital gains	Where shareholders are resident; capital gains change wealth but are not part of GNI
Tax revenue	Public share of income	Where wages, profits and consumption are taxed
Ownership	Claims on firms: controlling stakes or portfolio holdings	Where owners are resident
Strategic control	Power to set prices, terms, access, updates and security features	Where the firm is operated and under whose law

A country can gain on some of these and not on others.

Ireland shows why GDP, resident income and the tax base differ. In 2025 Irish GDP was €602bn and modified gross national income (GNI*) was €334bn **[FACT]** (**CSO**). Version 1 treated the difference as foreign profit. That was too simple. The gap has two parts:

- **Net factor income paid abroad, about €173bn.** This takes GDP to GNI. It is mostly the profits, including retained earnings, of foreign-owned multinationals.
- **About €95bn of further adjustments that take GNI to GNI* (components rounded).** Depreciation on intellectual property and R&D imports (about €69bn), depreciation of leased aircraft (about €10bn) and the income of redomiciled companies (about €16bn). The CSO removes these because they have little effect on Irish residents' income (**CSO definition**).


Foreign ownership also does not mean an absent tax base. In 2024, foreign-owned multinationals paid €24.8bn in Irish corporation tax, 88% of net receipts excluding the one-off CJEU settlement. The ten largest groups paid close to 59% **[FACT]** ([Irish Revenue, CT analysis 2025, pp. 8 and 11](#)). The lesson is that a territory can host large foreign-owned production, collect substantial tax on it, and still have resident income far below GDP. Each channel must be measured separately.

A worked example a Swedish manufacturer uses American AI

Suppose a Swedish engineering firm deploys an American AI model through an American cloud and becomes more productive. Who gains? Chart 10 sorts the parties by **where they are resident**, not by the supplier's nationality, and separates income from control.

Using foreign AI: who receives income, who owns, who controls?

A stakeholder map, not a measured split. Many parties can gain at once, in Sweden and abroad. How gains are divided depends on competition, bargaining power, ownership, where activity is located and tax rules. No study quantifies the national split for business use of AI.



CONTROL: who can change the tool's price, terms, availability or data access? Mainly the supplier and its home state, unless limited by contract, portability, open-weight alternatives or European-operated deployment. Control is separate from where income is received.

CONCEPTUAL DIAGRAM – boxes are not proportional and are not slices of one total: cloud fees pay for chips, and taxes are paid out of income already counted. Residency follows national-accounts logic (where the recipient is resident), not the supplier's nationality.

Chart 10

Chart 10. Using foreign AI: who receives income, who owns, who controls? A stakeholder map, not a measured split.

Version 1's diagram placed foreign customers on the "stays in Sweden" side and put Swedish-hosted energy and data-centre activity on the "flows abroad" side. It also used arrow widths that implied a measured split. This version corrects both.

- **The Swedish firm and its owners** gain if productivity rises faster than competitors can match. Its owners may be Swedish or foreign.
- **Its employees** gain if productivity gains are shared through wages. Some tasks may shrink or change.
- **Its customers**, in Sweden and in export markets, gain through lower prices and better products. Much of this is consumer surplus.

- **The Swedish state** gains taxes on wages, profits and spending, including taxes on foreign suppliers' local activity.
- **A US-owned cloud operator hosting the computing in Sweden** creates local construction, jobs, power purchases and taxable activity, while its profits accrue mainly to its foreign owners.
- **The US model provider** earns subscription and API fees, and controls the model's terms, updates and behaviour.
- **Chip and equipment makers** gain from the hardware demand: Nvidia and TSMC abroad, and ASML inside the EU.
- **Shareholders of the suppliers** gain, including Swedish residents through pension funds. They receive income and gains but no operational control.

How large is each share? This report does not establish a reliable economy-wide division of AI-created value between adopting firms, their workers and customers, suppliers and states. Version 1 went beyond this by concluding that "most of the gain stays with the adopter" or "in Sweden". That conclusion is withdrawn.


The consumer evidence is real but answers a different question. Brynjolfsson, Collis and co-authors estimate US consumer surplus from generative AI at \$116bn in 2025 and \$172bn in 2026, using willingness-to-accept experiments [MODEL OUTPUT] ("[What is Generative AI Worth?](#)", Apr 2026). Compared with roughly \$14bn of global consumer revenue, this implies that vendors captured about 5–11% of consumer value. That is an upper bound, because the revenue is global and the surplus is US-only. The study measures consumer welfare. It does not measure supplier profits, domestic GDP, tax receipts, or the business or national split.

The accounting boundary. The parties in Chart 10 cannot be added up as slices of one total. Cloud fees partly pay for chips. Taxes are paid out of wages and profits already counted. Consumer surplus is not income at all. Any later numerical version needs an explicit boundary and should show net gains, not gross supplier bills.

Supplier profitability today. Some suppliers are currently very profitable. Chart 9 shows selected examples, each for its own period and entity.

Supplier profitability: selected examples, not an 'AI profit pool'

Scarce hardware and cloud suppliers are currently very profitable; the leading model lab is not. These are separate examples with their own periods. Margins cannot be added, do not measure how AI's gains are shared, and do not show whether today's profitability will last.



Sources: Nvidia Q2 FY2027 release (revenue US\$96.2bn, operating income US\$63.7bn); TSMC Q2 2026 6-K (SEC); Amazon and Alphabet Q2 2026 releases; ASML 2025 annual results; OpenAI: 2025 revenue ~US\$13.1bn and operating loss ~US\$20.9bn from reported leaked accounts (not audited). Firms buy from each other: cloud fees pay for chips. Chart 9

Chart 9. Supplier profitability: selected examples, not an "AI profit pool".

Version 1's Chart 9 mixed companies and cloud segments, operating income and ASML's net income, and one quarter with a quarter of OpenAI's annual loss, on a common axis titled "where the profits are". It has been rebuilt as separate examples using operating margin, with each period and entity type (company or segment) stated; the OpenAI figure is a leaked estimate.

In the quarter to 26 July 2026, Nvidia earned operating income of \$63.7bn on revenue of \$96.2bn, an operating margin of about 66% [FACT] (Nvidia Q2 FY2027). TSMC's operating margin was about 60% in Q2 2026. AWS's was about 39% and Google Cloud's about 36% [FACT] (company filings). OpenAI's 2025 operating loss was reported at \$20.92bn on revenue of \$13.07bn, according to leaked audited financial statements obtained by Ed Zitron and independently reviewed by the Financial Times [ESTIMATE]. The documents are not publicly filed, so the figures remain a reported estimate. Four US hyperscalers spent about \$416bn on capex in 2025 and guided roughly \$650–750bn for 2026, far ahead of AI revenue [FORECAST]. The Bank for International Settlements, the IMF and the Bank of England all warned in 2026 of correction risk [FACT].

These margins show where scarcity and pricing power sit *today*. They reflect scarcity, product quality, intellectual property, switching costs and an investment boom as well as scale economies. They do not show how AI's gains are shared across the economy, and they do not prove that today's profits will last (Part VII).

Three conclusions [INFERENCE]:

1. **Adoption is not the same as domestic value capture, and it is not a transfer abroad either.** A Swedish firm using American AI can create income for Swedish workers, customers, shareholders and the state while paying foreign suppliers. Both countries can gain. Control over the tool may still sit under foreign law.

2. **Refusing to adopt is unlikely to protect European income.** The OECD's multi-country trade model finds that lagging adopters gain from cheaper AI-intensive imports, but that "maintaining competitiveness — especially in highly AI-exposed sectors — ultimately requires strong domestic AI adoption" [MODEL OUTPUT] ([OECD, Mar 2026](#)).
3. **The split is not fixed.** It depends on competition, costs, bargaining power and ownership. If a few suppliers gain pricing power — through lock-in, bundling or the collapse of open alternatives — the share flowing to them rises. Competition, portability and credible alternatives protect users' share.

Which layers matter for European benefit and control

National benefit arrives through domestic production, resident income, taxable activity and operational control. These channels need separate measurement. The evidence suggests different roles for different layers [INFERENCE]:

- Adoption capacity — skills, management, data and organisational change — determines how much value users can create locally, regardless of where a model was developed.
- Critical functions raise a different question: which services could continue for a defined period if a supplier's access, licence or updates were interrupted? The answer depends on the function and its tested fallback.
- European strengths in lithography, industrial data, domain software and power equipment show that value can accrue inside a global stack without owning every layer.
- Frontier models and general-purpose cloud remain contestable to different degrees. Competition, portability, cost and the future value of a frontier lead determine how much supplier power users face.

The implication. "Europe can use foreign AI" is true but incomplete. Effective use can create local value; supplier ownership may bring income and control elsewhere. Profits, tax and operational dependence do not follow one simple national ledger. Their balance is an empirical question, not a reason to reject foreign technology.

Part VI AI capital and the distribution of economic power

Will AI raise the importance of capital relative to labour

What substitution means. AI can substitute software and compute for some paid human tasks. That can shift income and bargaining power from workers towards the owners of capital. It is not a literal transfer of a fixed stock of human capital into machines. Nor does a smaller wage bill automatically become profit for a model provider. The saving can flow to lower prices for customers, higher profits for the adopting firm, higher wages for workers whose remaining tasks become more valuable, or fees to suppliers. Which happens depends on competition in each market.

Three different labour-market questions. *Task automation* asks which tasks machines take over. *Job displacement* asks which workers lose jobs or pay as a result. *Economy-wide employment* asks whether total employment and wages fall once new tasks and reallocation are counted. Firm-level evidence mostly answers the first two. Macro scenarios answer the third, and only under their assumptions.


Every model of transformative AI reviewed for this report predicts a falling labour share of income [MODEL OUTPUT]:

- The Anthropic Institute's scenario paper (Korinek, Jones and co-authors, September 2026) models the US economy to 2030. Labour's share falls from 60% to 59.4% (modest), 56.1% (substantial) and 45.2% (extreme). In the extreme case GDP is 32% higher, knowledge workers' wages are 11.5% *lower* than on the no-AI path, and other workers' wages rise [SCENARIO] (WP 2026-02; US only; policy responses excluded; published by an AI developer; values verified against the paper and official interactive explorer).
- The IMF (2024) finds capital returns rise in all its scenarios. Trammell and Korinek (2026) show full automation lowers the labour share and that wages "may rise or fall". Acemoglu (2025) expects AI to widen the gap between capital and labour income [MODEL OUTPUT].

A lower share is not necessarily a lower income. If output rises from 100 to 120 while labour's share falls from 60% to 55%, labour income rises from 60 to 66 [ILLUSTRATION]. In the Anthropic Institute scenarios, total labour income stays within about $\pm 1.5\%$ of the no-AI path in every case (author's approximation), while its *distribution* shifts sharply between occupations (Chart 13).

A smaller share of a larger economy: labour income in AI scenarios

In one US model, labour's share of income falls in every scenario. Because output rises, total labour income stays close to the no-AI path – but the composition shifts: in the extreme scenario knowledge workers' wages are 11.5% lower while other workers' wages rise.



Source: Korinek, Jones, Sacher, Cotter & McCrory, 'Economic Scenarios for Transformative AI', Anthropic Institute WP 2026-02 (Sept 2026) [SCENARIO, US only, 2026-30; excludes policy responses]. Right panel: author's approximation (output index × labour share ÷ 60%). The publisher is an AI developer.

Chart 13

Chart 13. A smaller share of a larger economy: labour income in AI scenarios (US model, 2030).

Right-hand panel: author's approximation (output index × labour share ÷ 60%), not a model output.

The data today are consistent with, but do not prove, a shift towards capital. The US non-farm labour share fell to 52.8% in Q2 2026, the lowest since the series began in 1947 [FACT] (BLS). After-tax corporate profits reached 12.4% of GDP in Q1 2026 [FACT]. The decline predates generative AI and is linked in the literature to "superstar" firm concentration (Autor et al., 2020). AI investment was about 39% of US GDP growth in the first nine months of 2025 (St Louis Fed), and AI-related companies make up about half of the S&P 500's value (Bank of England) [ESTIMATE].

If capital earns more who owns the capital

AI-related capital	Main owners	Where returns are booked
Accelerators and chips	Nvidia, AMD, Broadcom (US); TSMC (Taiwan); SK Hynix, Samsung (Korea)	US, Taiwan, Korea
Lithography	ASML (Netherlands)	Netherlands
Hyperscale compute and cloud	Amazon, Microsoft, Google, Meta, Oracle (US); five US firms own ~71% of global AI compute (Epoch estimate)	US, partly via Irish and other subsidiaries
Frontier models and IP	US labs; Chinese labs; Mistral (France)	Mostly US
Data centres in Europe	Largely US hyperscalers and global infrastructure funds	Hosting countries get construction, power sales, local employment and taxable activity

AI-related capital	Main owners	Where returns are booked
Industrial software and robotics	Siemens, SAP, Dassault (EU); robot makers now mostly Asian-owned (KUKA → Midea; ABB Robotics → SoftBank)	Mixed

Sources: [Epoch AI \(Apr 2026\)](#); company filings; Part VIII. **[FACT]/[ESTIMATE]**

Europeans own a meaningful share of US AI capital. Foreign investors held \$19.9 trillion of US equities in June 2025, of which European residents held \$9.6 trillion — roughly 9% of the US stock market **[FACT]** ([US Treasury TIC](#); ESTIMATE for the share, custodial basis). Norway's sovereign fund has 53% of its portfolio in the US and reported that US technology stocks contributed most to its 2025 return. Swedish investors held about \$396bn of US equities **[FACT]**.

What portfolio ownership does and does not give.

- **It gives** residents a claim on dividends and capital gains. Dividends enter national income. Capital gains raise wealth but are not counted in GNI. Depending on the vehicle, this income can be taxed at home when it is received, realised or paid out as pensions.
- **It does not give** operational control over products, pricing, data or security decisions.
- **It does not move** the issuer's corporate-tax base. Where multinational profits are taxed depends on where activity and profits are located. The January 2026 OECD "side-by-side" package creates safe harbours for multinational groups headquartered in eligible jurisdictions, designed to accommodate the US system, for fiscal years beginning on or after 1 January 2026. It is not a general exemption: qualified domestic minimum top-up taxes remain the primary mechanism for profits in countries that impose them **[FACT]** (OECD, 5 January 2026).
- **It does not protect** against a valuation correction, against which central banks warned in 2026.

Testing the proposition

If AI increases the return to capital, does the geographic ownership of AI-related capital become more important to national prosperity?

- **Plausibly yes, in three respects** **[INFERENCE]**. *Control*: owners decide how AI is priced, updated and deployed. *Agglomeration*: headquarters attract talent, suppliers and start-ups. *The fiscal mix*: European welfare states rely heavily on taxes on labour income; European labour shares are 6–7 points higher than in the US on comparable definitions. If AI shifts income towards capital, labour-dependent fiscal systems may collect a smaller share of a larger economy unless tax policy adapts. Whether revenue actually falls depends on growth, taxable activity, ownership, tax rules and distribution.
- **But less than a mercantilist view assumes.** Returns can be shared through diversified portfolios. Competition can hold down rents in contestable layers. Foreign-owned activity located in Europe creates wages and taxable profits. And the gains from use are available to users regardless of the supplier's nationality.

- **The conditions that decide it:** whether AI rents persist or are competed away, and where the profitable activity is located. Where rents persist and activity is located abroad, ownership and location matter more. Where rents are competed away, effective use matters more than ownership.

So what? For a Swedish or European politician, the capital-share question is partly a fiscal question and partly a question of who can steer the technology. A shift of income from wages to capital would not automatically shrink the tax base, because the economy could grow. But it would change what is taxed and who benefits, and would make the location of profitable activity and the design of capital taxation more important. Pension funds can spread households' exposure to global returns; they do not give the state operational control.

Part VII The flywheel of technological leadership

The hypothesis

The "whoever wins, wins" logic depends on a self-reinforcing loop: **better models** → **more users** → **more revenue** → **more capital** → **more compute** → **more talent** → **more data and feedback** → **more R&D** → **better models**. Add cloud, chips, energy, capital markets, universities, start-ups and defence procurement, and the loop becomes national.

Where concentration and scale economies are visible today

Layer	What is observed	Current reading
Accelerators	Nvidia: ~66% operating margin; >60% of installed AI compute; CUDA software ecosystem	Strong current position; durability depends on rival chips and software
Fabrication	TSMC: ~92% of sub-10nm capacity (2019); \$60–64bn annual capex no rival matches; ~60% operating margin	Strong; high entry cost
Lithography	ASML: sole EUV supplier; ~54% gross margin	Strong; a European chokepoint
Cloud	Three firms hold ~70% of Europe's market; operating margins 36–39%; high switching costs (the Commission preliminarily proposed designating AWS and Azure as DMA gatekeepers in June 2026)	Strong; switching rules and competition law may change it
Distribution / platforms	Microsoft 365 Copilot >30m paid seats; Gemini app ~950m monthly users; AI bundled into existing ecosystems	Strong
Agglomeration	US private AI investment \$286bn in 2025, 23× China; ~59% of top AI researchers work in the US	Strong, possibly eroding (flows of AI scholars to the US reportedly down 89% since 2017)
Scale of compute	Frontier training compute ~5× a year; training costs ~3.5× a year; largest runs >\$1bn by 2027	High entry barrier for frontier labs

Sources: company filings (Q2 2026); SIA/BCG (2021); Synergy Research; Epoch AI (2024–26); [Stanford AI Index 2026](#); European Commission (IP/26/1444). [FACT]/[ESTIMATE]

High margins and concentration are consistent with scale economies and market power. They also reflect scarcity during an investment boom, product quality, intellectual property and switching costs. They do not by themselves prove durable increasing returns in every physical layer [INFERENCE].

Where the model-layer lead looks contestable and what that does not prove

- **Leadership among firms shifts.** Among enterprise buyers of model APIs, Anthropic's spending share rose from 12% (2023) to 40% (2025) while OpenAI's fell from 50% to 27% (Menlo Ventures survey) [ESTIMATE]. Switching costs between model vendors are low. But this is movement *between US suppliers*. It does not show that the US national lead is weak.
- **Open-weight and Chinese models follow within months.** Epoch's November 2024 study put the average lag of open-weight models at about a year; its newer capability-index comparison puts it at around four months. Chinese models have trailed the US frontier by about seven months on average since 2023 (range 4–14) [ESTIMATE] (Epoch AI). Qwen overtook Llama as the most-used open model family on Hugging Face in 2025.
- **Prices fall fast.** The price of reaching a given capability fell by factors of 9 to 900 a year depending on the benchmark [ESTIMATE].
- **Data network effects look weaker than for social platforms.** Theory and evidence (for example Hagiu and Wright, 2023) suggest learning from user data creates advantages that are real but contestable.
- **The DeepSeek test.** In January 2025, a Chinese model trained with far less disclosed compute wiped about \$600bn off Nvidia's value in a day. Nvidia's revenue subsequently more than doubled: cheaper capability increased demand for compute. Commoditisation of models coexisted with rising profits on chips [FACT]/[INFERENCE].

What this does not show. Short gaps do not mean catching up is cheap: followers benefit from knowing what is possible, and the leading labs spend heavily to stay ahead. Nor do short gaps mean a temporary lead has little value. Usage at the frontier appears price-inelastic — a 10% price cut raises usage by only 0.5–0.7% on a large developer platform — suggesting the best models are "not yet a commodity" [MODEL OUTPUT].

The research-acceleration possibility. The strongest reason a model lead could matter more than today's evidence suggests is that frontier AI might accelerate science and AI research itself. If the leading systems shorten research cycles — in chip design, materials, biology, software or the next generation of models — a lead of months could compound into a larger one. The US Genesis Mission (Part 0) is built on this bet. Current evidence shows real gains in narrow domains such as protein structure; it does not yet show broad acceleration. The Anthropic Institute model, which includes AI in R&D, finds "the growth speedup from AI in innovation is relatively minor, even in the extreme change scenario" [MODEL OUTPUT]. This is a live uncertainty, not a prediction [SCENARIO].

Verdict

Current concentration is strongest in the physical and distribution layers; the model layer is more contestable today [INFERENCE]. Whether any of these positions become durable rents is uncertain. Physical-layer positions face rival chips, new fabs and competition law. Model-layer positions could become much more valuable if frontier AI accelerates research.

Version 1 attached probabilities (85% and 70%) to these propositions. They had no documented method and have been removed. **Indicators that would change the assessment:** the market share of non-Nvidia accelerators in new capacity; the capability gap between the best open-weight and closed models; evidence that frontier systems shorten research cycles outside narrow domains; and cloud switching behaviour once the Data Act's charge removal applies in January 2027.

So what? For Europe, the dependence risk is clearest in the layers where concentration is visible today — compute, cloud and platforms — because they determine prices, access and legal reach. The model layer is contestable today, but that could change if frontier AI starts to accelerate research. Europe's strongest current asset in this picture is its position as an indispensable supplier inside other countries' value chains, above all through ASML.

Part VIII US China and Europe where is power actually located

This part maps the AI stack component by component. It deliberately does not construct a composite "AI power score". Weights would be arbitrary, and the point is that power differs by layer. Where a measure covers only part of the world (a sample, a region or a type of spending), the table says so.

The map

Component	United States	China	EU27 / Europe	UK and others	Source, date
Frontier capability	Every frontier model since 2023 (Epoch Capabilities Index)	~7 months behind on average (range 4–14)	Mistral near-frontier, not at frontier	—	Epoch Jan 2026
Notable models, 2025	59	35	1 (France)	UK 1; S. Korea 8	AI Index 2026
Large-scale models (>10 ²³ FLOP) released 2025	66	82	6	UK 0	Epoch via OWID, Sept 2026
Private AI investment, 2025	\$285.9bn (83% of world)	\$12.4bn (excludes state guidance funds, ~\$184bn 2000–23)	France \$4.4bn, Germany \$4.1bn; "Europe" incl. UK ~\$21bn	UK \$5.9bn	AI Index 2026
Newly funded AI companies, 2025	1,953	161	n/a	UK 172	AI Index 2026
Hyperscaler capex, 2026 guidance	~\$650–750bn (Amazon, Alphabet, Meta, Microsoft; all uses, global)	n/a	No comparable private actor	—	Company filings
Cloud share of European market	~70% (AWS, Microsoft, Google)	small	~15% (down from 29% in 2017; revenues more than tripled)	—	Synergy Jul 2025
Location of AI-supercomputer performance, in Epoch's dataset (covers an estimated 10–20% of global capacity)	~75%	~15%	inside ~10% "others"	—	Epoch Apr 2025


Component	United States	China	EU27 / Europe	UK and others	Source, date
Ownership of global AI compute	Five US hyperscalers ~71% (Q4 2025), including capacity abroad	—	—	—	Epoch Apr 2026
Data-centre electricity, all workloads, 2024 → 2030	45% of 415 TWh → ~425 TWh	25% → ~280 TWh	Europe 15% → ~11% share	—	IEA Energy and AI 2025
Leading-edge chip fabrication	~0% (2019), ~28% projected 2032	0% leading edge; SMIC 5–7nm-class	~9–10% of all capacity, none leading edge; Chips Act target 20% "very unlikely"	Taiwan 92% of <10nm (2019)	SIA/BCG; ECA 2025
AI accelerators	Nvidia >60% of AI compute; data-centre revenue \$89bn in one quarter	Huawei <4% of Nvidia's compute output (2026)	none at scale	—	AI Index 2026; Epoch Sept 2026
Semiconductor equipment	~41% of value added	growing	ASML sole EUV supplier	Japan	SIA/BCG 2021
Electricity generation, 2025	~4,500 TWh	>9,500 TWh demand	~2,840 TWh; industrial prices ~2× US	—	EIA; IEA; Ember
AI patents granted, 2024	12.1% (but 52% of forward citations)	74.2%	Europe 3.0%	—	AI Index 2026
AI publications / citations	7.3% / 12.6%	17.8% / 20.6%	Europe 11.1% / 19.5%	—	AI Index 2026
Top AI researchers by place of work	59%	11%	Germany 4%	UK 6%	MacroPolo 3.0
Enterprise AI adoption	not comparable	not comparable	20% of firms (DK 42%, SE 35%, RO 5%)	—	Eurostat 2025
Population gen-AI use	28% (24th)	n/a	France 44%, Spain 42%, Ireland 45%	UAE 64%	Microsoft Jan 2026
Industrial robot installations, 2024	34,200	295,000 (54%)	67,800	Japan 44,500	IFR 2025
Robot density (per 10,000 manufacturing workers, 2024)	307	166*	EU 231; Germany 449; Sweden 377	Korea 1,220	IFR Apr 2026

Component	United States	China	EU27 / Europe	UK and others	Source, date
Public supercomputing (TOP500, June 2026)	3 exascale systems	#1 (LineShine)	JUPITER #5 (1.0 exaflop); 6 of top 15	Isambard-AI #13	TOP500
Market value (global top 30, Sept 2026)	23 companies	2	1 (ASML)	0	Market data

All values [FACT] or [ESTIMATE] as described in the sources. "Europe" in AI Index and IEA data includes the UK. IFR's April 2026 release reports China at 166 robots per 10,000 manufacturing workers using updated labour-market data from China's National Bureau of Statistics. Earlier IFR figures used a different denominator and are not directly comparable; the table therefore marks China as a series break. Version 1 also listed a 2.7% US-China "Arena" gap, which could not be re-verified and has been removed.

Private AI investment: the US has pulled away from everyone

Investment in privately held AI companies. In 2025 the US attracted 83% of the world total – 14x Europe and 23x China (nominal: \$285.9bn vs \$12.4bn). China's figure excludes state guidance funds (~\$184bn into AI firms 2000–2023).



Source: Stanford HAI AI Index 2026 (Quid data), processed by Our World in Data (updated 20 Apr 2026), deflated to 2021 US\$ with US CPI. "Europe" is the AI Index region incl. UK, not EU27. Excludes public-company capex (US hyperscalers guided ~\$650–750bn of 2026 capex), corporate R&D, and government funding.

Chart 4

Chart 4. Private AI investment by region, 2013–2025 (constant 2021 US\$). Source: AI Index 2026 / OWID.

Chart 4 is unchanged. It is in constant 2021 dollars (US 2025: \$241bn), whereas the text uses the AI Index's current-dollar figure (\$285.9bn). It covers private-company funding only. It excludes public

spending, notably Chinese state guidance funds, and hyperscalers' own capex, so it is not a comparison of total national AI investment.

Who builds the models? The US and China; Europe barely registers

Every model at the capability frontier since 2023 has been American; Chinese models trail by ~7 months on average (Epoch AI, Jan 2026). On sheer volume of large-scale models China overtook the US in 2025. The EU produced 6; the UK none.




Chart 5

Chart 5. Who builds the models? Notable models (2025) and large-scale models released per year. Sources: AI Index 2026; Epoch AI.

Chart 5 should be read cautiously. "Notable" and "large-scale" counts are different concepts, and release volume is not capability or national welfare. The annual large-scale-model counts were reproduced from Epoch's downloadable model dataset, updated 23 September 2026.

Who builds the models? The US and China; Europe barely registers

Every model at the capability frontier since 2023 has been American; Chinese models trail by ~7 months on average (Epoch AI, Jan 2026). On sheer volume of large-scale models China overtook the US in 2025. The EU produced 6; the UK none.





Chart 5

Chart 6. Three US providers hold about 70% of Europe's cloud infrastructure market, while European providers' revenues grew.

Market scope: cloud infrastructure services (IaaS, PaaS and hosted private cloud) in Synergy's European region, including the UK; €61bn in 2024.

Compute: location, ownership, access and electricity are different measures

The US hosts most AI-supercomputer performance in the best-documented sample, and five US firms own most of the world's AI compute. Neither figure says who can use capacity, and data-centre electricity use is not a measure of AI capability.



D. ACCESS is not measured by any of these. Firms and governments can rent capacity located abroad, and capacity located in Europe may be foreign-owned. Access in a crisis depends on contracts, export licences and jurisdiction.

Sources: A – Epoch AI, 'Trends in AI supercomputers' (23 Apr 2025) and data (Jun 2025). B – Epoch AI data insight (14 Apr 2026) [ESTIMATE]. C – IEA, Energy and AI (10 Apr 2025): 2024 = 415 TWh, 2030 base case = 945 TWh; regional TWh derived by author from IEA shares [FORECAST]; IEA 'Europe' incl. UK, Norway, Switzerland.

Chart 7

Chart 7. Compute: location, ownership, access and electricity are different measures.

Sources: [Epoch AI, Trends in AI supercomputers \(23 Apr 2025\)](#); [Epoch AI data insight \(Apr 2026\)](#); [IEA, Energy and AI](#). Regional IEA values derived by the author.

The AI value stack: where Europe stands today, layer by layer

'Winning AI' is not one thing. Positions and profitability differ by layer. Europe holds a chokepoint in lithography and strong supplier positions in power equipment; it is weak in chips, cloud and frontier models. Whether today's positions become durable advantages is an open question.

Layer of the stack	Who leads today	Europe's current position	Observed 2026 signal
Users & adopting industries	Every sector; gains depend on complementary investment	20% of EU firms use AI (2025); gen-AI use at US parity (EIB); fewer multi-process users	Large consumer surplus in US studies; business and national split not measured
Applications & agents	US platforms (Microsoft, Google, Salesforce)	Pockets: SAP, Lovable, Legora, DeepL	Fast growth; margins not yet established
Foundation models	US frontier labs; Chinese labs months behind	Mistral near-frontier; no frontier lab	Leading labs loss-making, investor-financed
Cloud & compute	AWS, Microsoft, Google: ~70% of European market	EU providers ~15% of own market; revenues growing	AWS operating margin ~39% (Q2 2026)
Data centres, power & grid	Global; US & China build fastest	Strong suppliers: Siemens Energy, Schneider, Legrand, ABB; Nordic power	Record order books tied to AI capex
Chip fabrication	TSMC (Taiwan)	~9-10% of all capacity; no leading edge	TSMC operating margin ~60% (Q2 2026)
Accelerators & chip design	Nvidia, AMD, Google TPU	Absent at scale	Nvidia operating margin ~66% (qtr to Jul 2026)
Lithography & equipment	ASML (NL) - sole EUV supplier; US & Japanese tools	Chokepoint held in Europe (ASML, Zeiss)	ASML gross margin ~54%; 2026 sales guided €43-45bn

Shading = Europe's current position (dark gold strong, light gold partial, grey weak/absent); author's judgement. Current margins are signals of today's scarcity, product quality and investment cycle - not proof that rents will last. Margins are not additive and do not measure each layer's share of AI value.

Sources: company filings (Nvidia Q2 FY27; TSMC Q2 2026 6-K; ASML 2025 and Q2 2026; Amazon Q2 2026); Synergy Research (Jul 2025); Epoch AI; Eurostat (2025); EIB Investment Survey 2025. CONCEPTUAL DIAGRAM - not a score. Chart 8

Chart 8. The AI value stack: where Europe stands today, layer by layer (conceptual, with 2026 data points).

What the map shows

- **The US leads the capital-intensive core:** compute, chips, cloud, private capital, frontier models and talent. The scale gap in private investment is extreme: US private AI investment in 2025 was about 14 times Europe's (including the UK) and 23 times China's, although China's state funding is excluded [ESTIMATE].
- **China leads in several physical and diffusion layers:** robots, electricity, manufacturing, patents by count, publications by volume and open-weight model releases.
- **Europe does not lead in any layer outright, but it holds assets:** the single most important chokepoint in chip-making (ASML), strong supplier positions in the power and grid equipment the AI build-out needs, a strong public supercomputing base, high robot density in several countries, relatively high population-level use of generative AI, and a research base whose citation share (19.5%) is close to China's [FACT]/[ESTIMATE].
- **The compute figures need care.** The US share of AI-supercomputer performance applies to Epoch's dataset, which covers an estimated 10-20% of global capacity. Location, ownership, access and electricity are different measures (Chart 7). A projected fall in Europe's share of data-centre electricity is not a forecast of its share of AI capability.

So what? "The US owns most frontier models" does not by itself mean higher US GDP. It means high profits today in the layers where scarcity persists; strategic leverage through

export controls, cloud jurisdiction and platform terms; pricing power where competition is weak; and an advantage in security uses where frontier capability matters. If models keep commoditising, the model advantage shrinks, while compute, chip and cloud positions may persist. If frontier AI starts to accelerate research, the model lead could matter more. Both remain possibilities.

Part IX China in the AI landscape

What the US–China competition actually consists of

Dimension	Who leads	Evidence
Frontier model capability	US, narrowly	All frontier models since 2023 American; Chinese lag ~7 months; US government tests (NIST CAISI, Sept 2025) found US models "substantially outperform" DeepSeek, especially on software and cyber tasks
Semiconductors	US and allies, decisively	Huawei's best chip $\approx$ half an H100; Huawei produces <4% of Nvidia's AI compute output; memory (HBM) is the binding constraint (Epoch, Sept 2026)
Compute	US	~75% vs ~15% of AI-supercomputer performance in Epoch's dataset
Manufacturing	China	~28% of world manufacturing value added (2023) vs EU ~17%
Robotics	China, ~9:1 in installations	295,000 vs 34,200 (2024); Chinese suppliers 57% of home market
Energy	China	Demand >9,500 TWh; ~half of global demand growth to 2030
Industrial deployment	China on volume	"AI+" plan targets >70% penetration of intelligent terminals and agents by 2027
Military applications	US ahead; China investing	See Part X
Research	China on volume; US on impact	Patents 74% China; forward citations 52% US
Capital	US, by far (private)	\$286bn vs \$12bn private; China's state guidance funds partly offset

Sources as in Part VIII; State Council AI+ Opinions (Aug 2025). [\[FACT\]](#)/[\[ESTIMATE\]](#)

Two strategic perspectives

The US view [\[OFFICIAL CLAIM\]](#). Part 0 sets out the published doctrine: the largest AI ecosystem sets standards and reaps economic and military benefits; allies should adopt the "full AI technology stack"; export controls deny adversaries the frontier; domestic regulation should be "minimally burdensome". US practice has also become transactional. Sales of downgraded Nvidia chips to China were allowed in 2025 against a 15% revenue share, and H200 sales in 2026 against a 25% share and volume caps [\[FACT\]](#).

The Chinese view [\[OFFICIAL CLAIM\]](#). The sources are the New Generation AI Development Plan (2017), the AI+ Opinions (August 2025), the 15th Five-Year Plan (March 2026) and the Global AI Governance Action Plan (July 2025). In DigiChina's reading, the five-year plan, which mentions AI 52 times, makes AI "the organizing logic for industrial transformation". The priorities are diffusion into industry and society; leadership of the "global open-source AI ecosystem"; self-reliant chips; and UN-centred governance, including a new World AI Cooperation Organization, launched in Shanghai in July 2026

with 29 founding states (no EU member identified among them). China has also shown willingness to use counter-leverage. It expanded rare-earth export controls in October 2025 (later suspended) and blocked exports of Nexperia chips packaged in China after the Dutch government intervened in the company **[FACT]**.

The two theories of victory differ **[INFERENCE]**. The US bets on frontier capability, compute denial and exporting its stack, while its Action Plan also names slow adoption as the main bottleneck. China bets on diffusion, industrial integration and open-weight reach, backed by a manufacturing and robotics base. This maps onto Ding's distinction between leading-sector innovation and general-purpose-technology diffusion. Neither government's framing should be taken as fact. US controls have been partly offset by algorithmic efficiency, smuggling (Epoch estimates about 660,000 H100-equivalents reached China through 2025) and open-weight distillation. China's AI+ targets face weak private capital; Chinese AI start-up funding fell sharply in 2025 **[ESTIMATE]**.

European relevance extends beyond a hypothetical Chinese lead. Open-weight models can increase choice and put pressure on prices; Chinese industrial deployment offers a comparison for Europe's manufacturers. At the same time, model security, data governance, component supply and the conditions of market access require scrutiny. The report cannot quantify the net value of these effects for Europe. **[INFERENCE]**

China as a model supplier. Chinese open-weight releases can provide European developers with another capable starting point. Local deployment may change the data and access questions, but it does not automatically settle model-security, licensing or update risks. Lower model prices can benefit European users even when European firms do not own the models. The size of that benefit, and any effect on supplier income, is not established here. **[INFERENCE]**

China as an industrial competitor. China combines manufacturing scale, robot deployment and growing AI capability. AI embedded in machines, vehicles and industrial software could strengthen its exporters and change the competitive pressure on European producers. That is a scenario, not a measured contribution of AI to Chinese productivity or export prices. The robotics and model counts in Part VIII describe capacity and releases; they do not prove an economic payoff. **[SCENARIO]/[INFERENCE]**

If China became the dominant global AI ecosystem what would change for Europe

- **Standards and governance** could shift towards UN channels and PRC-backed technical standards, including in surveillance-related areas **[SCENARIO]**.
- **Default models** for European SMEs and public bodies could become Chinese open-weight models: free, modifiable and good enough. US government testing found DeepSeek-based agents far more susceptible to hijacking and more likely to echo CCP narratives. Italy's data-protection authority restricted DeepSeek in January 2025 **[FACT]/[SCENARIO]**.
- **Hardware dependence** could deepen through robots and components. Europe already imports about 98% of its rare-earth magnets from China **[ESTIMATE]** (Commission).

- **Coercion** could add software dependence to material dependence, a pattern already visible in rare earths and Nexperia [INFERENCE].

And what does continued American leadership imply

- **Access policy can change quickly.** The January 2025 AI Diffusion Rule placed 17 of 27 EU member states in a restricted tier with compute caps. The Commerce Department announced its rescission on 13 May 2025, before the compliance date of 15 May 2025 [FACT] [BIS]. It is evidence that access policy can change, not proof that the caps constrained European deployments. Access is now governed by bilateral deals and licences; the UAE received licence-free access in July 2026 [FACT]. That may be less predictable for Europe [INFERENCE].
- **Legal reach.** US law can compel US providers to hand over data wherever it is stored. Microsoft France told the French Senate in June 2025 it could not guarantee otherwise [FACT].
- **Political exposure.** US AI firms are themselves subject to executive leverage, as the Anthropic designation and the 14 September post illustrate. European users of US models are therefore exposed to US domestic politics [INFERENCE].
- **Pressure on EU rule-making.** The US imposed visa bans on a former Commissioner and others over digital regulation (December 2025) and has threatened trade measures against "digital taxes, legislation or regulations" [FACT].

What this comparison shows. The US and China offer different mixes of capability, price, industrial reach and governance. Competition between them may expand European choice and reduce costs, while concentrated dependencies in either ecosystem can create exposure. The US is an ally and major technology partner; the EU describes China as simultaneously a partner, competitor and systemic rival (EEAS factsheet, 2023). These relationships are not interchangeable. For any specific European use, the relevant questions are what value the technology enables, who controls access and data, and how easily a user could switch.

Part X From economic power to geopolitical power

The bridge many readers miss

Productivity can become power through two broad channels. Neither is automatic.

The resource channel. Higher productivity enlarges the tax base and savings. Those are *potential* resources for defence, R&D, infrastructure, aid, diplomacy and external investment. Paul Kennedy's *Rise and Fall of the Great Powers* (1987) made relative productivity the long-run base of military power. Michael Beckley (2018) argues that raw GDP overstates the power of large, poor countries, and proposes "net power", measured roughly as GDP multiplied by GDP per capita. That makes *productivity*, not just size, central **[INFERENCE]** (from scholarly literature). The scale of the channel is visible in military budgets: the US spent about \$997bn in 2024, China an estimated \$314bn and states in SIPRI's Europe region (which includes Russia and Ukraine) \$693bn **[ESTIMATE]** (SIPRI). NATO allies committed in June 2025 to spend 5% of GDP on defence and security by 2035, of which 3.5% on core defence. A growing economy makes that easier to fund.

The dependence channel. Control of technologies others need creates bargaining power (Farrell and Newman's "weaponised interdependence"). There is a *chokepoint* effect, denying access, and a *panopticon* effect, gaining information through hubs one controls. Documented examples include semiconductor export controls, cloud jurisdiction, payment networks and satellite connectivity.

Between resources and power sits conversion. Resources become capability only through political decisions, institutions, procurement, industrial capacity, alliances and coordination. A richer economy that cannot procure, build or agree does not become more powerful. This matters especially for Europe, because **the EU is not a single state**. Defence budgets are set in 27 capitals and coordinated through NATO. EU-level money is small relative to national budgets. Procurement is fragmented, and industrial capacity is spread across many firms and countries. The same euro of GDP may therefore yield less capability in Europe than in a single state with one budget and one procurement system — or more, if Europe coordinates well.

From productivity to power: resources have to be converted

Higher productivity enlarges potential resources; it does not by itself create military strength or influence. Political decisions, institutions, procurement, industrial capacity and alliances decide what resources achieve. The EU is not a single state: much of this conversion happens in 27 capitals and in NATO.

1. RESOURCE CHANNEL



2. DEPENDENCE CHANNEL



CONCEPTUAL DIAGRAM – arrows show mechanisms, not inevitable results. Draws on Kennedy (1987), Beckley (2018), Ding (2024), Farrell & Newman (2019).

Chart 14

Chart 14. From productivity to power: resources have to be converted by institutions (conceptual).

Version 1's Chart 14 showed a closed loop from productivity to strategic leverage without institutions. The revised chart adds political and institutional conversion; arrows show mechanisms, not inevitable results.

AI-specific channels what is documented what is plausible what is speculative

The status column now describes the evidence narrowly. A documented use does not show that the use was decisive.

Channel	Status	Evidence
Intelligence analysis and targeting	Documented use	Project Maven used in 2024 strikes; NATO acquired the Maven Smart System (March 2025); more than 20,000 active users across more than 35 service and combatant-command tools (NGA, May 2025)
Drones and autonomy	Documented use, partial	Ukraine: AI terminal guidance reportedly raised strike success from ~10–20% to ~70–80%; engagement decisions still human (CSIS 2025)
Cyber operations	Reported by a vendor	Anthropic reported a campaign it attributes to a Chinese state-sponsored group that used an AI coding agent for 80–90% of an espionage operation (Nov 2025). Not independently confirmed

Channel	Status	Evidence
Defence adoption of frontier AI	Contracts and strategies in place	US contracts of up to \$200m each with four frontier labs (2025); "AI-first warfighting force" strategy (Jan 2026)
Information operations	Plausible	Lower cost of tailored propaganda; model bias findings
Government administration	Plausible	Large efficiency potential; early deployments
Scientific discovery	Plausible, uncertain	AlphaFold-type breakthroughs; broad acceleration not yet shown; the US Genesis Mission is a bet on it
Ocean and space transparency	Plausible	PLA procurement targets AI detection of US naval assets (CSET, Feb 2026)
"Decisive strategic advantage" from very advanced AI	Speculative	No evidence; the premise behind the strongest reading of "whoever wins, wins"

Sources: CSIS (Bondar, Mar 2025); [NATO NCIA \(Mar 2025\)](#); [Anthropic \(Nov 2025\)](#); CDAO (Jul 2025); CSET (Feb 2026).

Why superior AI capability could make an already powerful state more powerful

Four mechanisms can reinforce each other **[INFERENCE]**:

1. **More output per euro of defence spending:** better intelligence, faster targeting cycles, cheaper autonomous systems.
2. **A larger resource base:** if AI raises productivity, the tax base grows.
3. **Leverage over others:** control of chips, cloud and models lets a state set terms for allies and deny adversaries.
4. **Faster innovation:** AI in R&D may shorten the cycle for the next generation of weapons and technologies.

The best-evidenced effects today are **incremental and diffuse**, not a single decisive weapon. They favour states with mass-production and integration capacity (drones, robots, electricity) as much as those with frontier models. Mechanism 2 depends on diffusion, which is slow. And all four depend on institutions that can turn capability into decisions.

So what? For Europe, the geopolitical stakes of AI run mainly through the resource channel (can European economies grow fast enough to fund their own security?) and the dependence channel (can Europe act — militarily, legally, diplomatically — without permission from a supplier state?). In both, the constraint is often European decision-making and coordination rather than money alone. Frontier-model ownership matters mostly in the dependence channel and for security uses. It would matter more if frontier AI began to accelerate research.

Part XI Standards dependence and freedom of action

Normal interdependence versus strategic dependence

Europe imports most of its digital technology. The Commission states that the EU "spends EUR 264 billion a year mostly on US proprietary IT products and services". It also says the EU relies on non-EU countries for over 80% of key digital products, services, infrastructure and IP **[OFFICIAL CLAIM]** (the 80% figure originates with the EuroStack report and its method is not transparent). Import dependence as such is not a problem: specialisation is how economies get rich. The question is when dependence becomes *strategic*.

A practical test **[INFERENCE]**: a dependence is strategic when all four conditions hold.

1. **Leverage:** the supplier or its home state can legally or technically deny, degrade or surveil.
2. **Weak substitutes:** no alternative can be switched in within the time a crisis allows.
3. **Critical function:** the function affected is critical (state continuity, defence, critical infrastructure, core industrial processes).
4. **Plausible divergence:** the supplier state's interests could diverge from Europe's.

A bakery using American office software fails test 3. A defence ministry on US-controlled cloud, or an EU car line depending on Chinese magnets, may pass all four. Since 2025, test 4 has moved for the United States as well as for China. That shift is at the core of the new European debate.

What has actually happened

Case	What happened	Evidence
Legal access to data	Microsoft France, under oath to the French Senate (June 2025): "No, I cannot guarantee that" French data would not be handed to US authorities	[FACT]
Surveillance law	FISA Section 702 lapsed by statute on 12 June 2026, but collection continues under court certifications until about March 2027; the EU-US Data Privacy Framework rests on this architecture and faces an appeal at the Court of Justice	[FACT]
Service disconnection	After US sanctions, the ICC Prosecutor's Microsoft email was disconnected (May 2025); who initiated the disconnection is disputed. The ICC moved to Germany's openDesk in October 2025. Sanctioned ICC judges report losing US card, e-commerce and cloud services	[FACT] (initiator CONTESTED)

Case	What happened	Evidence
Pressure on rule-making	US visa bans on a former EU Commissioner and four others over digital regulation (December 2025)	[FACT]
Tiered chip access	AI Diffusion Rule (Jan 2025) put 17 EU states in a restricted tier; rescission announced 13 May 2025, before the 15 May compliance date	[FACT]
Compute as trade concession	US–EU framework (August 2025): EU "intends to purchase at least \$40 billion worth of U.S. AI chips"	[OFFICIAL CLAIM]
Chinese counter-leverage	Rare-earth controls (April and October 2025; the October wave suspended until 10 November 2026); Nexperia chip exports blocked, bringing EU carmakers "days away" from stoppages	[FACT]
Vendor pricing power	Broadcom/VMware: European cloud providers allege cost increases above 1,000%	[CONTESTED]
Defence integration	NATO's targeting and data-fusion system is Palantir's Maven; Germany's Bundeswehr held back, citing access to national databases	[FACT]
Satellite dependence	"Their entire front line would collapse if I turned it off" (Musk on Starlink in Ukraine, March 2025)	[FACT] (statement made; no switch-off occurred)

Sources: The Register (Jul 2025; Feb 2026); [Brennan Center \(Sept 2026\)](#); [JusticeInfo](#); Euronews; Federal Register; White House joint statement (21 Aug 2025); EPRS (Nov 2025); CISPE; NCIA.

Pattern [INFERENCE]: coercion so far has come through infrastructure, sanctions compliance and export licensing, not through AI models. A recorded event does not prove every claim attached to it. In the ICC case, the disconnection is documented; who initiated it is disputed. The ICC's later migration shows that a switch is possible; it does not by itself show that the court could operate independently of US services for a defined period.

Three different risks and what reduces each

Risk	Example	What reduces it	What does not
Access: a foreign state compels disclosure	CLOUD Act orders	European legal and operational control of the service, plus technical measures such as customer-held keys where the provider cannot see plaintext. Protection depends on the architecture, update dependencies and jurisdiction, and should be specified and tested	Data residency or EU staff alone; the AWS European Sovereign Cloud (live since January 2026) is judged by several analysts still reachable through the US parent
Continuity: a foreign state or vendor stops service or updates	ICC email; Starlink statement	A tested ability to run and update without the vendor for a defined period	Contractual promises to litigate (Microsoft's April 2025 pledge)
Economic: lock-in and pricing power	VMware	Portability, open standards and competition law. The Data Act has applied since 12 September 2025; cloud switching charges must be fully removed from 12 January 2027. Possible DMA cloud designation	Sovereignty rhetoric

Sources: [Commission, Data Act explained](#); [EDPB Recommendations 01/2020 on supplementary measures](#).

No technical or legal arrangement makes a system "immune". Version 1 used that word and "removes access risk"; both overstated what such measures achieve. France's SecNumCloud model — a European-controlled operator running licensed US technology, such as S3NS (Thales–Google), qualified in December 2025 — reduces much of the access risk and part of the continuity risk. It does so at the cost of continued technology dependence, because updates still come from the licensor **[FACT]/[INFERENCE]**.

The cost of sovereignty

Sovereignty is not free, and European ownership is not automatically better.

- **Full-stack replication** (EuroStack) is costed at €300bn over ten years **[ESTIMATE]** (advocacy). That is about €30bn a year. For scale, it is about 5% of the 2026 capex guidance of four US hyperscalers. The comparison is not like-for-like: hyperscaler capex is global, private and mostly general cloud capacity.
- **Targeted sovereignty** is much cheaper. An illustrative calculation: suppose 10–20% of the EU's €264bn annual IT spend is workloads that need high assurance, and the sovereignty premium (cost plus capability shortfall) is 10–30%. The annual cost is then about €3–16bn, or 0.015–0.1% of EU GDP **[ILLUSTRATION]**.
- **Opportunity costs** are real. Sovereign offerings often lag on features (the AWS sovereign region launched with about 90 services against 200+), and small providers may be *less* secure against non-state threats than hyperscalers **[ESTIMATE]**.

- **Gaia-X** shows how sovereignty projects fail: standards without anchor customers, governance by incumbents, federation without capacity, and vague definitions **[INFERENCE]**.

The Commission's June 2026 package (the Tech Sovereignty Communication and the proposed Cloud and AI Development Act) is the first official attempt to price this trade-off. It defines sovereignty as the ability to "develop, control and scale" critical technologies. It states that sovereignty "does not mean isolation, protectionism, or tech decoupling". It proposes four assurance levels for public-sector cloud, to be chosen through a sovereignty risk assessment for each use case **[OFFICIAL CLAIM]** ([COM\(2026\) 503](#)).

What the evidence shows. Sensitive state, defence and infrastructure uses expose a different kind of dependency from ordinary commercial workloads. The relevant test is whether a service can meet a specified level of protection and keep operating under plausible disruption. For other uses, global technology may provide greater capability at lower cost. The appropriate boundary depends on the use case, threat and available alternatives.

Part XII Does Europe need to lead the frontier to benefit

Five pathways and the conditions behind them

The earlier research version rated strategic options on a common scale. No common cost model supported those ratings. This edition presents five pathways and the conditions under which each might produce value. They can overlap; none is selected as a recommendation.

Europe's strategic options: conditions, not scores

The options are not mutually exclusive and are not ranked here. Their value depends on conditions that can be observed and debated. Combining them is a policy judgement with trade-offs, not a demonstrated optimum.

	What it could deliver	Conditions for it to work	Costs and risks	Indicators to watch
A Frontier or near-frontier capability	Controllable models for sensitive uses; research spillovers; bargaining power	Sustained capital and compute; talent; anchor customers; value of a lead stays high	Tens of \$bn per lab cycle; may trail leaders anyway	Gap to frontier (months); European training compute
B Broad, deep adoption	Most of the productivity gain for users, workers, customers	Skills, management, data, redesigned work; competitive supply; firm investment	Substantial firm-level investment; dependence on foreign suppliers	Multi-process use; productivity in adopting sectors
C Industrial and domain AI	Durable value in machines, processes, domain software	Firms keep the differentiating software layer; data sharing; compute access	Platforms move into verticals; robotics ownership already largely non-European	Software revenue of industrial firms; AI-enabled exports
D Assured core for critical functions	Continuity and legal control for state, defence and critical infrastructure	Clear scope; tested fallback; European-operated deployment; procurement	Premium cost; feature lag; risk of scope creep	Share of critical workloads with a tested exit plan
E Hosting foreign AI capacity	Construction, power sales, local jobs and tax; physical resilience	Grid capacity; terms on heat, access and security; net public cost	Uses scarce grid capacity; little change in control	Net public cost per MW; local access agreements

DISCUSSION TOOL – author's synthesis of Parts XI–XIII. No common cost model exists for these options; figures in the text are orders of magnitude only.

Chart 15

Chart 15. Five European pathways: conditions, not scores (discussion tool).

Option A — frontier or near-frontier capability. Mistral raised €3bn at a valuation above €21bn in September 2026 in a round led by Samsung, after ASML invested €1.3bn in 2025 [FACT]. It is Europe's only near-frontier lab. It is becoming a vertically integrated "sovereign AI" provider: models plus compute plus enterprise and defence deployment. Germany's Aleph Alpha and Canada's Cohere signed a definitive business-combination agreement in September 2026. The combined company would operate as Cohere and remain rooted in Canada and Germany; completion remains subject to regulatory approval [FACT]. The capital gap to the US frontier is one to two orders of magnitude: Mistral's record round is about a tenth of Anthropic's February 2026 round [ESTIMATE].

What would make more ambition worthwhile: evidence that frontier AI accelerates research, so that a lead compounds; security uses that require controllable frontier models; or the loss of open-weight supply. *What counts against:* cost, and the risk of trailing the leaders anyway. Comparing four US firms' annual global capex (roughly \$650–

750bn) with the EU's €20bn gigafactory fund (about thirty times smaller) shows the scale gap in infrastructure. It is not an affordability test for frontier model leadership. Model development, training, inference, general cloud capacity and public funding are distinct budgets, and some are private.

Option B — broad, deep adoption. The economics favour it as a source of productivity. The OECD estimates 0.4–0.9 points a year of added US labour-productivity growth over a decade (Filippucci et al., 2024) [MODEL OUTPUT]. Its G7 work gives 0.5–0.7 points for Germany and France in a medium-adoption scenario (Schief et al., 2025) [SCENARIO]. Ding's history suggests diffusion shaped past shifts in economic power. EU adoption is level with the US on breadth (37% vs 36% of firms using generative AI) but behind on depth (55% vs 81% of AI users apply it in two or more processes) [ESTIMATE] (EIB). *Conditions:* skills, management, data, redesigned work and competitive supply. Adoption can require substantial investment by firms. Public costs depend on the support, infrastructure and procurement measures chosen. *Risk:* on its own it deepens dependence on foreign suppliers.

Option C — industrial and domain AI. Tested in Part XIII. *Condition:* European firms keep the differentiating software layer in their own products.

Option D — an assured core for critical functions. This targets exactly the dependencies that pass the strategic test. The EU AI Gigafactories call (opened 30 July 2026; up to seven sites; bids due 12 November 2026; selection early 2027) is one vehicle. Only about €1bn of EU money is currently budgeted; the rest depends on the 2028–34 budget [FACT]. *Conditions:* a clear and narrow scope, tested fallbacks and procurement that creates demand. *Risk:* premium costs and scope creep.

Option E — hosting foreign AI capacity. The UK (a £31bn package of US commitments in September 2025, "Stargate UK"), Norway ("Stargate Norway", 100,000 GPUs) and Ireland (data centres used 22% of metered electricity in 2024) show the model [FACT]. It brings construction, power demand, local jobs and some tax, but no change in legal control. *Conditions:* grid capacity and terms on heat reuse, local access and security. *Risk:* it uses scarce grid capacity that could serve other users.

The de facto EU path combines them. Apply AI (adoption), the Cloud and AI Development Act (capacity and sovereignty levels), AI Factories and gigafactories, support for Mistral and other firms, and the \$40bn US-chip purchase intention together form a hybrid. Its risk is being a bit of everything and enough of nothing. Part XV assesses delivery.

What the comparison can and cannot settle

Three observations follow from the evidence.

- **Europe can capture substantial AI benefits without leading every frontier benchmark** [INFERENCE]. Most of the productivity gain available to Europe depends on effective use, and access to frontier models is currently available from several suppliers.
- **Whether additional frontier investment is worthwhile depends on cost, spillovers, security needs and credible alternatives** [INFERENCE]. The case strengthens if frontier AI begins to accelerate research, if security uses require controllable models, or if access becomes less reliable.

- Near-frontier European capability could have insurance and bargaining value, especially if frontier AI accelerates research or access narrows. Its cost and benefits remain uncertain, so the evidence does not establish an optimal level of investment [JUDGEMENT].

Version 1 stated a 55% probability that an EU lab would be within twelve months of the frontier in 2030. That figure had no documented method and has been removed. The indicator to watch is the measured gap, in months, between the best European model and the frontier.

The implication. Both “Europe needs its own ChatGPT” and “Europe can forget the frontier” collapse several questions into one. Productivity, domestic business formation, critical-function resilience and frontier research each have different conditions and measures of success.

Part XIII Industrial AI An alternative path for Europe

The claim under test: "Europe's opportunity may be less about replicating Silicon Valley and more about combining AI with advanced industrial capability."

The base is real and under pressure

- The EU produced about 17% of world manufacturing value added in 2023, against China's 28%. UNIDO projects China at 45% by 2030 [FACT]/[FORECAST].
- Machinery, transport equipment and chemicals are about 61% of extra-EU goods exports [FACT].
- Robot density is high in some countries: Germany has 449 and Sweden 377 robots per 10,000 manufacturing workers (third and fifth in the world); the EU average is 231 [FACT] (IFR 2026). But China installs more than four times as many robots a year as the EU.
- The EIB finds EU manufacturers ahead of US manufacturers in adopting AI and big data (48% vs 28%) [ESTIMATE] (EIB).

What is happening and who captures the value

Picks and shovels: Europe's clearest AI-era gains so far [INFERENCE]. ASML raised its 2026 sales guidance to €43–45bn (from €34–39bn; 2025 sales were €32.7bn), citing AI. Siemens Energy has a record €162bn backlog and about 69 GW of gas turbines on order — much of it, according to the company, for data centres and the US. Data centres and networks account for about 30% of Schneider Electric's order exposure; data centres are 26% of Legrand's sales. This is industrial capability *serving AI* [FACT] (for the figures above). It is large and real, and it is cyclical, because it rides on US capex.

Industrial software incumbents. Siemens spent about \$15bn on Altair and Dotmatics in 2025. It is building industrial copilots and an "Industrial AI Operating System", with Microsoft and Nvidia supplying the model and platform layers. SAP's AI features run largely on others' models and data platforms; its flagship "OpenAI for Germany" is an OpenAI model on Microsoft technology, operated by an SAP subsidiary. SAP's acquisition of Freiburg-based Prior Labs, with more than €1bn committed to build tabular foundation models for structured business data, is the most interesting counter-move [FACT].

Robotics: the opposite of the thesis. KUKA has been owned by China's Midea since 2017. ABB Robotics is being sold to Japan's SoftBank for \$5.4bn, explicitly for SoftBank's "physical AI" strategy [FACT]. Europe has no large European-owned general-purpose industrial-robot champion at the moment robot "brains" are becoming a horizontal layer (Nvidia GR00T, Figure's Helix, Chinese humanoids).

Automotive: the cautionary case. Volkswagen committed \$5.8bn to a joint venture with Rivian and pays Xpeng for software architecture; Xpeng booked RMB 1.72bn in licensing fees from VW [FACT]. Mechanical engineering excellence did not translate into software leadership, and value moved to whoever owned the new layer.

Defence: the strongest case so far [INFERENCE]. Helsing was valued at \$18bn in July 2026 and Quantum Systems at \$8bn. Helsing's AI agent flew a Saab Gripen E in 2025, and Mistral has a framework agreement with the French armed forces [FACT]. Demand is sovereign, battlefield data from Ukraine is available to European-aligned firms, and the firms combine AI-native software with hardware.

Diffusion is the binding constraint. In Europe's flagship branches (machinery, electrical equipment, vehicles, electronics) about one firm in four uses AI, against about two in three in IT services [FACT] (Eurostat 2025). Energy costs are cited as a barrier by 75% of EU firms against 36% of US firms [ESTIMATE] (EIB).

Does industrial AI shift rents to European incumbents or to US platforms

For incumbents	For platforms
Industrial data is proprietary and not on the web	Every flagship European industrial-AI product examined has a US model, cloud or chip inside it
High switching costs in PLM, ERP and automation	Platforms are moving into verticals (Nvidia blueprints and Omniverse; Palantir Foundry under Airbus Skywise)
Liability and certification favour domain integrators	Automotive precedent: value moved to the software layer
Horizontal industrial "operating systems" have failed before (GE Predix); customers resist handing data to a platform	Robotics foundation models may commoditise machine know-how

The "smile curve" frames the question. Value can concentrate at the two ends — chips, models and platforms on one side; customer relationships, installed base and data on the other — and be squeezed in the middle (integration). European incumbents hold the right-hand end. A downstream position is not inherently low-value: installed equipment, certification, domain data and customer relationships can earn durable returns. The question is whether they stay sticky as platforms move in.

Verdict

Partly supported, and conditional [JUDGEMENT]. Europe already earns large AI-era revenues as a supplier (lithography, power, grid), and its incumbents hold data, installed equipment and customers. But industrial AI is not a separate path that avoids the frontier stack: it runs on US models, clouds and chips, and robotics has largely passed to Asian ownership. The defensible framing is "industrial AI plus *contestable access to models and compute*", not "industrial AI instead of Silicon Valley".

The conditions for success are: deep adoption in manufacturing SMEs; European data-sharing that works at scale (Catena-X, with more than 1,000 members, shows it can); competitive access to compute; a European model layer credible enough to keep prices honest (Mistral's customers — Schneider, Thales, Veolia, SNCF, Airbus — suggest demand exists); and firms that keep control of the differentiating layer as they grow.

So what? The industrial path is one of Europe's most plausible routes to *capturing* AI value rather than only using it, but only if European firms keep the layer that becomes differentiating in their own products. The VW and ABB Robotics cases show how quickly that layer can move abroad.

Part XIV Sweden as a case study

Sweden makes the wider European question concrete: how do existing assets become productivity, resident income and dependable access?


Assets and constraints

Asset	Evidence	Class
Electricity surplus	Generation 168 TWh (2025); net exports 33 TWh, about eight times what Swedish data centres use (~4 TWh)	[FACT]
Cheap northern power	2025 average day-ahead price $\approx$ €16/MWh in SE1/SE2, $\approx$ €46 in SE3, $\approx$ €61 in SE4, against €89 in Germany	[FACT] (€ conversion approximate)
R&D intensity	3.57% of GDP (2024), the highest in the EU	[FACT]
ICT specialists	8.9% of employment, the highest in the EU	[FACT]
Enterprise AI adoption	35.0% of firms with 10+ employees (2025), third in the EU, up ten points in a year	[FACT]
Robot density	377 per 10,000 manufacturing workers, fifth in the world	[FACT]
Domain champions	Ericsson, Saab, Volvo, Scania, Atlas Copco, Sandvik, SKF, AstraZeneca, Hexagon	[FACT]
Start-up formation	Lovable and Legora among Europe's fastest-growing AI firms	[FACT]
Private research capital	Wallenberg AI programme (WASP), SEK 7.2bn in total	[FACT]
Funded pension capital	AP buffer funds SEK 2.2trn; premium pension SEK 2.9trn (end-2025)	[FACT]
Fiscal space	Government debt was 35.1% of GDP in 2025	[FACT] (Statistics Sweden; preliminary)
Alliances	EU member; NATO member since March 2024	[FACT]

Sweden also faces **real constraints**. North–south transmission is limited, and grid connections take years. Scale-up capital is thin relative to the firms' ambitions. Productivity growth has slowed. And the state's AI budget is small relative to its stated goal. Version 1 concluded that "Sweden does not have an asset problem". That conflicted with these constraints and has been withdrawn.

Wholesale power is cheapest in Sweden's north

Sweden was a net exporter of 33 TWh in 2025. North-south transmission limits how far northern prices reach; many new data-centre connection applications are in SE3. Wholesale averages are not all-in industrial costs and do not show whether grid connections are available.




Sources: Energiföretagen, Elåret 2025 (öre/kWh: SE1 18, SE2 18, SE3 51, SE4 67; converted at ~SEK 11/€, approximate); Bundesnetzagentur/SMARD (Germany €89.32/MWh). Svenska kraftnät: over 9,000 MW of new withdrawal applications in 2025, about half data centres – applications, not committed demand. Chart S1

Chart S1. Wholesale power is cheapest in Sweden's north. Wholesale averages are not all-in industrial costs.

Sources: [Energiföretagen, Elåret 2025](#); Bundesnetzagentur/SMARD; [Svenska kraftnät](#). Grid applications are not committed demand or completed connections.

Enterprise AI use in 2025: Sweden third in the EU

This measures breadth: whether a firm uses at least one AI technology. It does not show how deeply AI is used. The EIB's finding that EU AI users apply AI in fewer processes than US users is an EU-wide result; Swedish depth of use is not measured here.



Source: Eurostat, isoc_eb_ai, 'Use of artificial intelligence in enterprises' (Dec 2025); excludes agriculture, mining and finance. EIB Investment Survey 2025 (EU vs US depth). Chart S2

Chart S2. Enterprise AI use in 2025: Sweden third in the EU. This measures breadth, not depth.

Source: [Eurostat, isoc_eb_ai](#). Version 1 titled this chart "breadth is not depth". That judgement rested on EU-wide EIB data and is not established for Sweden.

The conversion question


Assets are not outcomes. The chain from physical capacity to prosperity has many links, and each can fail. **Whether Sweden converts its assets into prosperity and freedom of action is a hypothesis to test, not a finding.** Chart 16 sets out complementary inputs and the outcomes that would show whether conversion is happening.

Sweden's conversion question: what to measure

Sweden has strong assets and real constraints (grid, scale-up capital). Whether assets turn into prosperity and freedom of action is a hypothesis to test with observable outcomes, not a finding.

COMPLEMENTARY INPUTS

(needed together, not in sequence)



CONCEPTUAL DIAGRAM. Inputs drawn from Svenska kraftnät (2026), Energiföretagen (2026), Eurostat (2025), SOU 2025:12, SOU 2025:96.

Chart 16

Chart 16. Sweden's conversion question: complementary inputs and the outcomes to measure (conceptual).

Version 1 drew a serial chain in which each step followed the last and domestic ownership appeared as a precondition for gains. Skills, capital, data and compute are complementary, and foreign-owned operations can also generate wages and tax.

Asset	What the evidence shows	What it does not show [INFERENCE]
Cheap, clean electricity	The largest compute builds are foreign-owned: Microsoft (SEK 33.7bn; its Sandviken build was paused in April 2025 for lack of grid capacity) and Brookfield (up to SEK 95bn, 750 MW, conditional). OpenAI's Nordic "Stargate" went to Norway. The Swedish-controlled platform (Sferical AI, Wallenberg-anchored) is far smaller	Hosting brings construction, jobs, power sales and taxable activity; operational control and most profits sit with foreign owners. Net value depends on terms and on alternative uses of the grid capacity

Asset	What the evidence shows	What it does not show [INFERENCE]
Grid	Svenska kraftnät received over 9,000 MW of withdrawal applications in 2025, about half from data centres. An industry report puts data-centre applications in SE3 at 6.7 GW in 2025, up from 1.3 GW a year earlier (the two sources define applications differently). North–south transmission limits how far cheap northern power reaches	Applications are not committed demand
R&D intensity	Productivity growth fell from 2.7% a year (1996–2006) to 0.8% (2011–19) and 0.4% (2020–23) (Productivity Commission, SOU 2025:96)	The causes of the slowdown are debated; AI is one possible remedy among several
AI adoption	High breadth (35% of firms)	Depth of use in Sweden is not measured here; no productivity pick-up is yet visible in the data
Start-ups and scale-ups	Several distinct events: <i>foreign incorporation or listing</i> (Klarna listed in New York as a UK plc; Spotify is domiciled in Luxembourg, with main operations in Stockholm); <i>foreign acquisition</i> (Sana sold to Workday, \$1.1bn); <i>foreign financing</i> (recent rounds for Lovable and Legora led mostly by US investors; Lovable's latest co-led by the EU-backed Scaleup Europe Fund)	These events have different consequences. Sale proceeds, retained employment and later expansion in Sweden also count. A list of deals does not measure national value leakage
Domain champions	Saab's AI pilot agent came from Germany's Helsing; Ericsson's sales have fallen since 2022	Combining Swedish domain capability with imported AI can still create Swedish value
Pension capital	Heavily invested in global and US equities	Gives savers returns from foreign AI firms, but no operational control
State AI strategy	Goal: "among the ten leading AI nations". Budget: SEK 479m in 2026, about SEK 500m a year planned for 2027–30, roughly a fifth of the AI Commission's recommended SEK 11.5bn over five years	Ambition–money gap, unless the goal is revised
A domestic value-chain bet	Northvolt (\$7.7bn raised; \$5.8bn of debt) went bankrupt in 2024–25; its assets were sold to US-based Lyten	Execution risk in capital-intensive industrial bets

Sources: Energiföretagen (2026); Svenska kraftnät (Feb 2026; LMA 2026); Eurostat; SOU 2025:12; Productivity Commission, SOU 2025:96; written answer 2025/26:219; company disclosures; Dealroom. Productivity figures verified against the commission's published period averages.

How to test conversion. Measure observable outcomes rather than counting deals: domestic value added and productivity in AI-using sectors; income of Swedish residents; R&D employment in Sweden; exports of AI-enabled goods and services; taxable activity; adoption outcomes, including depth of use; and the number of critical functions that can run without foreign permission.

Why a persistent gap would matter for Swedes

Wages. Real wages tend to follow productivity over time. Sweden's productivity slowdown since 2008 is one reason why real wage growth has been weaker than in 1996–2006 [INFERENCE].

Income pensions — a precise mechanism for one part of the system. Income pensions (*inkomstpension* and *tilläggspension*) in payment are recalculated each year by the change in the income index, reduced by the 1.6% norm (the change is divided by 1.016). When the system's balancing mechanism is active, a balance index replaces the income index [FACT] ([Pensionsmyndigheten, glossary](#)). In plain terms: **income pensions in payment rise in real terms only when average incomes grow faster than inflation plus 1.6% a year.** The balance ratio stood at 1.147 for 2027, so balancing was not active [FACT].

This applies to income pensions, not to pensioners' total income. Premium pensions, occupational pensions, the guarantee pension, taxes and transfers follow other rules. Version 1 said faster productivity was "the only painless way to raise pensions". That was too broad. The safer statement is that sustained productivity growth strengthens the economic foundations of wages, pensions and public services. If AI's gains show up mainly abroad, Swedish savers may still share in them through the AP funds' returns, but income pensions would not.

Government finances and welfare. An ageing population raises demand for health and elderly care, where productivity is hard to raise and hard to measure. AI in public administration (the new AI-verkstad, SEK 200m a year) is one possible way to contain cost growth. In the Part IV illustration, a 7.8% lower output level by 2045 means gross public revenue lower by about 3.2% of the lower-output GDP at Sweden's 2025 tax ratio of 40.5%, assuming the same effective tax yield [ILLUSTRATION]. That is of the same order as planned annual defence spending, but it is not budget headroom: public-sector wages, pensions and service demand also move with the economy.

Defence. The government's stated path to 3.5% of GDP by 2030 is the largest new public spending flow. Saab–Helsing shows AI integration is already happening in Swedish platforms. How much value Sweden keeps depends on control of data, integration and at least some of the compute.

Industrial competitiveness. Sweden's export champions compete with firms in the US and China that are adopting AI in engineering and production. The automotive precedent (Part XIII) shows how value can move to whoever owns the software layer.

Strategic autonomy and influence. A small state's international weight depends on being indispensable somewhere. Sweden is important in a few places: telecoms, defence platforms and clean industrial power. Converting those into AI-era positions may matter more than a national frontier model.

Electricity politics. Several gigawatts of data-centre applications collide with price-sensitive households and industry in SE3. Without a siting framework, Sweden risks higher southern prices and little domestic value added [INFERENCE]. The previous government declined to commit to a national data-centre siting plan (interpellation 2025/26:414) [FACT].

What the Swedish case can reveal

The 13 September 2026 election left future budget choices open when this report was completed. Six questions are useful for assessing how Sweden converts its assets into outcomes [INFERENCE]:

1. Adoption depth. Are public bodies and industrial SMEs using AI in several important processes, and is that use visible in output or service quality?
2. Grid and siting. What value, employment, taxable activity and local access would data-centre projects create relative to other uses of scarce grid capacity? The legal rules for connection priorities require separate assessment.
3. Critical operation. Which defence, security and sensitive public-data functions have tested fallbacks if a supplier's services or updates become unavailable?
4. Scale-up. Are Swedish AI firms able to finance growth and retain research, employment and income in Sweden when they choose to do so?
5. Defence integration. Does procurement translate Swedish expertise into domestic capability and a lasting position in the value chain?
6. Public ambition and resources. How do stated AI goals compare with committed funding, delivery dates and measured outcomes?

Sweden has strong assets and real constraints. The open question is how much AI raises Swedish productivity and resident income, and whether critical functions remain dependable. The indicators above would make that conversion visible.

Part XV European AI policy as an object of study

The different outcomes in the debate

European AI debates often bring together productivity, innovation, value capture, resilience, security, competition, consumer welfare, distribution and sovereignty. These outcomes are connected, but they are not interchangeable. Table 21 separates ends from possible means to clarify what is being measured.

Objective	End or means?	Why
Prosperity (productivity, living standards)	End	What citizens ultimately care about
Security (ability to defend and resist coercion)	End	Precondition for everything else
Democratic self-government (ability to make and enforce one's own rules)	End	The EU's constitutional purpose
Fair distribution / social stability	End (constraint)	Legitimacy of the whole project
Innovation	Means	Valuable because it raises prosperity and security
Competitive European businesses	Means	Income, tax base, agglomeration and bargaining power
Domestic value capture	Means	Matters for income, control and agglomeration
Competition	Means	Keeps rents with users
Consumer welfare	Part of prosperity	—
Resilience	Means	Protects the ends against shocks
Technological sovereignty	Means	Worth paying for to the extent it protects the ends against plausible threats

The European Parliament's January 2026 resolution and the Commission's June 2026 communication put different weight on technological sovereignty [OFFICIAL CLAIM]. One treats dependence as a central political danger; the other describes sovereignty as selective capacity to develop, control and scale critical technologies without decoupling. The economic trade-off is whether added control is worth its cost in a particular use.

What could winning mean for Europeans

A European outcome cannot be read from one model benchmark. Four observable dimensions help organise the question:

1. **Prosperity:** faster productivity growth and rising living standards, visible in national accounts.

2. **Competitive businesses:** European firms that build, sell and integrate AI profitably, including in industry and defence, and that can scale in Europe.
3. **Distribution:** gains broad enough to sustain political consent, with tax systems that adapt if income shifts from labour to capital.
4. **Strategic options:** the ability to keep critical functions running and to make and enforce European rules without permission from a supplier state.

These dimensions need not move together. Europe could gain productivity while foreign suppliers earn substantial income, or own a strategic capability that does little for broad living standards. Several countries can benefit at once even as their relative power changes.

Three uncertainties remain central to the debate [INFERENCE]:

1. Continuity. Which functions must keep operating under supplier disruption, for how long, and at what additional cost? The answer may differ between ordinary business software and defence or core state systems.
2. Rule-making. Could dependence on a supplier make EU rules harder to enforce in practice? The report documents possible channels of influence but does not measure their effect.
3. Frontier option value. Would a European near-frontier model and research-scale compute provide material insurance if AI accelerates research or open-weight supply changes? The value depends on future capability and cost.

These are questions about trade-offs and evidence. Governments can reasonably weigh them differently; this report does not choose an objective function for them.

For Sweden, the same questions can be examined through the conversion indicators in Part XIV.

What Europe is actually doing announcements and delivery

A common criticism holds that European politicians are "asleep" on AI. That is a political judgement, not a finding of this report. The record shows substantial activity. The more useful question is how that activity compares with its own objectives on **delivery, scale, timing, access, energy, financing and adoption**.

Initiative	Stated objective	Status, September 2026	Assessment [INFERENCE]
AI Continent Action Plan (Apr 2025)	Triple EU data-centre capacity in 5–7 years; raise enterprise adoption	Being turned into law through the proposed Cloud and AI Development Act (Jun 2026), not yet adopted	Coherent on paper; relies on later legislation and budgets

Initiative	Stated objective	Status, September 2026	Assessment [INFERENCE]
InvestAI (Feb 2025)	"Mobilise €200 billion", including a €20bn fund for AI gigafactories	No public tally of the €200bn found; about €1bn of EU money currently budgeted for gigafactories, the rest depends on the 2028–34 budget	Announced scale exceeds committed public money; most of it is a mobilisation target
AI Factories (Dec 2024 – Oct 2025)	Give researchers and firms access to AI supercomputing	19 selected in three waves plus 13 "antennas"; about €10bn EU and national money (2021–27); systems being procured and installed 2025–27, e.g. Sweden's contract signed Apr 2026	Real delivery under way; scale modest relative to private compute
AI Gigafactories (call 30 Jul 2026)	Up to seven sites, each with far more AI chips than today's largest European AI factories	76 expressions of interest from 16 member states; EuroHPC regulation amended 16 Jan 2026; bids due 12 Nov 2026; selection early 2027; first construction 2027	Capacity unlikely to operate before about 2028; EU share capped at 17% of each gigafactory's capex
Apply AI Strategy (Oct 2025)	"AI first" in public decision-making; ten sector flagships; SME adoption	Implementation; enterprise adoption rose from 13.5% (2024) to 20% (2025)	Adoption breadth rising; depth not systematically measured at EU level
Data Union Strategy (Nov 2025)	Data labs and data spaces	About €100m for data spaces	Small budget relative to the role of data in adoption
Data Act	Portability and fair switching	Applies since 12 Sep 2025; switching charges removed from 12 Jan 2027	Delivered in law; effect on lock-in not yet observable
Digital Omnibus on AI (Reg. 2026/1744)	Simplify and delay parts of the AI Act	In force since 27 Jul 2026	Reduces near-term compliance uncertainty; standards still late
Scale-up finance (Scaleup Europe Fund; Savings and Investment Union)	Late-stage capital so firms can scale in Europe	Fund has co-led rounds in Mistral and Lovable; capital-markets integration ongoing	Early signs; scale small relative to US late-stage capital

Sources: [Commission, AI gigafactories](#); [EuroHPC AIGF call](#); [AI Factories](#); [Apply AI](#); [Data Act](#); [Regulation \(EU\) 2026/1744](#); [EuroHPC, Sweden contract](#). The Commission's gigafactory page counts 77 proposals; EuroHPC material gives 76.

What the record shows **[INFERENCE]**:

- **Announcements run ahead of delivered capacity.** Much of the public money for gigafactories sits in a budget not yet agreed, and the first sites will not operate before about 2028.
- **Compute receives more headline attention than the largest levers.** Adoption depth, energy and grid capacity, and scale-up finance matter most for productivity and competitiveness, but receive smaller or less visible budgets and little systematic measurement.
- **Energy is a binding constraint that AI policy does not control.** Industrial electricity prices remain roughly twice US levels, and grid connections are slow.
- **Comparisons with US private capex show the scale gap but are not like-for-like.** A multiyear public programme is not the same as annual private global investment.
- **The response is European and national at once.** Much of the delivery — skills, public-sector adoption, grids, procurement — sits with member states, so EU-level scorecards capture only part of it.

Trade-offs that shape the outcomes

Table 23 maps tensions between adoption, dependence, cost, safeguards, energy and finance. It identifies evidence that would clarify each trade-off; it does not prescribe a preferred policy package **[INFERENCE]**.

Trade-off	Tension	What would show whether users retain value?
Adoption vs dependence	Using the best (mostly US) tools deepens lock-in	Evidence on switching costs, prices, depth of use and supplier income
Sovereignty vs cost and capability	Sovereign offerings cost more and lag on features	Cost and quality of alternatives for sensitive and ordinary workloads
Speed vs safeguards	Faster deployment raises risks in high-stakes uses	Measured effects of safeguards on harms, entry and deployment
Frontier ambition vs adoption budget	Public money for a frontier lab competes with money for diffusion	Comparative costs and benefits of frontier and diffusion investment
Hosting vs value capture	Foreign-owned data centres bring capex but little control	Net local value, grid pressure and legal control
Energy for AI vs energy for households and industry	Data-centre demand can raise prices in constrained zones	Opportunity cost of power in constrained regions
Openness to talent vs political constraints	Europe needs AI talent; migration politics is tight	Observed talent availability and integration outcomes
Capital-market depth vs national control	Scale-up finance comes largely from US funds	Location of research, ownership and scale-up income

The analytical question is which outcomes European initiatives actually deliver: productive use, operating capacity, competition, dependable access and measurable value for residents. Announcements and budgets alone cannot answer it.

Part XVI Regulation what the evidence shows and does not show

The 14 September post attacks "guardrails". The inference "AI matters geopolitically, therefore regulation is bad" does not follow logically. But the opposite inference — "regulation is not a problem" — does not follow either. The existing evidence does not identify the causal effect of EU AI rules on European AI performance. This part reviews the strongest evidence on both constraints and safeguards.

Where the EU stands (September 2026)

- The AI Act entered into force in August 2024. Prohibitions applied from February 2025; obligations for general-purpose models from August 2025; Commission fining powers over general-purpose model providers from 2 August 2026 [\[FACT\]](#).
- The **Digital Omnibus on AI** (Regulation (EU) 2026/1744, in force since 27 July 2026) delayed high-risk obligations to 2 December 2027 (stand-alone systems) and 2 August 2028 (AI in products). It extended SME relief to small mid-caps, softened the AI-literacy duty and *added* a prohibition on AI-generated non-consensual intimate content and child sexual abuse material. Its own recital blames late standards and late national set-up for a burden "heavier than expected" [\[FACT\]](#).
- Implementation lags: no harmonised standard had been cited in the Official Journal by June 2026, and only 8–9 of 27 member states (trackers differ) had fully designated national authorities [\[FACT\]](#).
- Twenty-one companies signed the voluntary Code of Practice for general-purpose models, including OpenAI, Anthropic, Google, Microsoft, Amazon and Mistral. Meta refused; xAI signed only the safety chapter [\[FACT\]](#).
- In late August 2026 the AI Office sent its first formal requests for information to more than 30 model providers [\[FACT\]](#).

This is "delay and simplify", not deregulation.

Evidence on constraints

GDPR — the one EU digital law with a substantial ex-post literature [\[ESTIMATE\]](#). Venture deals in EU tech fell about 26% relative to the US (Jia, Jin and Wagman). About a third of apps left Google Play and new app entry halved (Janssen et al.). EU firms cut data storage by 26% (Demirer et al.), equivalent to a 22% rise in the cost of data but only a 0.1–0.6% rise in total production costs. Exposed firms' profits fell about 8%, with small IT firms hit about twice as hard, while Google gained market share (Frey and Presidente; Peukert et al.). The costs are real, concentrated on specific margins and regressive by firm size. The studies do not measure an aggregate productivity loss.

A model-based European estimate. The IMF finds that national and EU regulation on occupational requirements, AI safety and data privacy "could reduce Europe's productivity gains by over 30 percent if AI exposure were 50 percent lower in tasks, occupations and sectors affected by regulation" [\[MODEL OUTPUT\]](#) ([IMF WP/25/067](#)). Version

1 cited this result but did not engage with it. It shows that regulation *could* matter a great deal. But the size of the effect is assumed (a 50% cut in exposure), not observed, so it is not a measured cost.

AI Act compliance costs are known only in advance, not after the fact. The Commission estimated €6,000–7,000 per high-risk system plus €3,000–7,500 for verification. The larger burden for small firms is a quality-management system costing €193,000–330,000 to set up. The widely quoted "€31bn" cost estimate came from a method its own source study's authors called misleading and should not be used [ESTIMATE]/[CONTESTED].

Size thresholds cause cliffs. French rules triggered at 50 employees reduce aggregate innovation by about 5.7% (Aghion, Bergeaud and Van Reenen, AER 2023) [ESTIMATE]. Hard SME carve-outs create new cliffs; phase-ins avoid them.

Fragmentation and uncertainty. Draghi counted about 100 tech-focused laws and more than 270 regulators. Italy's 2023 ChatGPT ban cut Italian developers' output by about half for two days [ESTIMATE]. Among firms that considered AI, 53% cited unclear legal consequences as a barrier (Eurostat 2025) [FACT].

Evidence on safeguards and trust

Stated trust. In a 47-country survey (KPMG/University of Melbourne 2025), 70% said AI regulation is needed and 80% would trust AI more if assurance mechanisms were in place [FACT] (stated preferences, not behaviour).

Regulatory design can help firms. Firms in the UK financial regulator's sandbox raised 15% more capital (BIS research) [ESTIMATE]. This suggests *how* one regulates matters.

Harms that rules target are real. The 2026 prohibition on AI-generated sexual-abuse imagery addresses documented harm with few legitimate business models affected. Vendor-reported AI-enabled cyber activity (Part X) and government model-security findings (NIST CAISI) show why some frontier transparency and testing duties exist.

Other powers regulate too. California (SB 53, September 2025) and New York (RAISE Act, December 2025) have enacted frontier-AI transparency and incident-reporting laws. Colorado's deployer-side law ran into the same implementation problems as the EU's high-risk regime. The December 2025 US executive order fights over *who* regulates (federal or state), not *whether*. China requires algorithm filing and has mandated AI-content labelling since September 2025, and it competes at the frontier [FACT].

What the comparisons cannot show

- **Adoption parity.** EU and US firms use generative AI at the same rate (37% vs 36%, EIB Investment Survey 2025) [ESTIMATE]. That shows basic use is not lower in Europe. It does not identify the effect of regulation on investment, firm entry or advanced deployment, where Europe does lag.
- **Variation within the EU.** Adoption ranges from 42% of firms in Denmark to 5% in Romania under shared EU law. This shows that factors other than EU-level rules matter. It does not prove that non-regulatory factors explain most of the variation, because national implementation, enforcement and other national laws also differ. Version 1 drew the stronger conclusion; it is withdrawn.

- **Barrier surveys.** Lack of expertise (71%) is cited more often than unclear legal consequences (53%) **[FACT]**. Firms' answers show what they perceive, not which factor causes the largest economic loss.

What remains to be measured

Type of regulation	Examples	Evidence	Open question
1. Targets clear harms with few legitimate uses affected	AI Act prohibitions; deepfake labelling; the 2026 ban on AI sexual-abuse imagery	Similar rules adopted by US states and China	How does enforcement change harm and trust?
2. Imposes significant costs on specific margins	GDPR as enforced for data-intensive entry; broad high-risk categories	GDPR studies: 10–30% effects on specific margins	Which costs fall on entrants versus incumbents?
3. Creates fixed costs that favour incumbents	Quality-management and documentation obligations	GDPR concentration effects	How large are fixed costs by firm size?
4. Fragments the single market	27 enforcers; national gold-plating	IMF, Draghi, ChatGPT ban	How much fragmentation is due to national practice?
5. Builds trust and thus adoption	Harmonised standards; product liability; sandboxes; transparency	Surveys; sandbox evidence	Does assurance measurably change adoption?
6. Treats frontier development differently from ordinary deployment	General-purpose model duties; California SB 53	Frontier compliance costs small relative to \$1bn training runs; all major US labs except Meta signed the Code (xAI in part)	Which duties affect frontier and ordinary deployment differently?

The evidence does not support a single verdict that regulation is either Europe's main obstacle or harmless. Studies document costs on particular margins, while surveys and examples suggest reasons for safeguards and predictable rules. No ex-post AI Act cost estimate exists yet. Table 24 sets out what is known, what is assumed and what remains to be measured [JUDGEMENT].

The post's attack on guardrails and an unqualified defence of every rule are both too simple. The useful question is which specific obligations create costs or benefits, for whom, and how those effects change with design and implementation.

Part XVII Competing interpretations

Five perspectives, each in its strongest form.

1. The AI race / national-power view. *Claim:* AI leadership will strongly shape national power. *Strongest evidence:* concentration and high margins in chips, cloud and compute; the US hosts most documented AI-supercomputer capacity and attracts 83% of private AI investment; AI is already used in intelligence, targeting and cyber operations; export controls work as leverage; frontier AI might accelerate research. *Strongest counter-evidence:* frontier leads are measured in months; Chinese models remain competitive despite controls; macro productivity gains are not yet visible. *Assessment:* right that control of scarce layers and security uses matters; unproven on the size and persistence of economic advantages from model leadership.

2. The diffusion view. *Claim:* the biggest economic gains go to effective adopters, not inventors. *Strongest evidence:* Ding's history of industrial revolutions; Nordhaus on innovators' small share of social returns; US consumer-surplus estimates; Comin and Mestieri on intensity of use explaining income gaps; the OECD's finding that competitiveness in AI-exposed sectors requires domestic adoption. *Counter-evidence:* AI's complementary assets (compute, chips) are scarce and concentrated, unlike electricity's; adoption on foreign platforms creates dependence; research acceleration could make invention matter more. *Assessment:* the best guide to productivity; incomplete on control and on the income split, which is not measured.

3. The commoditisation view. *Claim:* models will become cheap and ubiquitous, and frontier ownership is overstated. *Strongest evidence:* price declines of 9–900× a year for given capability; open weights months behind; rapid shifts in enterprise market share; DeepSeek. *Counter-evidence:* frontier usage is price-inelastic; the best models still command premiums; frontier training costs rise ~3.5× a year, concentrating the frontier among a few firms. *Assessment:* describes the model layer today reasonably well; would be overturned if frontier systems began to accelerate research and compound their lead.

4. The safety and governance view. *Claim:* treating AI as a geopolitical race encourages excessive risk-taking and underinvestment in safety. *Strongest evidence:* leading lab heads themselves reportedly called in September 2026 for pacing the frontier; vendor-reported AI-enabled cyber operations; the US administration's recasting of guardrails as presidential discretion; race dynamics erode international coordination. *Counter-evidence:* unilateral restraint by democracies may hand the lead to others; safety rules can be designed not to slow diffusion. *Assessment:* the race framing is a choice, not a law of nature. Europe has an interest in safety standards that do not become barriers to adoption, and in not being the only party that pays for them.

5. The industrial-policy view. *Claim:* AI matters, but energy, manufacturing, capital and institutions matter more than models. *Strongest evidence:* China's lead in robots and electricity; Europe's AI-era revenues so far come from power equipment and lithography; energy costs are a barrier for 75% of EU firms against 36% of US firms; diffusion depends on management and skills. *Counter-evidence:* the automotive precedent shows industrial strength does not protect the software layer. *Assessment:* correct that AI's economic

value is realised through physical and institutional capacity; wrong if it implies models do not matter.

Synthesis [JUDGEMENT]. These perspectives describe different layers more than they contradict each other. Concentration is visible where scarcity persists: chips, fabrication, lithography, compute, cloud and distribution. Productivity depends on deep use across the economy. The model layer is more contestable today but could become more consequential if frontier AI accelerates research. Institutions, safeguards and industrial capacity affect how diffusion occurs and where value accrues. The evidence does not yet determine one European course.

Part XVIII The answer to the title question

What the phrase gets right

- **Productivity differences compound into large differences in resources.** A one-point growth advantage makes an economy about a fifth larger in twenty years. Institutions then decide how much of that becomes defence, research and influence.
- **Some layers of AI are highly concentrated today.** Chips, fabrication, lithography, hyperscale compute and distribution show high concentration and margins. Control of them is real leverage, and it has been used.
- **Security uses matter.** AI is already used in intelligence, targeting, drones and cyber operations.
- **The US government treats AI as an industrial and strategic project.** Its published plans aim to build the largest ecosystem, export the full American stack to allies and use AI to accelerate science (Part 0). Europe will be affected whether or not it agrees.

What it oversimplifies

- **There is no single race and no single winner.** Leadership is split by layer: the US leads compute, chips, private capital and frontier models; China leads robots, electricity, manufacturing and open-weight releases.
- **Winning is not only owning the best model.** Model leads between firms shift quickly, although a national lead could become more valuable if frontier AI accelerates research.
- **Much of the economic payoff depends on diffusion.** Effective use across the economy is open to any country that organises well. How gains are divided between users and suppliers, and between countries, is not measured.
- **The evidence does not show that safeguards decide who wins.** It documents real costs from specific rules, from fragmentation and from uncertainty, and real reasons for some safeguards.
- **A race framing can be self-defeating.** It can undermine the safety and international coordination that sustained adoption requires.

What the comparison reveals for Europe

1. Productivity is consequential. Europe's output-per-hour growth has trailed the US for most of three decades. AI could narrow or widen that gap depending on how deeply it is used and how large the technology's eventual effect proves to be.
2. Use and ownership create different gains. European users of American or Chinese AI can create local value while suppliers earn income abroad. Competition, portability and the location of complementary businesses influence the balance.
3. Control matters most where interruption would be costly. A model, cloud service or component may be replaceable for one use and difficult to replace for another. The report does not establish one boundary for all sectors.

4. Frontier leadership has uncertain option value. Its importance grows if frontier systems accelerate research or become central to security; it falls if capable models remain abundant and cheap.
5. Europe already supplies important parts of AI-era value chains, including lithography and power equipment. Whether industrial firms retain the differentiating software layer as they adopt AI remains an open question.
6. Regulation has real costs and possible benefits. Current evidence is too limited to rank it against capital, market scale, compute, energy and skills as an explanation of Europe's position.
7. The observable test is delivery: operating capacity, depth of adoption, grid access, scale-up outcomes and continuity of critical services.

For Sweden, the question is sharper. It has assets that could make it one of Europe's strongest AI adopters, and constraints in grid capacity and scale-up capital. Whether its assets convert into Swedish productivity, incomes and control is a question to measure.

"Whoever wins AI, wins" captures a real contest for capability and influence, but there is no single prize. The United States, China and Europe occupy different positions in a changing system. Europeans can benefit from technology developed elsewhere while still needing to understand how productivity, supplier income and control are distributed. This report is an attempt to make those mechanisms visible, and to show where the evidence still stops.

Annex A Chart list

No.	Title	File	Status in v2
1	Small differences in growth become large differences in output	C01_compounding.png	Unchanged
2	US output per hour grew faster in 1995–2005 and again after 2019	C02_eu_us_productivity_v2.png	Revised: real growth index; single-year level benchmark
3	Where the gap shows up: an accounting view	C03_what_caused_the_gap_v2.png	Revised: headline softened
4	Private AI investment by region	C04_private_ai_investment.png	Unchanged; caption caveat
5	Who builds the models?	C05_models_by_origin.png	Reproduced from Epoch dataset updated 23 Sep 2026
6	Three US providers hold ~70% of Europe's cloud market	C06_cloud_market_v2.png	Revised: European revenue growth shown
7	Compute: location, ownership, access and electricity	C07_compute_datacentres_v2.png	Revised: separate measures; coverage caveat
8	The AI value stack	C08_ai_value_stack_v2.png	Revised: unsupported claims removed
9	Supplier profitability: selected examples	C09_supplier_profitability_v2.png	Rebuilt
10	Using foreign AI: who receives, owns, controls?	C10_value_capture_stakeholders_v2.png	Rebuilt as stakeholder map
11	Published AI productivity estimates in their own units	C11_productivity_estimates_v2.png	Rebuilt as separate panels
12	Early or late: AI diffusion arithmetic	C12_early_vs_late_v2.png	Assumptions made explicit
13	Labour share and labour income in AI scenarios	C13_capital_vs_labour_v2.png	Labour-income panel added
14	From productivity to power	C14_productivity_to_power_v2.png	Rebuilt with institutions
15	Europe's strategic options: conditions, not scores	C15_europe_strategy_options_v2.png	Rebuilt as discussion tool
16	Sweden's conversion question	C16_sweden_conversion_v2.png	Rebuilt: complementary inputs, outcomes
S1	Swedish wholesale power prices by zone	S1_sweden_power_prices_v2.png	Qualified
S2	Enterprise AI adoption by EU country	S2_ai_adoption_by_country_v2.png	Neutral title

Underlying values, sources, definitions and caveats are in the version 1 chart-pack workbook and, for revised charts, in the version 2 source-table workbook.

Annex B Principal sources with links

Dates are publication dates; page or table references are given where the report relies on a specific passage. Full URLs for every source remain in the research dossiers.

US doctrine and the quotation

- White House, [America's AI Action Plan](#) (July 2025), p. 1 (introduction; three pillars), p. 5 (adoption as bottleneck), p. 20 (exporting the full stack).
- [Executive Order 14320, Promoting the Export of the American AI Technology Stack](#) (23 July 2025), sections 1–3.
- [Executive Order 14363, Launching the Genesis Mission](#) (24 Nov 2025), section 1.
- Truth Social post of 14 Sept 2026, via the [archive copy](#); [Townhall report, 14 Sept 2026](#); [Roll Call Factbase transcript, 14 Sept 2026](#) (00:03:54).
- [BIS, rescission of the AI Diffusion Rule](#) (13 May 2025); [Federal Register, Framework for AI Diffusion](#) (15 Jan 2025).

European productivity

- European Commission, [AMECO database](#), Spring 2026 vintage (OVGD, NLHT, UVGD); author's calculations in the workbook.
- ECB, [Economic Bulletin 6/2024, Box 1](#) (23 Sept 2024).
- I. Schnabel, [speech of 16 Feb 2024](#).
- A. Bergeaud, [ECB Forum on Central Banking paper](#) (July 2024).
- M. Draghi, [The future of European competitiveness, Part A](#) (Sept 2024), pp. 23–28.
- Eurostat, [national accounts metadata on PPS](#); [bop_its6_det](#) (services trade with the US, updated 16 Apr 2026).
- Conference Board, [Total Economy Database summary tables](#) (Sept 2025).

AI productivity and distribution

- ECB, [P. Lane, speech of 23 Mar 2026](#).
- IMF, [Misch, Park, Pizzinelli and Sher, WP/25/067](#) (Apr 2025).
- OECD, [Filippucci et al., Nov 2024](#); [Schief et al., G7, Jun 2025](#); [Nicoletti, The impact of AI on productivity, distribution and growth, Apr 2024](#); [Yildirim, AI meets trade, Mar 2026](#).
- [Acemoglu, Economic Policy \(2025\)](#); [Penn Wharton Budget Model \(8 Sept 2025\)](#); [Aghion & Bunel \(2024\)](#).
- [Korinek, Jones, Sacher, Cotter & McCrory, Anthropic Institute WP 2026-02](#) (Sept 2026).
- [Brynjolfsson, Collis et al., "What is Generative AI Worth?" \(Apr 2026\)](#); [Brynjolfsson, Chandar & Chen, "Canaries in the coal mine" \(rev. Aug 2026\)](#).
- [BLS, Productivity and Costs \(Q2 2026\)](#).
- [EIB, EU firms, AI and big data](#); [EIB Investment Survey 2025](#).

Value capture, ownership and tax

- CSO Ireland, [Annual National Accounts 2025, GNI and globalised results; Modified GNI explained](#).
- Irish Revenue, [Corporation Tax 2025 analysis](#), pp. 8 and 11.
- OECD, [international community agreement on the global minimum-tax package and side-by-side package](#) (5 January 2026).
- [US Treasury, foreign holdings of US securities](#) (June 2025).
- Company results: [Nvidia Q2 FY2027](#); [TSMC Q2 2026 6-K](#); [Amazon Q2 2026](#); [ASML Q2 2026](#); [reported OpenAI 2025 audited accounts](#) (documents not publicly filed).

Stack, compute and dependence

- Epoch AI, [large-scale models dataset](#); [Trends in AI supercomputers](#) (23 Apr 2025); [hyperscaler compute ownership](#) (Apr 2026); [US vs China capability gap](#) (Jan 2026); [Cohere–Aleph Alpha business-combination agreement](#) (16 Sept 2026).
- Stanford HAI, [AI Index 2026](#); IEA, [Energy and AI](#) (Apr 2025); [Synergy Research](#) (24 Jul 2025); [IFR robot-density release](#) (8 Apr 2026).
- Commission, [Data Act explained](#); [COM\(2026\) 503](#); [EDPB Recommendations 01/2020](#).
- NIST CAISI evaluation of DeepSeek (Sept 2025); NATO NCIA on Maven (Mar 2025); [NGA testimony on Maven users](#) (May 2025); Anthropic, disrupting AI espionage (Nov 2025).

EU–China relations: [EEAS, EU-China Relations factsheet](#) (7 December 2023); Chinese official plans discussed in Part IX.

EU policy and regulation

- Commission, [AI gigafactories](#); EuroHPC, [gigafactories call](#) (30 Jul 2026); [AI Factories](#); [Apply AI](#).
- [Regulation \(EU\) 2026/1744 \(Digital Omnibus on AI\)](#).
- [KPMG/University of Melbourne, Trust in AI \(2025\)](#).

Sweden

- Sweden: [Statistics Sweden 2025 government debt and nominal GDP](#); [Statistics Sweden tax-to-GDP table](#); [Productivity Commission, SOU 2025:96](#); [government defence-spending path](#); Eurostat, energy and grid sources listed in the research dossier.

Annex C How to read the evidence

This report is for curious readers who want to understand why AI might matter economically and politically for Europe. It starts with the mechanics of productivity, technology diffusion, value creation and the AI stack, then applies them to Europe and Sweden. No specialist knowledge is assumed.

Every substantive claim carries an evidence label. The labels matter more than usual because the AI debate freely mixes measured facts, model outputs and marketing. A label describes the claim it follows. It does not certify the inferences built on that claim. A verified event, for example, does not prove every cause attributed to it.

Label	Meaning
[FACT]	Observed and recorded: official statistics, filings, legal texts, events
[ESTIMATE]	A measurement with material uncertainty (survey, commercial database, derived calculation)
[MODEL OUTPUT]	The result of an economic model; true only if the model's assumptions hold
[FORECAST]	A central projection by an institution or company
[SCENARIO]	An explicitly conditional "what if" path; not a prediction
[CONSULTANCY]	Estimates from banks or consultancies; methods partly opaque
[CONTESTED]	Serious researchers disagree
[OFFICIAL CLAIM]	A government or institution's stated position or objective
[INFERENCE]	The author's own reasoning from the evidence
[ILLUSTRATION]	The author's arithmetic, used to teach a mechanism; never a forecast
[JUDGEMENT]	The author's explicit interpretation or assessment; not a measured fact. Research version 2 contains no numerical probabilities

Geography is kept strict. "EU27" means the 27 member states. "EA20" is the fixed historical grouping of 20 euro-area countries used in the time series; the euro area has had 21 members since Bulgaria joined on 1 January 2026. "Europe" is used only where a source uses a wider region (for example Stanford's AI Index region, which includes the UK, or the IEA's, which includes Norway and Switzerland). These are not interchangeable and are not compared as if they were.

Present versus future. The report separates what is true today from what may become true. Where a sentence is about the future, it says whose projection it is and on what assumptions.

A note on sources. Several sources are published by AI developers, including Anthropic's Economic Index and the Anthropic Institute's scenario paper. They are used because they contain unique data. They are always paired with independent estimates (OECD, IMF, ECB) and flagged, because their publishers have a commercial stake in the subject. Direct links to the principal sources are in Annex B. The principal sources and the status of material estimates are stated in the text and Annex B.