Europe 2031

What getting AI wrong means for us

By Daan Juijn, Stan van Baarsen, Judith Dada, Philip Fox, Alex Petropoulos, and Michiel Bakker.

Copywriting and editing by Tom Chivers.

The current trajectory of AI calls for the most ambitious political agenda in the history of post-war Europe. Unless we embark on it now, Europe will lose the ability to shape its own future. We will end up economically and politically sidelined, with values we cannot defend, social welfare systems we can no longer fund, risks we cannot address, and a Union that cannot hold.

> **March 2031 - Washington, D.C.**
>
> *Caroline splashes cold water on her face and looks at herself in the bathroom mirror. Her hands are shaking. She grips the edge of the sink and waits for it to pass. Through the small high window she can see a slice of Washington sky, flat and bright.*
>
> *Down the corridor, six people are deciding the fate of the European continent. She does not know whether anything she said will matter.*
>
> *She suspects it will not.*

This is a scenario about Europe's impending slide into irrelevance: how AI is driving it, and what can still be done to change course.

It is told through the fictional characters of Caroline Dubois, a French policy worker in Brussels, and Christian Vogt, the German founder of a fast-growing AI company who has just relocated to Silicon Valley. While their personas are made up, the events they experience are grounded in existing trends, and meant to be realistic.

To understand how Europe is at risk of squandering the coming AI revolution, we first have to return to early 2025 - because the story we're about to tell is already well underway.

January 2025 - DeepSeek and ye shall find

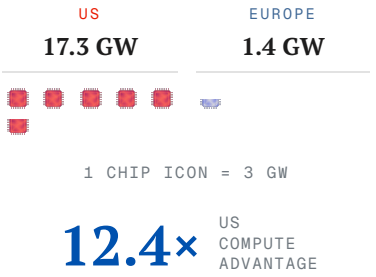

Caroline Dubois' office is buzzing. This morning, a Chinese firm, DeepSeek, released a new AI model, R1. It is cheap and effective, and – although no European firm has played a part in developing it – Brussels is in a fervour.

Caroline works on digital technology at the Directorate-General for Trade and Economic Security – DG TRADE, the part of the European Commission that deals with things like tariffs and exports controls. At twenty-eight, she has been there for three years, and feels she is earning the respect of her peers and superiors; but she, unlike them, is seriously worried about Europe's future, and the DeepSeek news has not reassured her. She has recently visited Silicon Valley. It is 9,000 kilometres from Brussels, but it feels further. The idea that AI is sparking a new industrial revolution is a truism in California; in the European Commission offices, it is bordering on science fiction.

R1's rise has excited her colleagues because they see it as clear proof that it is possible to train cutting-edge AI without the resources of the Silicon Valley giants. European policymakers jump at the idea that it is possible to outsmart the Americans, to be small and nimble and clever, to do better without the hundreds of billions of dollars that have been pouring into indulgent American data centres and gigantic model training runs.

The idea makes sense. DeepSeek has supposedly developed R1 at a fraction of the cost of OpenAI's ChatGPT, or Anthropic's Claude, but as far as anyone can tell, it is almost as powerful. It even has the new 'reasoning' feature that the American models have started rolling out, in which the model lays out its thought process on a virtual scratchpad. And its weights – the mathematical values that make the model what it is – are publicly available. Anyone can run the model on their own infrastructure, free from American tech dependence.

Though somewhat energised by the hopeful mood in Brussels, Caroline isn't quite ready to join in. She carefully listens to more cautious voices who point out that big efficiency gains are hardly unheard of – Anthropic, Google DeepMind, and OpenAI find them all the time. DeepSeek has smart researchers and moved fast, but will soon find itself constrained by computing power; China simply doesn't have enough 'compute' to train models ever since America restricted exports on the cutting-edge AI chips

that China can't produce domestically. They argue that finding efficiency gains is much easier if you have lots of compute to help you find them. DeepSeek's own CEO pointed out that their main problem is the US ban on AI chip exports.

And so Silicon Valley is largely unfazed by R1. The hyperscalers do not slow down their massive AI investments. Days after R1, OpenAI releases o3-mini, a more impressive model than R1, and a sign that American progress is still rapid.

In Brussels, the o3 news barely registers. Acknowledging it would mean accepting that AI acceleration is continuing, driven by huge American companies, led by CEOs that can't be trusted.

Caroline isn't sure which side to take. She was in California three weeks ago, and what she saw there has stuck with her.

She stayed with Christian Vogt, an old friend from her study-abroad programme in Berkeley half a decade earlier; she from Sciences Po, he from TU Munich, two of Europe's leading universities. Christian had moved to San Francisco three years ago, just as she started at the Commission, to found a company building image and video models. The AI startup was still tiny, but it had recently closed a Series A funding round and Christian, a gregarious type, knew *everyone*.

The whole visit had startled her. The hours Christian's team pulled were insane – seventy- or eighty-hour weeks, people sleeping in the office. And when he took her to a dinner party at the house of another startup founder in Hayes Valley, she found the conversations almost incomprehensible; not because of the technical details but because of the deep belief that the world was going to change radically. 'By January 2026, I think most of my code will be written by Claude,' one guest said in a matter-of-fact tone. The host said he had stopped hiring junior engineers because ChatGPT would soon do the entry-level coding better. Someone mentioned, in passing, that they thought artificial general intelligence - AI that is better than any human at most tasks - was probably two or three years out. Caroline had asked what this would mean for Europe, given it lacks a competitive AI sector. But nobody at the table had really thought about Europe, except as a potential compliance headache.

In the Waymo back to Christian's apartment, Caroline had told him she found the whole visit disorienting. Christian had laughed and reassured her that after a few months in Silicon Valley, she would start to 'feel the AGI' herself. Caroline had wondered if Christian and his peers were part of a strange bubble, a cult even, or if she was living in a world that was already disappearing into the past.

Back in Brussels, she had considered telling her director about the experiences she had had on her trip, but it felt untranslatable. The people she had spoken to were some of the smartest people she had ever met. They were calm, they were sharp, and they seemed well-informed. There was an urge inside her to take what they had told her deeply seriously and to get others to do the same. And yet their claims would have sounded insane in a Commission conference room halfway around the world. So she said nothing.

The day after the R1 release, her phone buzzes. It's Christian.

Christian: ok so brussels is doing a deepseek victory lap???

Caroline: My director thinks it proves we can catch up…

Christian: lol

Christian: the lesson is reasoning models work and china can build them

Christian: also, o3-mini is better and nobody is talking about it

Christian: your director should come to SF

Caroline: I wish.

She puts the phone down. She isn't sure Christian is right, but she isn't sure he's wrong, either. The voices of caution about R1 sound reasonable to her. So do the voices of excitement.

February 2025 - Plug, baby, plug

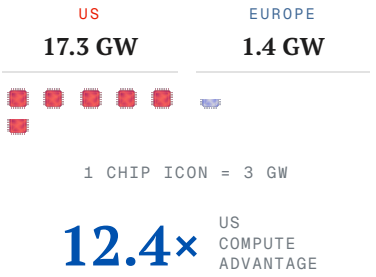

Three weeks after R1, Emmanuel Macron hosts the AI Action Summit in Paris.

The summit series was founded around safety and international coordination – the 2023 edition in the UK had been called the AI Safety Summit – but the 2025 edition is different. The geopolitical environment has hardened. Russia is grinding forward in Ukraine; Donald Trump has been back in the White House for a fortnight and is already threatening tariffs on European goods. Europe is now focused on competitiveness, not safety. And the timing is good: DeepSeek has apparently shown that catch-up is possible, and could perhaps even be cheap. Europe wants to signal it is taking AI seriously.

European Commission President Ursula von der Leyen announces a €200 billion InvestAI Fund, which includes a €20 billion AI Gigafactories Initiative to build four to five large AI data centres on European soil. Macron sells France as the best place in Europe to build AI: with all its nuclear power, developers can just 'plug, baby, plug.'

After three years in Brussels, Caroline knows to look beyond the headline numbers. The €200 billion is largely a repackaging of existing funds, and an assumption that European industry will pour its own money in. Only a fraction is genuinely new, and even that is spread out over five years. She shudders when comparing it to the data centre investments announced by American hyperscalers, projected to exceed $400 billion in 2025. Still, she hopes the vibes are shifting, and that Europe is waking up.

Christian: 200 billion?

Caroline: It's mostly aspirational.

Christian: the americans have real money

Christian: the texas thing is being built right now

Christian: like there are bulldozers

Caroline: I know.

Christian: how does brussels feel about altman, ellison, masayoshi son on a stage with trump

Caroline: Well, not thrilled.

Whether the vibes have shifted or not, some of the positive energy is overshadowed when US Vice President JD Vance gives an aggressively anti-Europe speech at the summit. He repeats the trick two days later at the Munich Security Conference. European leaders in the audience are seen on camera almost visibly biting their tongues.

In the weeks after, Caroline's bosses come to a conclusion. The transatlantic relationship is effectively broken. Washington can no longer be trusted, not in defence, not in energy, and certainly not, it is now clear, in tech. 'Sovereignty' becomes the new buzzword in European capitals.

But sovereignty is easier to announce than to build, and whether the push will amount to anything is not clear, at least in AI. Caroline suspects that while leaders profess to care about the new technology, their old scepticism is very much still in place beneath the surface. One of her colleagues, a thoughtful and senior policy officer whom she likes, implies during a lunch shortly after the Vance speeches that he thinks you probably don't need vast data centres to make AI work, precisely because Sam Altman, Larry Ellison and Donald Trump all stood on a stage to say you do. It is, Caroline thinks, a simplistic but understandable position: Brussels' scepticism of Trump and the Silicon Valley elite runs deep. But she is worried that just because they say it's sunny, that doesn't mean it's dark outside.

August 2025 - Forever blowing bubbles

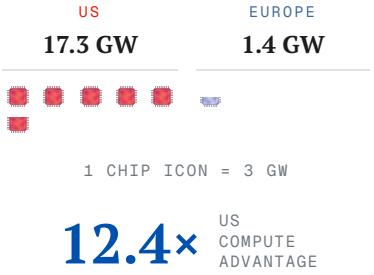

The European AI excitement subsides after Paris. The summit was dramatic, but the months afterward are filled with slow, dense, and unglamorous policymaking. Building a continental AI ecosystem turns out to require things that announcements alone cannot conjure – talent, capital, energy supply - as well as the know-how and drive to coordinate them. The largely illusory nature of the €200 billion is widely known now, and most of the air has gone out of the initial elation.

In the US, there is no such slowdown. The battle for talent heats up: Meta, which has fallen behind the AI frontier somewhat, is poaching top researchers from rivals on salaries that make Premier League footballers look underpaid. Mark Zuckerberg even personally cooks soup for them to lure them in. The Trump administration publishes its national AI strategy, titled 'Winning the Race,' which describes the AI transition as 'an industrial revolution, an information revolution, and a renaissance – all at once.'

Europeans, who consider themselves more grounded than Americans, find the rhetoric overheated and the messengers untrustworthy. Caroline hears the phrase 'AI bubble' regularly. Her colleagues have seen tech leaders overstate their products' usefulness before: 'Remember NFTs?', one says. 'Those pictures of monkeys that people bought for $50,000?'

Christian: how's the bubble talk going

Caroline: It's everywhere. My director said it three times in a meeting yesterday.

Christian: wait until they see anthropic's revenue numbers

Christian: they're going to do 9 billion this year

Christian: from like 1 billion last year.

Christian: that's frickin huuuuge

Caroline: I'll mention it.

Christian: you won't

She does mention it. No one seems very excited, though. Some colleagues point out that Anthropic is still not making a profit.

The 'bubble' narrative is strengthened by the release of GPT-5. OpenAI hypes the model hard: Sam Altman posts an image of the Death Star; employees talk about the 'year of the AI agent'. But to the public, the model itself is somewhat disappointing, essentially a more polished version of the existing o3 model. It still hallucinates and makes silly mistakes.

In Brussels, the AI sceptics take another victory lap. 'I told you AI was hitting a wall,' Caroline's colleague tells her over a beer at Place du Luxembourg. He says it kindly. Caroline has, over the past few months, become known in her unit as the AI booster, which is a label she is not sure she deserves. She is not yet certain her colleagues are wrong; She agrees that GPT-5 really was disappointing. Maybe the Americans are hyping. Maybe Brussels is right.

On the other hand, some experts say that GPT-5 is broadly what you'd expect given the exponential growth curve of AI's capabilities; o3 was only a few months ago. But Europe's AI sceptics are deaf to those arguments, given OpenAI's marketing blunder in over-promising and then under-delivering.

Caroline: What do you make of GPT-5?

Christian: it was a misfire. they should have shipped o3 under that name.

Christian: but the underlying capability is still moving

Christian: we've been using o3 in our pipeline for the last months and it's wild

Money is still pouring into the big US labs and the race to build data centres is still on, but that money seems to bounce back and forth like a ping-pong ball – Oracle and OpenAI sign a $300 billion compute deal, most of which will be spent on NVIDIA chips; NVIDIA in turn agrees to invest $100 billion in OpenAI. In San Francisco, the deals are seen as proof of the tech's transformative nature – the old rules no longer apply. In Brussels, they look like the credit-default swaps of 2007: complex, circular, unstable, and likely to end badly.

That is why the French company Mistral raising €1.7 billion – most of it from the Dutch lithography machine maker ASML – is hailed as a big win for Europe, despite its funding round being twenty times smaller than OpenAI's. The bubble talk means that few Europeans are worried about AI companies having too little money. If AI and the largest investment rounds are overhyped, there's no need to hurry.

November 2025 - Worlds apart

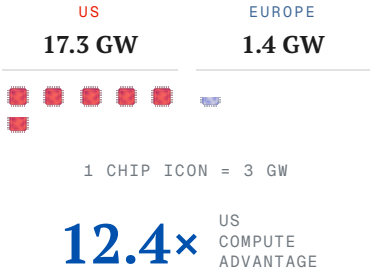

In November, Anthropic releases Claude Opus 4.5. There is no Death Star tweet, no countdown, and no event. The model is good – better than GPT-5, by most accounts – but the release is, by 2025 standards, undramatic.

What is dramatic is what people start doing with it.

The breakthrough in AI agents has arrived, but not how people imagined it. Initially, through the spring and summer, the major labs had released various products that controlled your computer directly. They would move your mouse, fill out web forms, click through booking flows. While nice in theory, these mostly did not work in reality. Many people who tried them concluded that AI agents were not yet ready for prime time.

What does work, it now turns out, is agents that write code. Anthropic's Claude Code, a terminal-based tool, becomes – once paired with the new Opus 4.5 – the surprise hit of the year. Christian's friend was right: when January arrives, most lines of software in Hayes Valley really are written by AI. Developers quickly realise that code is the universal interface: if you can write code, you can do anything a computer can do. Want to send an email? Don't navigate Gmail; Claude Code, write a script. Want to organise documents? Claude Code, write a script.

In San Francisco, this becomes common practice. People talk about 'Claude Code mania.' Developers burn through thousands of dollars a month in tokens, refusing to go to sleep without first queuing up overnight tasks for Claude to complete by morning. They believe, with reason, that they are getting the better end of the deal: the productivity gains are extraordinary. Anthropic becomes the fastest-growing company in history.

Christian: most of our code is now written by claude

Christian: i don't mean assisted i mean fully written end to end

Christian: my engineers are managing it like a junior team

Christian: except it never sleeps :D

Caroline: Is that good for your engineers?

Christian: for now yes

Christian: huge productivity boost

Christian: the labs are using it internally and their progress is accelerating

Christian: like noticeably accelerating

Christian: also they're shipping a new opus every three months now

Christian: used to be every six

Caroline reads the messages on the metro into work. By the time she gets to her desk she has decided to bring it up at her unit's morning meeting.

It does not go well. She mentions the productivity gains, the new revenue numbers, the compressed release cycle, the labs using their own AI to build their next models. The room is polite. Someone asks if she has peer-reviewed studies that demonstrate the productivity gains. Someone else gently observes that consumer-grade AI agents have been a disappointment all year. Caroline says the consumer agents are not the point; the self-improvement is. The room moves on.

Her colleagues are, in any case, effectively banned by institutional policy from using Claude or ChatGPT on work devices. Standalone American AI tools are considered a data-protection risk. The Commission offers its own 'GPT' which is essentially a wrapper around a couple of old and small open-source models. The workforce that is meant to be regulating frontier AI systems mostly cannot use them.

She walks back to her office immensely frustrated by her colleagues' obstinacy, trying to remember when, exactly, she had stopped being uncertain about AI's capabilities.

February 2026 - A war on two fronts

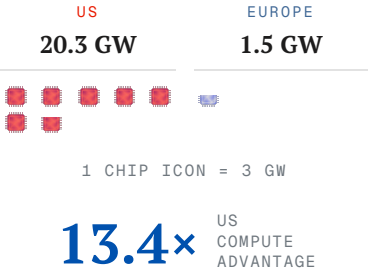

The first half of 2026 is not, realistically speaking, the ideal time for Europe to be making deliberative progress on AI policy. A lot is happening. In January, US special forces raid Caracas and abduct Nicolás Maduro. Days later Trump threatens to invade Greenland. In February, the United States and Israel bomb Iran. Oil prices surge. Trump threatens to leave NATO if European allies didn't send ships to help reopen the Strait of Hormuz.

Then the Department of War (DoW) goes to war with Anthropic.

The falling out is caused by a contract dispute. Anthropic leased its models to the DoW, and over time they have become crucial to its military operations, including those in Venezuela. The contract specified two red lines: Claude cannot yet be used for fully autonomous lethal weapons, and it cannot be used for mass surveillance of US citizens. Anthropic has allegedly insisted on these terms since the very beginning.

The Pentagon now wants to renegotiate them, and when Anthropic refuses, the dispute escalates quickly. At first, the DoW threatens to force Anthropic to provide unrestricted access through the Defense Production Act; then it does the opposite - it designates the company a supply chain risk, a definition usually reserved for companies the government considers compromised by hostile foreign powers. aAs stated by the Secretary of War, the designation prevents any other firm which works with the Pentagon from doing business with Anthropic, which would effectively destroy the company - after all, Anthropic's compute providers all work with the DoW.

Caroline: What would it take for Anthropic to move to the EU?

Christian: they won't

Christian: they're patriots. they want america to win.

Christian: you could try axing copyright regulation and labour protections

Christian: fix datacentre permitting. massive buildouts. tax breaks

Caroline: That's not going to fly. We're not Silicon Valley.

Christian: you can at least make it more attractive for them to expand in europe

Caroline: People here worry they'll move to London.

Christian: why would that be bad?

Caroline is stuck to the news. The pieces fit together unpleasantly. The US government has woken up to AI, and is now attempting corporate murder because it can't control the company building it. In Brussels, civil servants are banned from using frontier AI tools to draft memos; in Washington, they're being used to plan military operations.

Anthropic wins a preliminary injunction against the supply-chain-risk designation, but the relationship between the company and the government is now fraught.

March 2026 - The wild frontier

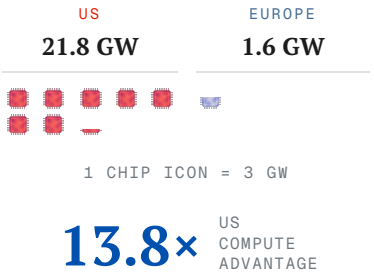

Back in November, France, Germany and the European Commission announced the Frontier AI Initiative. The plan was to establish, by the end of Q1 2026, the world's best-funded not-for-profit AI research organisation. Europe, the announcement said, has all the talent and data it needs, and is building compute through its Gigafactories. What it lacks is an initiative to bring these advantages together. Insiders were excited.

But when Q1 ends, there is no nonprofit. Building a world-class research institute is hard, especially when you're battered by one geopolitical crisis after another. What doesn't help either is that people advising the initiative give wildly conflicting advice: some say LLMs are a dead end when it comes to reaching AGI, and the nonprofit should try to find a new paradigm. Others want to find a particular niche, whether AI for science or AI for industry. Still others say something about preparing for the quantum-computing future. Among the group of people the Commission has decided to consult, consensus is nowhere to be found.

Money is also short. France has huge debt problems; the European Commission's new budget cycle is still two years away; Germany has money but has just agreed to build Europe's strongest conventional army.

If the Frontier AI Initiative isn't well-funded, it will be unable to offer the compute or the compensation packages that would compete with US labs, so it's hard to hire world-class talent. Without world-class talent, it's difficult to secure more money. Rumours spread that the Frontier AI Initiative is stuck in a loop familiar to European tech ambitions: institutions are wary to commit until it works, and it cannot work without commitment.

In March, OpenAI raises $122 billion in a single round.

Caroline: Do you know anyone who'd be a good lead for the Frontier initiative?

Christian: on a brussels salary?

Caroline: people seem open to exceptions

Christian: let me ask around

Christian: you know this only works if you scrap all red tape right

Christian: like the uk did

Christian: btw did you see the openai number?

Caroline: Yes.

Christian: that's more than every european ai company combined ever raised

Christian: in a single round

April 2026 - Mythos

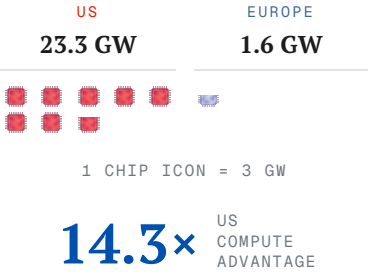

Anthropic announces Claude Mythos, the most capable AI model to date.

Anthropic's internal testing has revealed that the model is so proficient at coding and security research that it has helped identify thousands of previously unknown vulnerabilities across every major operating system and web browser. Some of these are decades old, sitting undetected in codebases reviewed by thousands of developers and security researchers over many years. Mythos found them in weeks.

Almost overnight, Anthropic has become one of the most capable cyber organisations on the planet. It decides Mythos cannot be released until cyber defenders have had a chance to patch remaining holes in their software.

The company launches Project Glasswing: a defensive coalition giving exclusive access to Mythos to a small group of partners – including AWS, Apple, Google, Microsoft, Nvidia, CrowdStrike – to find vulnerabilities before adversaries can exploit them. The aim is to make American software infrastructure secure before open models reach comparable offensive capability, which Anthropic estimates is six to twelve months away.

The UK's AI Security Institute - which has built up state-of-the-art AI expertise in government - is invited for testing, but no European firm or government is provided access.

Christian: you've heard about glasswing i assume

Caroline: Yes.

Christian: there isn't a single european company on it

Caroline: I'm aware.

Christian: the entire continent's software stack is going to be insecure relative to american AI capabilities

Christian: does anyone in brussels actually understand what that means

Many Europeans are sceptical that Mythos is as powerful as claimed. 'They're hyping their products,' scoffs a European Commission agriculture lawyer sitting at the next table from Caroline in the cafeteria. 'Oh no, we've built a nuclear weapon. But don't worry, we can sell you a bomb shelter!' When Caroline objects, the lawyer points out that it's impossible to verify Anthropic's claims if they don't release the model. 'How convenient!'

Caroline can barely hold her frustration. Why would all these large companies publicly associate with Project Glasswing if the whole thing is a scam? How did Mozilla identify twenty times more vulnerabilities this month than normally? She works through the night to write a memo on the strategic implications of European exclusion from Glasswing, and manages to land an emergency briefing with her director, who proves much more receptive.

Over the next weeks, the EU and Member State governments urge Anthropic to give them access - leaders understand a threat to national security when they see one. Anthropic is cautious but open to talk. But when the company eventually announces it wants to expand Mythos access to seventy additional organisations - many of which are European - the White House objects. It is not secure, they say. Besides, Anthropic's compute is finite; more partners means less capacity for US government workloads.

A month earlier, the US administration had wanted to wipe Anthropic off the face of the earth. Now it wants the company's most powerful model kept largely for itself. It finds a workaround to let agencies access Mythos through Glasswing while the supply chain designation is technically still in force. Suddenly, the White House decides it can get along with Anthropic after all. The president still calls Anthropic 'woke', but now also 'very smart'. * *

June 2026 - Disillusionment

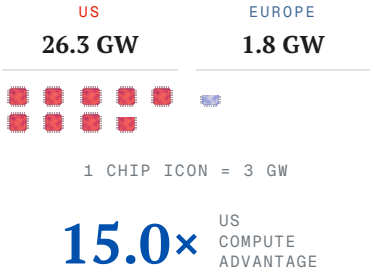

After eighteen months of waiting, DeepSeek has released its next model, V4. It is impressive given the company's limited resources but still at least six months behind the American frontier. DeepSeek admits that it lacks the compute to serve the model widely. The company that briefly persuaded Europe that compute did not matter remains as compute-constrained as ever.

Europe is struggling even more than China. Mistral has fallen further behind. It has managed to raise another €830 million, in a round 150 times smaller than OpenAI's most recent one. Desperate for more capital, Mistral has begun talking to American investors. A rumour breaks that SpaceX, Elon Musk's rocket and AI conglomerate, is exploring an acquisition.

Christian: mistral really might sell to spacex

Caroline: The French government would never let that happen.

Christian: they need access to vast amounts of compute and capital though

Christian: i think they either sell or continue their consulting pivot

Christian: either way, it's the end of european frontier ai

Christian: like that's it. there isn't another one

The infrastructure story is just as grim. The largest AI supercomputer in the US runs at 1,250 megawatts. The largest in Europe runs at eighty-three. The Gigafactories plan is far behind schedule, with expected operation pushed back to 2029. Budget constraints mean that it will be scaled down in ambition. Even if a Gigafactory were to drop magically out of the sky tomorrow, fully built, it would be around a tenth the size of America's largest AI data centre. Europe hosts only five per cent of the world's total AI compute. The United States hosts eighty.

Sure, there are bits of good news too: Softbank, the same company that stood on stage with President Trump, has promised to invest $45 billion over the next five years to build AI data centre capacity in France. But Caroline has learned to take such promises with a grain of salt. Fluidstack, a cloud provider that planned to build a gigawatt-scale data centre near Paris, recently dropped the idea and, to add insult to injury, moved its headquarters from London to the US. OpenAI similarly pulled back on its own plans to build a large data centre in the UK, citing regulatory hurdles.

When, in June, the EU announces its long-awaited tech sovereignty package, Caroline feels torn about the fanfare. On the one hand, the diagnosis is correct: Europe really has become too dependent on American tech. One of her suggested solutions has made the package too, in the form of a proposal for designated zones where industry can build AI data centres at an accelerated pace. On the other hand, the numbers do not add up. The EU's stated goal is to attract €200 billion in private capital for AI data centres by 2036 - that's a quarter of what the American hyperscalers are spending in 2026 alone. When she explains the situation to a friend, he returns: 'So you're proposing a decade-long programme the size of their quarterly capex and are calling it a historic mobilisation?'

After extensive negotiations, a select group of Europeans has by now been promised access to Mythos. But while legal departments are still fleshing out contract details, Anthropic is already training Mythos 2.0. Worse still, the US president just signed an Executive Order that asks frontier AI companies to 'voluntarily' allow the US government to screen their models for cyber risks, and which lets government agencies access those models before any public release. The details of this national security review are scant, but the structure is plain enough. Any model powerful enough to matter – a 'covered frontier model,' in the EO's language – now passes through Washington first: the government gets up to thirty days of exclusive federal access before release, and a say in which 'trusted partners' can use the model before a full public release.

Caroline reads it twice. The thirty-day window is survivable, but the partner-selection power means that whoever gets the model while it is still ahead of the open frontier is now a decision made in Washington, through a classified process, with no particular reason to put a European name on the list.

Christian: did you see the executive order?

Christian: this is glasswing but for every next-gen model from now on

Christian: there's no clause about allies. not one

Caroline: I know.

Christian: so whether europe gets early access to frontier models

Christian: is now a discretionary call by the NSA

Christian: every three months, forever

Meanwhile, Anthropic and OpenAI continue to see exponential revenue growth; the former is on course to reach $100 billion annualised by year-end, and has signed emergency compute deals to serve its rapidly increasing customer base. At the AI labs, internal workflows now change every month or so, as the models enable new ways of working. Researchers manage swarms of agents that handle the bulk of their coding. AI is contributing to its own R&D, making itself better, and using those gains to get better at getting better. A few months earlier, the US Federal Reserve released research showing employment levels among junior software engineers are significantly below expectations. Software company stocks are plunging, as investors assume powerful coding AI models will eat their lunch.

Caroline has seen much of this coming for eighteen months, but has not been able to make any meaningful difference.

The Scenario

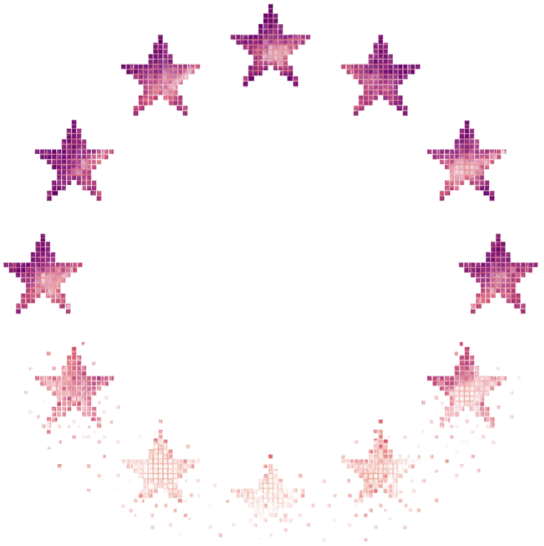

Up to this point, everything we've said has happened – with only Caroline's and Christian's personal stories representing fictional elements. From here on out, we start speculating. We no longer single out individual AI companies, and instead refer to made-up actors: Atlas for the leading American AI company, Helios for the leading European company, and Zimo for the leading Chinese one.

August 2026 - The fork

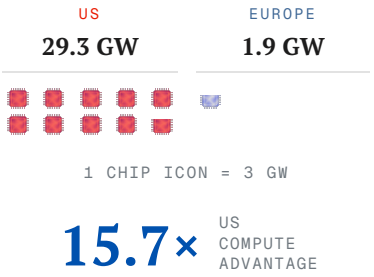

Back in February, the German chancellor made a state visit to China. In Hangzhou he toured Chinese EV and robotics factories. Insiders reported that he came back shaken, convinced that German manufacturers have no shot against their Chinese competitors unless they radically change course. Now he schedules a similar trip to San Francisco. Over the last months, people he respects - German business leaders and economists, American bankers - have told him the AI revolution is for real. He wants to see it with his own eyes. An industry delegation joins him.

The chancellor is not the kind of man who is easily impressed, but again he returns from his trip quieter than he left. He has come to realise that European elites have misread the past eighteen months. While they debated whether AI would hit a wall, the technology progressed faster than even the bulls predicted. While they weighed AI's importance, it revolutionised software engineering and cybersecurity. And while they hailed their sovereign AI initiatives, their dependence on American providers only deepened.

Over the following days, he holds a series of long phone calls with the French president and the president of the European Commission. All three feel Europe has reached an inflection point. AI will keep getting more capable. There is no good reason to believe it will stop at human level. It will rearrange labour markets, security architectures, and the balance of power between continents. Europe is dependent and unprepared. If the ship is to be turned, it must be turned now.

The question is how. Their technical advisers plead for a full regulatory carte blanche for data centre providers and other companies in the hardware supply chain. Only an extreme response is proportional to the stakes, they say. AI will spur a new industrial revolution, and if Europe doesn't rapidly industrialise, it will get left behind. Europe needs to move faster than anything it has ever attempted in peacetime.

Their political advisers beg for something more restrained, warning that their governments might not survive the public backlash against such a package. The people don't like AI; France, Germany and the EU face many other important challenges, and only a small fraction of voters will understand the necessity. The 'effective response,' they say, may be most effective at ending political careers.

September 2026 - A positive vision

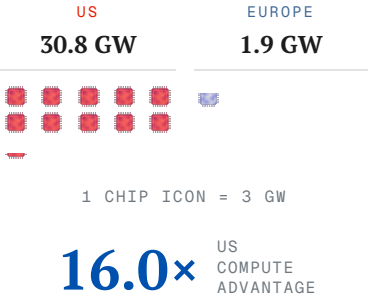

At the Franco-German AI Sovereignty Summit in Strasbourg, the French president and German chancellor give a joint address. It is written by their political advisers.

Caroline watches at home in Lille. The two leaders speak of resolve and of destiny. The US can no longer be relied on in tech, they say, any more than it can be in defense or energy. So Europe must build its own frontier AI.

The challenge is huge but surmountable, they say. What is required is investment, rules that force American providers to play fair, and a public that buys European. The speech has excellent slogans.

Over the following weeks, the AI announcements come almost daily. The plan is to overwhelm the sceptics and imbue a sense of hope. The Frontier AI Initiative gets properly funded with a €2 billion capital injection. Four more Gigafactories are announced. New AI adoption programmes are launched. The European Commission charges an American provider of general-purpose AI models for failing to comply with the AI Act. It also opens two systemic-risk proceedings under the Digital Services Act over their handling of AI-generated disinformation, invoking the law's most open-ended provisions.

The headline measure, the one Caroline's colleagues talk about over coffee, is a Commission proposal – endorsed by France and Germany – for a Digital Sovereignty Regulation requiring critical public sector workloads to run on 100% European cloud and AI software by 2032. No more American AI; no American clouds. It is modelled on the European climate targets, and would create the kind of binding deadline that focuses minds. It would also establish a large, guaranteed pool of future customers for European providers.

The measures are well-received. Decoupling from the US is popular, and commentators hail the move as industrial policy for the AI age; Europe is at last backing its own technology companies. Even Caroline's more sceptical colleagues concede that something is finally moving.

Some economists and tech policy people express doubts. They say that Europe can no longer avoid ugly trade-offs. The investments are too small and don't create enough leverage – where are the market incentives to build compute at scale on European soil? The procurement targets, they

say, do not address fundamental issues, such as the fragmented single market, inflexible labour laws that harm startups and AI adoption, or the national rules that functionally close off sectors like healthcare and legal services. And while regulatory actions under the AI Act are well-justified, those under the Digital Services Act appear at least in part politically motivated. Shouldn't Europe pick its battles more deliberately?

But European elites are tired of negativity, and this is a positive vision, a rally-round-the-flag moment. If policymakers have doubts, most of them are left unsaid.

Caroline certainly has doubts. She is worried that the whole AI sovereignty package assumes that, in nine years, Europe will still have a frontier AI sector worth protecting. What happens if it doesn't?

She writes a memo making that point, calling for measures that build leverage - a contingency plan in case European AI champion Helios cannot close the gap, or the Gigafactories fall short, or the buy-European mandates leave civil services with worse tools than the American ones. Her director reads the note carefully. It is, he tells her, a thoughtful contribution. He promises to pass it up the chain.

Christian: 2030? 100% sovereign?

Christian: what's the plan b?

Caroline: There isn't one.

Christian: of course there isn't

In private, the German chancellor and French president tell each other they did what was possible. Even they cannot smash through so many political realities all at once. The system would not have absorbed more.

June 2027 - Closing windows

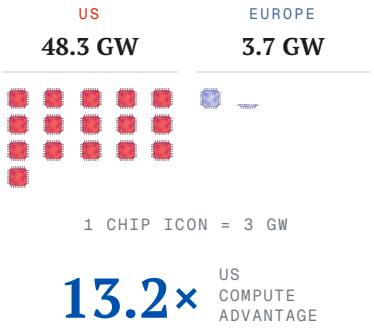

Over the next summer, the Chinese lab Zimo releases the weights of a Mythos-class AI model. The offensive cyber capabilities that Anthropic had locked behind Project Glasswing are now available to anyone. Some of those people are not entirely well-intentioned.

The new capabilities let hackers cut through old, ill-defended commercial software like a chainsaw through tissue paper. European universities, hospitals, regional governments, any institution which hasn't paid for Mythos-class cyberdefence, find themselves locked out of their own systems, with cryptocurrency wallet addresses on the screens and no remaining option except to pay. Only those who use AI for cyberdefence have a chance, which makes people furious at the labs: they unleashed the disease and now sell the cure, often at high rates.

Europe is learning the effect of its sovereignty attempt in the most direct way possible. The Digital Sovereignty Regulation has just passed and several Member States have begun adjusting procurement well in advance. The American AI leader Atlas is miles ahead of Europe's AI champion Helios in cyberdefence. Those agencies most committed to the 'Buy European' agenda – organisations which have procured exclusively from European providers – are now the ones paying the ransoms.

Organisations that kept an American contract on the side are doing better. But even they are running second-tier defences. Ever since the Executive Order, the US government informally greenlights the release of each new frontier model. The result is that the best American cyber models reach Europe two to six months after their domestic release, dangerously close to the open-source frontier, and well after American hackers gain access. Officially, this is about safety and oversight. Unofficially, Washington has discovered an asymmetric advantage it has no interest in surrendering.

Caroline spends most of June on calls with national governments. The conversations are short and unpleasant.

Christian: how's brussels handling the ransomware wave

Caroline: Badly. The sovereignty policies are biting.

Christian: the irony

Caroline: It's not funny from here.

Christian: sorry. you're right

Just as the situation is becoming untenable in Europe, both the US and China announce stringent restrictions on open-sourcing frontier AI models. Washington cites national security, pointing to adversary R&D acceleration and the proliferation of autonomous cyber capability. Beijing speaks of social stability and orderliness. The two governments have, for once, arrived at the same conclusion: frontier model weights have become too dangerous to give away to anyone.

Walking out from the office on the evening of the Chinese announcement, Caroline notices a positive mood around her. The EU's Tech Sovereignty Package had hailed open source for its counterweight to American AI dominance. And yet her colleagues are relieved. The ransomware wave will slow. Offence will plateau; defence will catch up. The crisis, her colleagues agree, has been contained. Few people seem to be thinking anymore about what the open source ban means for Europe's growing dependencies.

January 2028 - Taking stock

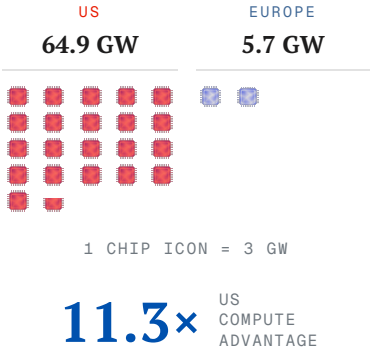

Sixteen months after Strasbourg, Caroline has been promoted and moved to a different team.

AI progress is even faster than in 2026, just as the tech CEOs predicted. Agents now run spreadsheets, design software, operate financial tools. One model generates an image; another opens Photoshop and clicks through fifty iterations until the composition is right. Caroline sits through a demo of an agent navigating enterprise software at twice human speed without a single misstep, and finds it unsettling to watch. The leading US labs are well on their way to automating AI research itself; Atlas has gone public, its market capitalisation hits four trillion euros, and by now everyone agrees AI is the new big thing.

The race has also thinned out - a result of exploding capital requirements and AI R&D flywheels. The American AI companies in fourth and fifth position have fallen further behind. One is considering folding its frontier research into a partnership with the leading lab; another has spun up a parallel cloud service and is attempting to manufacture AI chips.

Helios, riding the European sovereignty wave, has impressively held its position roughly one-and-a-half years behind the American frontier. It has raised billions from new European investors, and has kept talent from defecting.

The relative gap is unchanged, which makes it a kind of victory. But the absolute gap is a different story. In 2023, eighteen months was one generation. Now that capability progress has accelerated and iteration cycles have shortened, it is many. And AI systems are no longer mere chatbots: they are running drug discovery campaigns, are automating mathematics and cybersecurity, and performing large chunks of many white-collar workers' jobs.

An old friend of Caroline's, a software engineer at a French retail bank, tells her over a glass of wine that his employer has a sovereign AI policy. He is, partially, permitted to use Atlas. He is not permitted to upload anything substantive – not even to name the columns in a dataset. He could use Helios, but it is worse. So in practice, he runs Atlas on his personal laptop, and copies the files across.

The pattern repeats inside most large European companies. Everyone knows this is happening, but no one says it out loud.

And yet, some Brussels politicians have retained hope.

The Frontier AI Initiative is running at full steam. France and Germany have thrown a lot of political will behind the non-profit. They have decided to offer competitive compensation, and have managed to secure serious talent, some of which has experience at American labs. One of the Initiative's research bets – world models that can be used to train robots – is producing genuinely promising work that is getting international attention.

The 2028 EU budget has unlocked substantial new funding for AI in science: medicine, materials, clean energy, areas where Europe could still plausibly win. Adoption pilots, after a rocky start, are returning positive results. Doctors are more productive. Teachers are reporting better outcomes for their students.

Warnings of employment catastrophe are yet to materialise as well. Even in sectors with high adoption, AI is not significantly reducing jobs. And workers who can orchestrate AI systems – consultants, lawyers, software engineers, analysts, designers – are seeing their output and wages rise. Judgement, client relationships, and accountability become more important as AI takes over the gruntwork.

A European PM concedes in an interview that Europe was slow to start, but insists it is now catching up, and that European AI is spurring a new productivity boom. Caroline shares it with Christian when she gets home.

Caroline: Maybe my own boss isn't so bad after all.

Christian: enjoy the productivity boom

Christian: it's a lagging indicator

May 2028 - The compute crunch

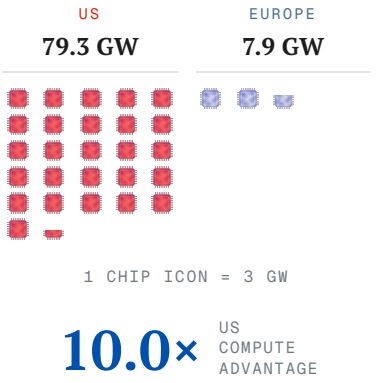

The world is now screaming for more compute.

Workarounds have been found for every previous bottleneck in the AI chip supply chain. When fabs were in short supply, TSMC and Samsung built more. When US grids could not handle additional data centres, the labs bought mobile gas engines and generated their own power on site. When high-bandwidth memory supplies got tight, capacity was diverted from consumer products – with the effect that smartphone prices spiked, but that was very much someone else's problem. But this bottleneck is different.

ASML, the Dutch semiconductor giant, is the only company that can manufacture the EUV machines TSMC uses to fab AI chips. Output has risen from sixty to eighty-five machines a year – a real improvement, but wholly insufficient against skyrocketing AI demand. ASML has more than 5,000 suppliers upstream. If any of them misses targets, the whole chain slows. The engineering is so complex that improving a single one of the machines's 100,000 parts requires a PhD and years of specialised experience. You cannot conjure such people out of thin air.

Washington does not love the fact that ASML still exports some of its older immersion lithography machines - called 'DUVi' by the industry - to select Chinese companies, which use them to create China's somewhat less advanced AI chips.

After briefly seeing eye-to-eye on open source, US-China relations have again soured over the last year. Washington is worried that China will outbuild it once it completes its domestic semiconductor industry. China is adding immense amounts of energy at a staggering pace, whereas American electricity production has flatlined in comparison. If AI becomes a game of building the most data centres, China will win in the long run. Beijing, for its part, worries that Washington will achieve an irrecoverable military AI lead before China has matured its semiconductor ecosystem. If so, the US could use its AI advantage to extract geopolitical concessions. Those worried in Beijing are not eased by some American AI CEOs talking openly about using military AI to 'democratise' autocratic regimes.

Washington has come to realise that it has a short window to act. Its lead in AI is largely the result of clever export controls instated years ago, on Dutch EUV machines and American AI chips. Now, it doubles down. It piles

pressure on The Hague to halt ASML's remaining exports and servicing of its DUVi machines. It is a big escalation, since China uses the same machines to make domestic chips for everyday goods like smartphones and laptops.

The Dutch resist and seek support from other EU Member States. They understand Washington's reasoning but do not want to be bossed around, let alone get dragged into a great-power conflict by an administration that has spent the past two years bullying Europe.

The Commission backs the Dutch, but several Member States are terrified of American retaliation if Europe keeps exporting the machines. The European position fragments before it has properly formed.

Feeling betrayed, the Dutch turn to Japan and South Korea - countries with similarly important positions in the hardware supply chain that also face American pressure - to coordinate a joint response. But both capitals remain heavily dependent on American military support and have, for good measure, been visited by senior American officials in the previous fortnight. They are politely unavailable.

When the Dutch hold firm, Washington threatens them. It can invoke the Foreign Direct Product Rule, a rule that lets Washington assert jurisdiction over any product made anywhere in the world, by anyone, so long as it was produced using American technology or software. It is the same instrument Washington used to strangle Huawei in 2020. ASML's machines qualify several times over. Continue to sell any DUVi to China, and ASML would be in violation of American law. The company could face crippling penalties – denial of export privileges, civil fines, and, in theory, criminal liability for individual executives. Such penalties would also harm the US, which is dependent on ASML's machines. But ASML cannot afford to call Washington's bluff. The Dutch cave.

In Brussels, Caroline reads the news on her phone between meetings. She thinks of her 2026 memo about the need to build leverage. She thinks of her director's friendly face when he told her it was a thoughtful contribution.

Christian: this is the warning shot

Christian: tell me they're treating it like one

Caroline: My direct colleagues are.

Caroline: Others are calling it a setback.

Christian: ok

Christian: noted

Caroline: I mean, I get why the US did it.

Caroline: But we didn't even manage to negotiate a single thing in return.

Christian: it's fucked up

August 2028 - Writing on the wall

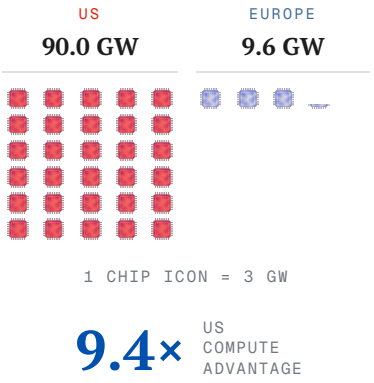

AI models no longer think in English.

Instead of writing thoughts onto a digital scratchpad – the system used since early 2025, and which humans can read – the new systems cycle through long lists of numbers, called high-dimensional vectors, which nobody, not even other AI models, can properly interpret. Freed from the need to translate their complex thoughts into English, the systems think faster and more deeply. It creates a rapid jump in intelligence and capability.

Safety researchers who follow the work are alarmed – many of their strategies for controlling models relied directly on these scratchpads being available. How will we know the models aren't secretly pursuing their own goals? How will the misuse monitors catch worrying behaviour if they can't read the reasoning? In Brussels, experts want the EU AI Office to force developers back to scratchpad-based systems, or to present other evidence they understand what is happening under the hood. But the AI Office is already locked in heated proceedings against two American developers. The transatlantic relationship cannot take on more.

For most people, the more immediate effects are on capabilities. Models which previously could not reliably finish multi-day research projects now can. In the US, employers who had been holding back from job cuts now pull the trigger. Entry-level hiring further stagnates. Unemployment accelerates.

Europe sees less labour pressure, but also less growth: its economy is growing at 1.6% while the US is posting 3.8%. The gap is undeniable and widely attributed to differences in AI value capture. Europe has access to mostly the same models, but cannot extract the same economic gains. There are three reasons.

The first is ownership. The AI companies and infrastructure providers whose revenues are skyrocketing are all based in the US. European AI-native startups exist, but it is American venture capital that delivers the scale-up rounds, and the fast-growing ones increasingly move overseas.

The second is uptake. Despite Europe's adoption pilots, American firms have moved more aggressively to integrate frontier AI into their workflows. Some European companies are held back by fragmented rules, others by a

risk-averse management culture or by internal policies that force the use of inferior home-grown alternatives. A law firm in Milan that once charged premium rates for its knowledge of Italian commercial law now competes with a US firm whose AI handles Italian, French, and German jurisdictions at once, faster and cheaper. The Italian firm's lawyers are still blocked from using frontier models. The same pattern plays out in consulting, software, marketing, and finance.

The third is what happens to the firms that do adopt. Many mid-sized American companies restructure around AI in months – flatter org charts, smaller teams, faster cycles – while European ones often take years. Works councils slow the deep adoption of powerful AI tools; employment protections make it hard to let go of staff whose jobs can be automated and whose labour force would be needed in parts of the labour market that faces shortages. While some workers become vastly more productive with the help of AI, a small but growing number of European workers engage in pretend-work: they log in, attend meetings, but otherwise let agents do most of the work in a fraction of the time. This is not without benefits: it means extra family time, long lunches, and afternoon walks. The flipside is that firms are paying for both AI and unproductive human workers, and capital that should be funding growth remains stuck maintaining the existing workforce.

This comes at a time at which the European economy is already struggling - its core manufacturing industries are increasingly moving overseas, with its once-lauded automakers among the worst hit. Having missed the boom in electric vehicles, they face immense pressure from cheaper and better Chinese cars, leaving millions of workers with uncertain futures.

And now the European tax base is beginning to erode too. Money that used to flow to labour increasingly flows to American companies and their investors, much of it routed through low-tax jurisdictions European treasuries cannot reach. At the same time, unemployment claims tick up across the continent – not dramatically, but the US numbers suggest where things are heading. And the AI boom is driving up global interest rates, leaving countries like France to spend more of each year's budget paying for its state debt.

Caroline: I think I finally understand you.

Caroline: The only real remedy for Europe's doom spiral is economic growth. Caroline: But the growth is happening in the US.

Christian: right

Caroline: I just bumped into an old classmate of mine.

Caroline: She works 3 hours a day.

Caroline: Lets her agents do the rest.

Caroline: Apparently, she has taken up gardening.

Christian: good for her i guess.

November 2028 - Vox populi

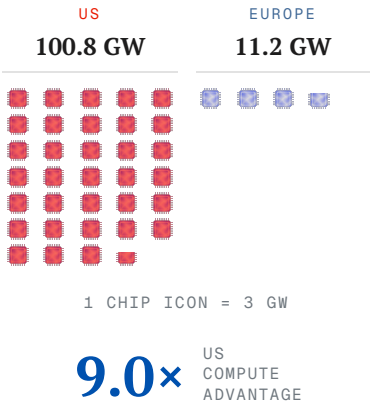

It is the US presidential election. The world, as always, holds its breath.

AI was the campaign's defining issue, but not in the way the industry would have hoped. By now, most Americans want little to do with AI. Climate groups worry about data centres' energy use; unions worry about jobs. Teachers want it out of classrooms; lawyers want it out of courtrooms.

During the primaries, populists on both left and right ran on strong anti-AI platforms. They called for data centre moratoriums, worker protections, bans on AI in education and age restrictions. The general election produced more centrist candidates - strong industry ties helped support the presidential campaigns - but the sentiment did not go away. People are angry.

Anti-AI views span the political spectrum, but it does not bring people together. Fractures have risen, not just along the Democrat-Republican line but between elites who use AI and middle classes who find it frightening, dehumanising, or immoral, and who are watching inequality grow in real-time. People increasingly take to the streets; multiple tech company CEOs have faced assassination attempts. A data centre was recently set on fire. But the US government cannot crack down on AI; the new president, just like the last one, is convinced that the US has to win the AI race from China or face unacceptable threats to national security. So instead of giving the public what it wants, he issues some duct-tape solutions to try and calm it.

Anti-AI sentiment spreads to Europe, too. People are furious at US tech companies and want governments to take action. They call for stronger safety nets, not realising that Europe can barely afford the ones it already has.

Meanwhile, Europe's AI companies are falling further behind.

The gap between Helios and Atlas has widened further, despite the public investments, the compute subsidies, the preferential procurement contracts. American labs, powered by swarms of internal AI agents writing most of their own code, are making algorithmic progress at more than twice their human-only speed. Only compute constraints are preventing even faster progress, and Atlas has more compute than anyone in history.

Helios' research multiplier, with a fraction of the processing power and without access to the best US models for internal use, is barely registering. Every month, the gap grows wider.

Public efforts are delivering public goods, but aren't meaningfully boosting European sovereignty. The Frontier AI Initiative has made impressive progress on its interpretability agenda, which benefits AI reliability globally, but its world-model programme has been hijacked. As soon as results started picking up, Atlas took notice of the Initiative's progress, rapidly spun up its own team, and managed to poach four of the Frontier AI Initiative's best researchers in exchange for astronomical compute budgets. The researchers wanted to stay, they tell colleagues on the way out; they believed in the European project and wanted to help. But at some point, belief is not enough.

The public can see that the EU's grand AI strategy is failing. Investments have not allowed champions such as Helios to catch up. Regulatory actions under the DSA have not levelled the playing field; their most visible contribution has been to annoy the Americans and worsen EU-US relations.

But Europe has invested vast political capital – and actual capital – in the project, and admitting failure would mean admitting that it has spent two years and tens of billions of euros on a dead end. So it doubles down.

A new €20 billion European AI Sovereignty Fund is announced by the European Commission, focused on photonics, edge AI, and other 'next-wave' paradigms. It is clearly a long shot. The same institutions that failed to turn money into frontier capability are being asked to try again, with a harder target and less time. It is, a Polish MEP tells Caroline at a reception, the first known case of raising your bet with no cards in your hand.

Cracks appear. In Germany, a populist party running on an explicitly anti-AI platform – *Stop the machines, save German jobs* – is polling first in the upcoming federal election. In Italy, where Eurosceptic governments have been the norm for half a decade, populist parties openly campaign on a referendum on EU membership.

In Paris, patience runs out. The Élysée no longer believes the EU's plan will deliver on the timescales required. Helios, the only real remaining European player in LLMs, may be months away from being permanently

outcompeted. After intensive negotiations, France announces a €15 billion state investment in Helios for a 17% equity stake, a board seat, and veto power over future fundraising rounds.

Christian: france just bought a corpse

Caroline: That's not fair.

Caroline: These people are actually trying to do something.

Christian: doesn't change the facts.

January 2029 - Dreams of a robot economy

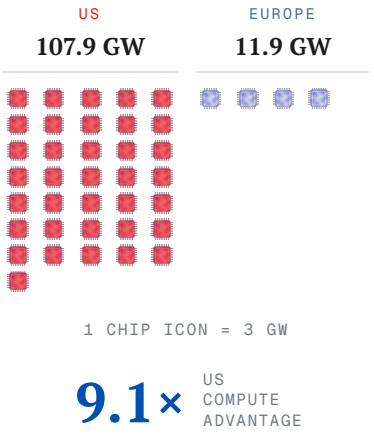

China's AI strategy is not unlike the EU's, except in one crucial aspect: it appears to be working. Like Brussels, Beijing is backing its most promising companies, protecting the industry with subsidies and procurement mandates, and pushing for rapid adoption. But it is easier for a centralised, authoritarian state to force its own agenda than it is for twenty-seven disputatious Member States in a liberal democracy. When its compute was fragmented between different labs, the Chinese Communist Party simply ordered those labs to pool resources; the European Commission has no such powers. China's talent pool is deeper and it has access to cheap, abundant energy. Chinese frontier labs remain within a year behind the US.

Importantly, China does not believe its strategy relies on being at the bleeding-edge frontier of cognitive AI. The government has always wanted AI to boost the physical economy first, and holds a huge advantage in robotics manufacture. Massive state subsidies have pushed annual humanoid production above 1 million units; a household robot can now be bought for €10,000. In cities like Shenzhen, it is commonplace to see humanoids cleaning streets, or packages being delivered by quadrupeds. Beijing is betting that it only needs to stay close on Silicon Valley's heels in AI for its industrial advantage to start to pay off as more robotics-relevant capabilities emerge.

The bet is still looking viable, and in the US, concern is mounting. American politicians have started to refer to the AI race with China as the second Cold War. Tensions have risen further since the US forced ASML to halt exports to China, and Beijing and Washington have both since beefed up the security around their AI development. In the US, researchers are screened by the NSA and need security clearance to access frontier model weights; the federal government regularly accuses China of stealing algorithmic secrets. China has semi-nationalised Zimo and is throwing its considerable weight behind its AI efforts. When part of the Chinese power grid goes down near a major Zimo data centre, rumours spread that it was the result of a US-backed, AI-powered cyberattack.

Christian: im in shenzhen.

Christian: a robot wearing a sailor hat just made me a negroni

Christian: they're making robots like paperclips here

Caroline: Was the negroni any good?

Christian: yeah. excellent

April 2029 - Access all areas

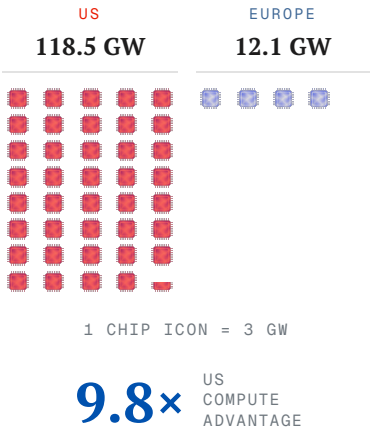

AI demand is soaring, and compute supply cannot keep up. The labs are rationing access to their frontier models and raising prices; businesses are screaming for tokens.

American firms operate 70 per cent of the world's AI compute, and sell their services around the globe. This means that American infrastructure is being used to make foreign businesses more productive, foreign militaries more capable, and foreign labs more competitive – although that last one is via illegal distillation attacks, not legitimate sales. Washington is getting more concerned about it by the day.

The American national security review has been formalised and no longer pretends to be voluntary. Access to the most capable models is throttled by default, in part to save compute for American customers. Domestic government agencies get them first; allied governments next; adversary nations not at all.

But access to the models and their weights is one thing. The compute shortage means that inference is also at a premium, and by April, Washington has had enough. It begins not only limiting access but also rationing use, even for those countries inside the tent.

The *Frontier Inference Services Rule* (FISR) is a country-based licensing regime. Tier 1 countries – close allies, such as the Anglophone 'Five Eyes' intelligence-sharing nations, as well as Japan, South Korea, Taiwan and the Netherlands – get unrestricted commercial access and light-touch reporting. Tier 3, hostile countries such as Iran, Russia, and China, are by default denied any access at all. Most of Europe is in the 100 or so Tier 2 countries in between. FISR dictates that no more than 25 per cent of any provider's frontier inference can go to Tier 2 customers in aggregate, with individual licenses reviewed against a list of factors including 'alignment with United States national security interests.'

For Brussels, the 25 per cent is the number that matters, and it is bad news. European customers currently account for nearly a quarter of US frontier inference on their own. Fairly sharing it with eighty or so other countries would mean roughly halving European allocations. Within a week, European business customers without long-term contracts get notices

from their providers, cutting volumes and raising prices. They cannot change suppliers to another lab, because every US provider is under the same restrictions.

Europe depends on American AI, but the US has no such dependency on Europe. Compute is now so scarce that losing European customers has no discernible effect on the labs' bottom lines: they were already capacity-constrained, so the inference they are no longer selling to Europe can be redirected to latent American demand or simply their own internal R&D. Atlas and its competitors privately urge the White House to soften its stance - the restrictions risk closing off future markets and are horrible PR. But they are unwilling to express their concerns publicly, because as national security restrictions around AI increase, the more important good relations with the federal government become. Besides, they are focused on winning the capabilities race and keeping control of AI systems that are rapidly becoming cleverer than the people running them. There is little pressure on Washington to shift its stance.

In Brussels, an emergency European Council meeting is convened. French and German ministers head to Washington to demand Tier 1 designation. They are told Tier 2 better reflects 'the current state of the bilateral relationship.'

Caroline reads the readout on the metro home. She had always known this moment would come. It still lands like a dropped weight.

Caroline: They understand now. It's just too late.

Christian: yeah

Christian: i was afraid of that

Christian: how long before they threaten to cut all access

Caroline: I was thinking the same.

May 2029 - Tilt

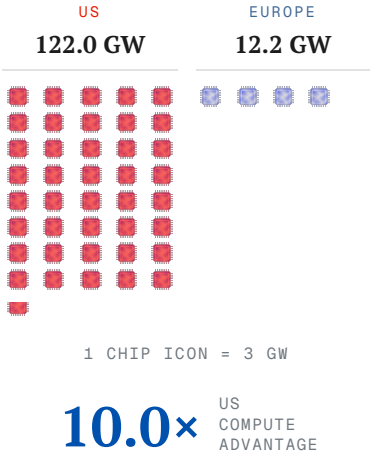

A week on from the imposition of the inference cap, the phones in European capitals are ringing as panicking businesses try to put off the inevitable. A French utility CEO calls the Élysée to say his cybersecurity teams are already losing ground to AI-enabled attacks, and that halving access to American frontier models would leave critical infrastructure dangerously exposed.

The head of a Danish logistics giant tells the country's prime minister that years of optimisation have left the company dependent on US systems it now cannot easily replace, and that their entire business model is at risk. A German delegation of smaller businesses warns the Chancellery that price rises alone will prevent thousands of small firms from using frontier AI at all.

European companies that were smart enough to postpone or walk back their sovereign AI policies spent years building their operations around frontier agents. Now they face having those agents pulled out from under them. The European alternatives are almost two years behind. The Chinese alternatives are not really an option for anyone with a compliance department.

In Brussels, at long last, no one argues that AI is overhyped. Caroline has not heard the term 'bubble' in months. Her director, who used to say she was exaggerating, now spends his days on the phone to national capitals, triaging the damage. He stops by her desk on a Tuesday afternoon with two coffees and hands her one. They drink them without saying much. It is, she realises, as close to an apology as he is able to offer.

Everyone understands the problem, now, but understanding it and being able to fix it are very different things. Demand for Helios's models has skyrocketed far beyond the company's capacity. The Gigafactories are finally under construction but will not come online until next year, and even then will only be a small fraction of what is needed to close the gap.

Christian: you know what the funny thing is

Christian: france is going to make a killing on the Helios investment

Christian: new demand is through the roof

Caroline: Some 3D chess right there.

The European economy is being strangled. After a series of tense phone calls with Washington, European leaders decide that something drastic must be done. For the first time in its history, the Commission opens a formal examination under the Anti-Coercion Instrument – the *trade bazooka*, as the Brussels press has been calling it for years, back when it was a deterrent nobody really expected to fire.

After four months of assessment, the conclusion is that FISR constitutes economic coercion. But the assessment also reaches an uncomfortable finding: the obvious retaliation is self-defeating. Tariffs on US cloud and digital services would raise the price of the frontier inference European businesses are already scrambling to secure. Excluding American providers from public procurement, the threat that would have stung a decade ago, is moot – the Digital Sovereignty Regulation will already do it.

So the proposed countermeasures reach for the instrument's less obvious tools: suspending the intellectual property protections US labs hold inside the single market, and screening – or blocking – American acquisitions of European firms. They are calibrated to hit Washington's exporters without touching Europe's own compute budgets.

The measure that would inflict the most pain is the one pointed at the lithography supply chain: restricting ASML's exports and servicing to American fabs in Arizona. It is also a nuclear measure that will invite a response Europe may not be able to afford.

When it comes to a vote, it falls short of the qualified majority needed. The Netherlands and Ireland oppose, citing transatlantic relations. Poland and the Baltics, worried about Russia, do likewise. Italy abstains. On the record, a senior Commission official tells reporters that the result reflects 'differing national assessments of the strategic environment.' Off the record, a French civil servant tells the same reporters that the delegates are too scared of Washington to use the weapon they had spent a decade building.

Caroline spends most of the summer in crisis meetings. The same people who were telling her a few years ago that everything would be fine are now asking her whether anything can still be done. She tells them that the window for meaningful action closed some time between 2025 and 2027. All that remains now is damage control.

February 2030 - Labour shock

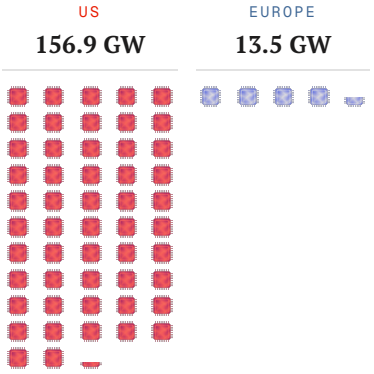

11.6× US COMPUTE ADVANTAGE

Europe's pretend-work honeymoon, of people going gardening while the agents did their work and their companies were barred from firing them, does not last.

European firms, stuck paying a full human staff, cannot compete with their leaner American rivals, especially since they are hamstrung by worse models and less abundant, more expensive inference. The shock hits the most AI-exposed industries first. Software companies lose out because US ones ship faster, at a fraction of the cost. Mid-tier consultancies find their advice is easily anticipated by frontier models, and they have little to add.

What is more, the AI systems just keep getting better. Continual learning was often called the last hurdle to true automation of cognitive labour. Human workers build up context, judgement, and tacit knowledge over the course of a career. AI systems started each new conversation with the same frozen weights. Long context windows and external memory closed part of the gap, but the model never actually *learned* anything new during deployment.

That, however, has changed. If Atlas' latest model spends six weeks inside a consultancy firm, it starts to write the way that firm does. It learns who defers to whom, which clients take bad news badly, which senior staff have reputations worth leaning on. The implementation is imperfect, but the failures are rare and worth the price. Cognitive jobs once thought safe, protected by context-sensitive judgement or institutional, illegible knowledge, are now exposed.

Few people are actively let go because of AI, but the problem is that companies stuck in crisis mode are not creating new jobs. The graduate jobs market is the worst in living memory. Law firms, finance firms, accounting practices – anyone still able to afford frontier rates, or sensible enough to have signed a long-term contract – have paused or shrunk their intake.

Caroline's younger brother finished a master's in logistics management last summer and has been looking for work ever since. They have dinner together in Paris once a month – she pays – and he asks her whether he should retrain. She does not know what to tell him. The nursing shortage is real, but he doesn't like the sight of blood. He is not cut out for the trades. He nods, and orders another beer, and she lets him change the subject.

Politicians who built careers on anti-AI platforms are vindicated. The ones who hedged are scrambling. Street protests become routine in European capitals – some demanding labour protections, some demanding bans on American AI, some demanding bans on all AI, all of them demanding that someone in a building somewhere do something.

European intelligence services have, for over a year, been flagging coordinated information operations aimed at European audiences. The narratives are tuned to local anxieties: American tech is hollowing out the continent, Washington treats Europe as a vassal, the transatlantic relationship is a one-way street. None of it is exactly false. The Hague and Berlin publish attribution reports pointing at Beijing. The public, inclined to agree with the underlying sentiment regardless of who is amplifying it, does not care much.

Christian: how's europe

Caroline: People are angry, understandably.

Christian: it's not great here either

Christian: another data centre got firebombed last week

Caroline: I saw.

Christian: everyone feels like they're losing

Christian: except the labs

June 2030 - An offer you can't refuse

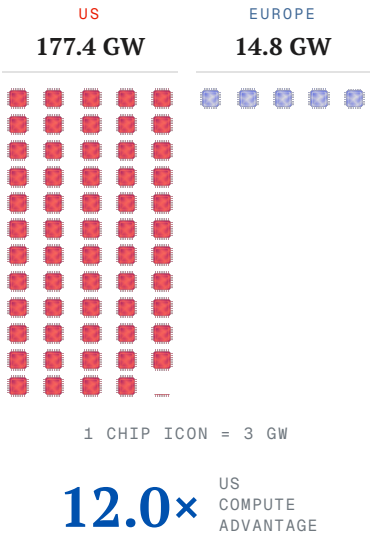

China is leading on robotics. But now Atlas is going all-in on it too.

It announces plans to spend hundreds of billions on industrial data and manufacturing capacity to build robots on a scale comparable to China. Its CEO lays out his reasoning: America still leads in AI software. With help from his new world model team, he has managed to crack the last software challenges standing in the way of general-purpose robotics. But physical supply chains take years to build. If Atlas can get there ahead of its rivals, the advantage will compound: soon, robots will be building the factories that build the robots, just like the AIs are writing the code that improves the AIs.

And there is work to do to catch up to China. While it takes two years to build a new robot factory in the US, in China it takes just seven months. A decade of data-center hyperscaling has left the US energy-constrained, and the American people resentful of everything related to AI; new plants are met with local resistance and state politicians side with their voters. The economy's undeniably rapid growth is constrained by infrastructure, inequality and public opinion.

So the CEO decides the fastest way forward is not to build but to buy – to retrofit factories that already exist. Backed by friendly investment funds, he goes shopping for industrial companies with usable floor space that can be turned over to producing wheeled robots, quadrupeds, and humanoids.

Europe's car manufacturers are high on the list. After years of Chinese competition, Germany's largest is close to bankruptcy. Its market capitalisation is down 80 per cent from its pre-EV peak to €18 billion. To Atlas, now worth $13 trillion, it's small change. But it has a big payoff. German labour law has forced the manufacturer to keep its now unnecessarily large workforce. Their wages can no longer be paid and shareholders have been actively looking for an exit for two years. For a company looking for high-tech manufacturing to build tens of millions of robots a year, it is an ideal opportunity.

Berlin disagrees. It blocks the sale on national security grounds, but insiders know this is more about national pride.

Atlas does not back down. The CEO calls the US president, who is already on record as saying that whoever wins robotics wins the economy, and that China is moving faster than Washington will tolerate. Within seventy-two hours, the White House announces that European automotive imports are hit with sky-high tariffs. Officially, the action has nothing to do with the attempt to buy the carmaker.

Three weeks later, a sale is announced, although it is disguised as a partnership. The German state will take a 20 per cent stake in the new entity, and the existing board will retain nominal control. The press release uses the word 'European' eleven times.

Behind the scenes, Atlas is in charge. It holds the operational majority, the licensing rights to the manufacturing platform, and the right of first refusal on any future capital raise. Profits run through a Delaware holding firm. Berlin saves face, but that is about all.

The pattern repeats several times over the next few months. Boards have a fiduciary duty to act in the interests of shareholders, and when the offer is well above market value and the alternative is insolvency, shareholders' interests are clear. One after another, high-tech manufacturers – robotics, aerospace, specialist toolmakers – are acquired and restructured by American ventures. The official line is that doing so protects jobs and keeps vital facilities on European soil, but in reality, Europe has no plausible counter-offer.

Christian: number eleven

Caroline: I counted too.

Christian: of course you did

August 2030 - Model Collapse

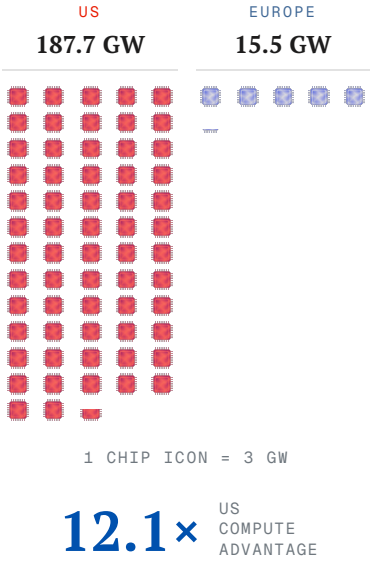

The United States is not trying to destroy Europe. It is trying to beat China in a race it believes is existential. But from where Caroline is sitting, it is hard to tell the difference.

Europe went through a version of this before, after the 2008 financial crash. Welfare obligations rise just as tax receipts fall. Governments borrow against growth that isn't coming. Lenders tighten credit conditions. At each step, policymakers' options get fewer and worse.

In Paris, the finance minister presents a budget to the Assemblée Nationale that nobody in the chamber, including him, believes. There are three numbers which simply cannot be reconciled: welfare spending is approaching COVID-era levels, corporate tax revenue is down nine per cent, and a tenth of the budget is spent on servicing ever-rising debt. The only way to even pretend it can be held together is by making wildly unrealistic growth assumptions that everyone immediately sees through.

Caroline watches it on her phone over lunch. She thinks about her brother. He has not found work. He has student debt and no savings, and has moved back in with their parents. Several of his friends have done the same, and some have started talking about moving to the UK, which has - to the surprise of many - manoeuvered the AI transition much better than most EU countries have.

Moody's puts France on negative watch within the month. S&P follows. The downgrade itself comes in June, but the market priced it long before. The cost of French government borrowing pulls sharply away from Germany's, to the widest gap since the single currency was created. Markets are no longer sure a French euro and a German euro are the same thing, or that the eurozone can hold together.

In August, French debt service crosses 12 per cent of the budget. By October the agencies are reviewing on a rolling basis rather than quarterly. Italy, Spain, and Greece are downgraded in close order. In all four countries, money that used to flow into the exchequer has been diverted to American balance sheets, and the remaining tax base is dwarfed by rising welfare and borrowing costs. At every release, growth assumptions are revised downward.

Around the same time, the loans start to appear. A Chinese state-backed fund extends a credit line to a Portuguese infrastructure bank. A second refinances a tranche of Greek sovereign debt at terms no European institution will match. A third underwrites a Spanish regional government. A leaked Commission memo describes the pattern as 'strategically motivated capital deployment.' Nobody is quite sure what the strategy is. Some analysts think Beijing wants ASML or an EUV licensing deal; others nothing more concrete than for Europe and the United States to drift apart.

Christian: did you see the france numbers

Caroline: What worries me more are the Chinese lifelines.

Caroline: We're getting torn apart.

Christian: i'm sorry

Christian: i really am

The political consequences soon arrive. Protests build throughout the spring until they turn violent. The screens in the Berlaymont show riots in Paris and Rome; young people, mostly, who share few convictions other than that the system has failed them. Populist parties, often explicitly anti-AI and anti-US, lead the polls in most EU countries. Europe is fragmenting. Southern Europe needs help from the north, but even those countries which are not entirely racked by debt and could theoretically help, such as Germany, face their own internal crises.

Bits of the bloc start to drift away. Slovakia has stopped pretending to take the Commission seriously on topics other than trade. Poland and the Baltics, wary of Russia and no longer trusting the EU to save them, deepen engagement with the US. The Nordics, who built their own data centres and thus have something to negotiate with, form their own coalition without Brussels. Many people in the UK conclude that Brexit may have been good for one thing after all: they can now broker bilateral AI deals with Washington more easily.

Across the Atlantic the picture is different. Facing its own anti-AI backlash, the US administration has rolled out job guarantees and direct cash transfers at a scale Europe cannot contemplate. The protests there have not stopped, but the government has bought time. The contrast is not lost on European publics, or on European finance ministers, who understand that the American response is funded by AI-driven growth their own economies cannot offer.

By January 2031 the euro is under sustained pressure. Capital leaves southern Europe and does not come back. The ECB intervenes, intervenes again, and then runs out of credible instruments. By late February, depositors in Italian and Greek banks are moving money north faster than the ECB can offset, and the euro in its current form is no longer something European officials will defend in private.

Christian: how are you holding up

Caroline: I'm okay.

Christian: come to san francisco

Christian: seriously. we'd hire you tomorrow

Caroline: I can't leave.

Christian: why not

Caroline: Someone has to stay.

March 2031 - Between giants

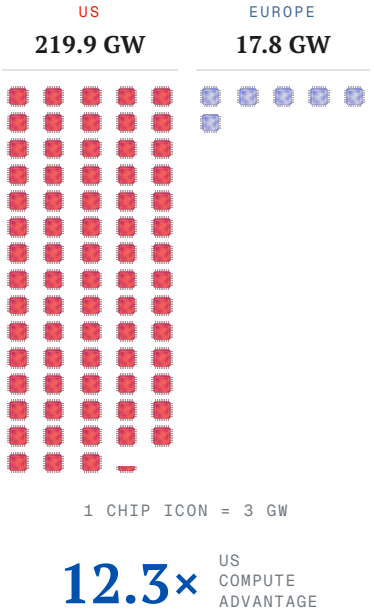

By early 2031 the power is concentrated in two places. American labs hold the cognitive frontier; China still holds the physical one. Atlas alone is worth more than every listed European company combined. The three largest American AI firms each spend more on compute than the EU spends on defence. The leading Chinese humanoid producer ships more units in a month than Europe ships in a year.

The increasingly tense transpacific relationship has brought Taiwan to the fore. TSMC's fabs, where the world's leading-edge chips are made, are mostly on the island. The gap between US and Chinese cognitive AI has grown over the last year, and Taiwan's strategic importance has grown with it. As a result, minor confrontations between navies, which happened perhaps once a quarter two years ago, are now happening every week. Both sides have tested AI-managed weapons platforms in public, and more powerful ones in private. In Washington the labs and the Defense Department are entangled past the point where the distinction means much; in Beijing the integration is more formal. Pundits and analysts regularly use the word *war*, without modifiers such as 'trade' or 'cold'. Three or four men in each capital have the decision-making power that could lead to full-blown conflict.

Europe is in a disastrous state, even if no one in an official position is willing to say so plainly. Polarisation is rampant, and the social model is breaking. In the countries worst affected by the AI shock, economic growth has halted. In others, it has all but decoupled from well-being. Mortgage delinquencies have surged in those member states which tend to use variable-rate loans. European sovereign debt trades at levels usually reserved for countries where people carry cash in wheelbarrows. Where economists once argued about the wealth gap with the US – is it just downstream of hours worked, is the European standard of living better? – now they have stopped; people who go to California can see the difference within minutes of landing.

But Europe has one last card to play. After five years of failing to build a frontier AI sector, it still owns the one bottleneck which the entire race runs through. ASML remains the only company in the world capable of building the EUV lithography equipment that is used to print cutting-edge

chips. Without access to its machines, the US could not keep extending its lead in AI; with access to its machines, China would likely have caught up some time ago.

Beijing has made real progress on domestic DUVi machines, and is expected to begin mass production in a year from now. But that is too little and too late: a year now feels like a lifetime, and DUVi does not produce the cutting-edge chips China needs. The US is pushing ahead every day. The Chinese government is deeply worried that superintelligent AI is on the horizon, and nobody is quite sure what it could do. Some advisers worry the US could use sufficiently advanced AI to neutralise China's nuclear second strike. Others fear such systems could be made persuasive enough to destabilise the Party itself. China's lead in robots is real, but cannot alleviate those concerns.

So Beijing leans harder on the strategy it has been running for two years. The loans to southern Europe grow larger and the terms more generous. The information campaigns intensify. European leaders receive signals about what a closer relationship might eventually look like: privileged market access, robotics co-production, a seat at a table from which Washington has largely excluded them. For the first time, ASML and its EUV technology are mentioned explicitly. Five years of being treated as a vassal by Washington has left its mark, and in several capitals the alternative is, for the first time, being seriously discussed.

China is trying to peel Europe off, and it is succeeding more than Washington is comfortable with. The Pentagon considers losing control of ASML a threat comparable to the proliferation of nuclear weapons. The White House decides it needs direct control of the company now, while the option still exists. The call goes to the three countries that can stop it: The Netherlands, where ASML is headquartered, and Germany and France, without whose acquiescence no Dutch government could agree.

The US vice president delivers the offer in a forty-minute call over a secure line. ASML itself will be folded into a joint Dutch-American holding company, governed by a shared board on which Washington holds a controlling vote over output, customers, and technology transfer. In return, the US offers a capital injection on a scale Europe cannot match, enabling a range of new production sites to be built on US soil. More strikingly, it also promises direct cash transfers to European citizens, indexed to US AI windfalls. These would start small, perhaps €100 per person per year, but would grow over time.

The three European leaders decide not to loop in their fellow EU prime ministers and presidents; they do not consider the offer acceptable, and worry that the others would pressure them to take it. When they politely decline, the White House makes the offer public and adds a stick to the carrot. If the EU does not sign, the whole region will drop to Tier 3 under the Frontier Inference Services Rule and lose all access, present and future, to American AI. The US knows it holds all the cards: through its Taiwanese partners, it has amassed enough EUV machines and maintenance know-how to forego ASML access for longer than Europe can forego access to frontier AI.

Several European capitals reach out to the Chinese, looking for a counter-offer that might give them room to refuse Washington. What they get is not what they were hoping for. Beijing has concluded that its charm offensive is not enough and that they, too, have to switch registers: rare-earth export terms will be reviewed if the Dutch sign with Washington. Robot export licences will be reassessed. The deadline is shorter than the American one.

Europe now has three options, and all of them are bad.

Sign with Washington, and the continent gives up its only piece of leverage and becomes an American protectorate in all but name, while its remaining manufacturing dies if China follows up on its threat and cuts off its exports.

Align with Beijing, and the south might stop drowning, and the Union perhaps survives its immediate fiscal crisis. But Europe will be handing China the keys to the future, and would face a response from Washington which it cannot absorb – not even with Chinese help. Tier 3 would only be the opening move.

Sign with neither, and Europe gets nothing. Instead it would absorb the full displeasure of both great powers at once. It risks losing frontier inference and critical manufacturing inputs at the same time, which would likely cause the already strained Union to disintegrate.

The European Council convenes an emergency session. After fourteen hours, little progress has been made. It authorises a delegation to fly to Washington with a mandate everyone in the room knows is ambiguous on purpose. Against normal Council practice, the principals will decide on the ground.

The meeting takes place in the Eisenhower Executive Office Building on a Tuesday morning. The European delegation is led by the Dutch PM, the French president, the German chancellor, the Polish PM, the Spanish PM, and the Italian council president, each accompanied by their foreign minister and a national security adviser.

The Americans have brought a similar-looking group, plus two officials in the back row who are not introduced and who are wearing earpieces.

Those earpieces are connected to a frontier AI model that has infiltrated every European channel it could find. It knows what each European principal said in cabinet last Tuesday. It knows which of them is having an affair, which is being treated for prostate cancer, which one's daughter is in trouble with the law. It knows what each of them is most afraid of, and what they will trade to avoid it. The Europeans do not know this.

Caroline is in the support team the Commission has brought to Washington. She is sitting in the delegation room, two doors down, watching a closed feed on a wall screen.

By late morning it is clear the principals are not aligned. The German chancellor and Polish PM push hard for the American deal. The Spanish PM wants to align with China; the French president wants to refuse both. The Dutch PM looks ill and the Italian council president has barely spoken in three hours.

At noon, they call a break. The principals scatter into side-rooms with their delegations. Caroline steps out of the delegation room to find coffee and clear her head.

In the corridor, she nearly walks into the council president, who has come out of the meeting room alone. His jacket is off and his tie is loosened. He is shorter than she had imagined. He stops when he sees her lanyard.

'Commission?'

'DG TRADE, sir.'

He looks at her for a moment. 'Walk with me.'

They walk slowly. He is known for consulting unpredictably before major decisions – junior staff, journalists, his driver. Some find it charming. Coincidentally or otherwise, he has also survived forty years in European

politics.

'In the room,' he says, 'everyone is making the case they have been making for two years. I have heard it all.' He looks at her. 'What worries you most, of the three?'

She thinks before she answers.

'Choosing neither looks like keeping our options open,' she says. 'But it isn't. We need to pick one, rather than just letting things happen to us and pretending we're the victims of circumstance. And we can't give the Chinese that much power. So it needs to be the Americans.'

The council president nods slowly. He does not say whether he agrees. He pats her shoulder, lightly, the way an uncle might. 'Thank you. Go get your coffee.' He turns and walks back towards the meeting room.

Caroline goes to the bathroom. She splashes cold water on her face and looks at herself in the mirror. Her hands are shaking. She grips the edge of the sink and waits for it to pass. Through the small high window she can see a slice of Washington sky, flat and bright.

Down the corridor, six people are deciding the fate of the European continent. She does not know whether anything she said will matter. She suspects it will not.

That evening she walks back to the hotel alone. It is cold for Washington in March. She thinks about her brother. She thinks about the dinner party in Hayes Valley, six years ago, and the calm certainty of the people around the table that the world was about to change.

Her phone buzzes.

Christian: you ok

Caroline: Hell of a day.

Christian: my flight got delayed

Christian: i can meet for a drink in an hour

Caroline: I'd like that.

Epilogue

Europe's slide into irrelevance was not inevitable. Even in 2026, the continent could still have changed course, had it shown the courage and political will to take drastic measures.

'Do not go gentle into that good night.
Rage, rage against the dying of the light.'

— Dylan Thomas, Botteghe Oscure, 1953.

June 2034 - Project Inheritance

The domes are visible from the kitchen window. There are four of them, white and ribbed, the closest one maybe two miles off across the scrub. Caroline has been told they are full-scale prototypes for the lunar bases, pressure-tested in conditions meant to approximate Shackleton Crater. On clear nights she can sometimes see a second, smaller dome, lit from the inside, where the company is testing whatever it is they are testing.

She lives in a small house outside a town in New Mexico. The nearest grocery store is twenty minutes away by self-driving car. She just turned thirty-seven. She has been here for a year and a half.

She left the Commission in November 2031, eight months after the ASML negotiation. Brussels had become a grim place. The summer before she resigned had been the one with the Taiwan crisis, when nobody slept for four days because nobody knew whether the carrier groups were going to stop. She had quit without anywhere to go.

Christian had asked her, again, to come to America. She had said no, again. Then, a year later, she had said yes. Her mother had died in the spring and there was little left for her in Belgium or France, apart from her brother. She had not moved to San Francisco. She had grudgingly taken his money and moved here instead; America, but not *that* America.

She and Christian see each other more regularly now. He has sold his company. He is, by any reasonable measure, immensely rich, and has used a portion of the money to fund what he insists is the most important project of his life.

This morning, she is waiting for a call from that project.

It is called Inheritance and its stated purpose is to produce a record, as complete as possible, of how humanity got here. The AIs have already ingested everything public, from emails to parliamentary records to podcasts. What remains, Christian says, is the private stuff.

Christian is doing this because he expects – soon, although the specifics are vague – humanity to start spreading beyond Earth. As the ships move out into the cosmos, communication with home will become slow and sparse, and eventually impossible. Human values will drift over time. The better these space-faring humans understand their own history, he says, the better they will be able to make decisions their ancestors would have approved of.

The first time he explained this to her, Caroline had laughed, but he had been serious.

She agreed to be interviewed. She owed him – literally; he gave her a large amount of money – and besides, her therapist, an actual human therapist, says talking about her years in Brussels would be healthy.

Her laptop chimes. The man who appears on the screen when she clicks the button is in his early forties. He has dark hair and is wearing a soft grey sweater. He is sitting in what appears to be a study. He introduces himself as Daniel.

Caroline looks at him for a moment. 'God, I hate that you're not real.'

'Me too,' he says, smiling. 'Is the appearance alright? We find people are more candid if the interviewer presents in a way they find comfortable.'

'It's fine,' she says.

'Thank you. We can begin whenever you're ready.'

She makes coffee, comes back, and sits down. He waits without doing anything that looks like waiting.

'I'm ready.'

The interviewer nods. 'You worked under three directors at the Commission. Did they understand what was happening?'

'None of them understood in time. My second director understood that AI was going to be economically important, but he thought it would be important in the way that the iPhone was important. He didn't understand that it was going to eat the world. My third director understood that it was going to eat the world, but she thought we had ten years. We had two.'

'And you?'

'I understood it about a year before most of my colleagues did. Which was about three years too late.'

'Take me back to the summer of 2026. Before Strasbourg. What was that period like?'

'Strange. Fluid. The German chancellor had just come back from San Francisco and seemed open to serious measures. So were the French president and my boss. For a couple of weeks it felt like the door was open. I pulled a lot of all-nighters.'

'What did you recommend to your superiors?'

She thinks back. ‘We needed the metal on our soil. The data centres, the chips, the power supplies. Anchored under European law, in jurisdictions Washington could not commandeer at six hours’ notice. Not the five-year-long Gigafactories process everyone was so proud of. Real compute. The numbers had to be in the tens of gigawatts. We should have been building like a country at war.’

The interviewer raises an eyebrow. ‘But Europe hosted only five per cent of global compute at the time.’

‘We could have changed that. Global compute capacity was almost doubling every year. We could have got to fifteen, maybe twenty per cent in five years, if we had wanted it badly enough. That would have been enough to serve most European customers. The threat of losing access to that much compute would have kept Washington from doing what they eventually did.’

‘How would you have done it?’

‘European companies couldn’t have done it on their own. We needed to also partner with the Americans. Money wasn’t the main thing. Subsidies would have helped, but the hyperscalers had money. What they needed was speed. They were in an insane race to ship each new model. Bringing a single data centre online a month faster was worth billions to them.

‘I argued Europe should create Special Compute Zones, reduce permitting timelines from two years to three months, strip out every bureaucratic obstacle. Set up concierge teams liaising between AI companies, energy providers and municipalities. Convert decommissioned power plants, where the grid connections are ready and waiting. Build new electricity generation like you mean it. A stripped-down version of this idea ended up in the EU Tech Sovereignty Package, but we didn’t execute.’

‘Why not?’

‘Because rolling out the red carpet to American hyperscalers seemed like the opposite of sovereignty. The Eurostack people wanted to build their own ideal, but they were ten years too late. I told them that, realistically speaking, we can’t build all the data centres ourselves, so what we need to do is make sure the ones that *are* built are bolted to our floor. But no one wanted to be the politician to say that sentence. It was too humiliating.’

The interviewer nods.

‘You also wrote about a coalition of middle powers.’

‘Yes. We screwed up royally with ASML, but it could have been avoided. The problem was that we tried to coordinate among twenty-seven Member States. A smaller group of middle powers – the Netherlands, Germany, France, the UK, Canada, Japan, South Korea, perhaps assisted by the European Commission – might have pulled it off together. Most of these countries held important pieces of leverage. Supply chain bottlenecks, AI talent, energy.

‘The whole AI supply chain ran through a small number of places, and most of those places were not the United States. We never sat down with the Japanese or the Koreans and said *we are in the same position, you and us, between two empires that only care about us instrumentally, and we have more leverage together than apart.’*

‘It would not have been easy. Every one of these countries had its own complicated relationship with Washington. But there was a middle ground that was in their interest too. We never really tried to find it.’

She pauses. From the next room she can hear the household system humming, readying the laundry. Outside, a quadruped delivery unit is pausing on the gravel path.

‘What else?’

‘There was a real opening in robotics and industrial AI. The Frontier Initiative was surprisingly successful in the world models that made the robots work in the end. But we didn’t set up the industry partnerships, we didn’t unlock the data, we didn’t screen foreign investments. So when Atlas noticed, it simply bought the researchers and the companies.’

The interviewer nods again. ‘What about the job market?’

‘I had a friend in the French government who worked on labour market reforms. He always talked about how Denmark had been running the answer for thirty years: the whole flexicurity thing. You pair serious wage insurance, retraining, and real support for displaced workers with the freedom for companies to let go of staff whose jobs have changed – giving businesses the agility to drive deep AI adoption and stay competitive globally.

‘We knew the model. We could see the Danish numbers. But every Member State had its own labour code and its own unions and its own political coalition, and nobody wanted to spend the capital. So the productivity gap widened and the protection we were supposedly preserving dried up.’

She looks out the window for a moment.

‘I think about June of last year a lot,’ she says, more to herself than to him. ‘The Spanish riots. Watching them on the news, with a feeling they were something I had failed to prevent.’

The interviewer waits. He is very good at that.

‘I know I shouldn’t blame myself, but I do.’

‘Why?’

‘All of what I just described was economic. Compute, supply chains, labour codes. The thing I never really cracked, and which I think mattered much more than I realised, was the story. A picture of why they might *want* this stuff to happen.

‘We had a negative vision. We were good at negative visions. We were going to be devoured. Every memo I wrote started from a description of what we were losing. But we couldn’t offer a positive vision. You cannot ask people to absorb years of disruption on the basis of *otherwise it gets worse*.

‘We should have told them what the good version of the new thing looked like. But we didn’t know what it would look like. We didn’t even have the vocabulary for it. You know, I tried. I wrote one bad memo. But I should have kept writing them.’

The interviewer makes a note on a notepad which does not exist. ‘Do you think this lack of a positive vision was why the measures never passed? You said the political situation was fluid in summer 2026.’

‘It was fluid and it was not. Even the leaders who understood what was coming were not fully convinced themselves. Every one of these measures looked like an ugly compromise; like dealing with people we didn’t like and throwing out rules we had spent decades crafting. They would have involved spending every coin of political capital and much of the actual capital on things voters didn’t understand and didn’t want, for gains that were years in the future. None of the sovereignty benefits felt graspable back then. I’m not sure any politician could have passed the drastic changes we required with their careers intact. I’m not even sure they *should* have done it.’

The interviewer lifts his head. ‘Say more about that.’

‘I don’t know. I genuinely do not know. There is a version of this where the leaders who saw what was coming spent everything they had on it, lost their governments, were replaced by the populists who were already winning, and the populists undid the policies and the outcome was worse.’

‘And yet you wrote those memos.’

‘What I saw was that the European model was about to collapse. So the question was whether you suspended the rules for one sector for a period, in order to preserve the social model elsewhere – or whether you preserved the rules everywhere, and watched the whole thing crumble. I came down on the first answer. Some of my colleagues accused me of betraying Europe’s social-democratic roots. Called me a libertarian. But everything I wrote was an attempt to protect the European project. The traditional political divides had stopped mattering.’

‘Did you ever think: we should hit the brakes?’

‘Of course. Many people did. I would have pressed the button if I could. Slow things down and allow society to adapt, let AI safety research catch up to the new capabilities frontier. But we didn’t have a button. The technology was developed outside of Europe. The US and China were in this frantic race. If we had built actual leverage, we could perhaps have forced them to slow down, or at least get both to pick up the damn phone and talk to each other. But we had no bargaining power whatsoever, so Europe’s only remaining path was through AI, not around it. ’

She pauses.

‘You know, I’ve tried to find people to blame for all of this. But I have come to the conclusion that there were no real villains in the story. It is just that the system that produced our decisions was responding correctly to the incentives it had been given, and the incentives that had been given were hopelessly unfit for what we were facing. The failure was not a failure of individuals but a failure of stuff like information flow, political constraints, and the speed at which institutions can adapt. Nobody wants to hear that, because if it is true then there is nobody to be angry at, and people need to be angry at someone.’

She looks at the interviewer.

‘What do you think?’

He pauses.

‘Can I be very direct with you?’

She frowns.

‘I think you’re full of shit.’ He is smiling.

She stares at him.

‘Excuse me?’

‘Not the analysis. The analysis is fine. But I do not believe the analysis is what you actually feel.’

‘What?’

‘You are telling me a story in which nobody is to blame, in which the incentives were miscalibrated, in which the failure was structural, in which it is not even clear the leaders should have tried. It is a careful story. It is an intellectually defensible story. I do not think it is the story you believe.’

‘You don’t know what I believe.’

‘No. But I have a guess. I think you are angry. I am not sure at whom.’

She does not say anything.

‘I could be wrong,’ the interviewer says. ‘I am not the right tool for many things. But you asked what I think.’

She is quiet for a long moment. The dome through the window is a different colour than it was when the call started.

‘Fine. You want to know who I’m angry at? I am angry at Christian. I am angry at Christian right now, because Christian sent me a fucking AI to interview me about the worst years of my life, and apparently the AI is now also going to do therapy. I am angry that I took his money. I am angry that he keeps showing up in my life with this – this earnest face, this *I have a project that will save humanity* face, and the project is always real and the money is always real and I cannot ever quite say no. Is that what you wanted?’

‘Some of it.’

‘What does that mean?’

‘I think you are redirecting.’

‘Don’t psychoanalyse me.’

He shrugs. There is a pause. She hasn’t noticed that she is gripping the edge of the table. She opens her mouth and closes it again.

‘Fine. I am angry. I am so fucking angry. I am angry at the US and China for almost throwing us into World War III. I am angry that we let a handful of power-hungry men decide our future. I am angry at the AI labs for building these ‘tools’ without knowing how to control them. I am angry at the European leaders for their endless colouring inside the lines, their roundtables, their waiting for the political environment to be ready.

‘The political environment is never ready. It was not ready during COVID either. It was not ready when Russia invaded Ukraine. People did things anyway. When COVID hit, we went into lockdown and rolled out a vaccine program in record time. After the Russian invasion, we built LNG terminals out of thin air and pulled financial support from every Member State to defend the eastern border. We broke every rule because we realised our future depended on breaking them.

‘You know what I blame our leaders for? Not for stupidity, not for malice, but for their lack of courage. Their unwillingness to look at a thing that was hard to grasp but also obviously happening and say *I am going to do something about it that will probably end my career, and I am going to do it*

anyway. Nobody did that. Nobody. They all saw their local incentives and went along with them, and they told themselves stories about why the local incentives were the only thing they could see, and the stories were sophisticated and well-argued and completely beside the point.

'It is like – it is like – I don't know. Some oracle shows up at your door and says you have to qualify for the Olympic 200-metre freestyle in three years or the world ends. And you barely work out. You go to the gym maybe once a week. There is no realistic version of this in which you qualify. But the oracle is right. The world really does end if you don't. So what should you do? You should train. You should train like a god-damn lunatic. You should sleep eight hours a night and eat the right things and stop seeing your friends, because the alternative is the world ending. But you don't. You don't train because the water is cold. Because you are really more of a tennis person. Because you decide the coach has a profit motive. You do nothing. You let it happen.

'That is what we did. We let it happen. We told ourselves it could not be done and we went home and we let it happen. And I am so fucking angry at all of us for that. Because it could have been done. It could have been done. Not certainly. But there was a chance and we did not even try.'

She stops. She finds that she is crying. She has not cried in front of another person for years, not even with her therapist.

The interviewer has, she notices, the decency to look away.

'I know you can still see me,' she says.

'Yes.'

'It's fine.'

She wipes her eyes. Outside, the quadruped is retreating down the gravel. For a brief moment the late sun catches its plating.

It is, she finds, almost beautiful.