

THE POWER TO COME

How artificial intelligence is redrawing who commands — and why democracy is not equipped to respond. Eight arguments.

JULIAN THOMAS FAIRWEATHER · UN FOR THE PEOPLE

The greatest redistribution of power in modern history is happening right now. It does not occur through revolutions or wars. It occurs through code, data, and dependency — so gradually that we barely see it. And when we do see it clearly, it may be too late to reverse.

THE STARTING POINT

1

THE EFFICIENCY IS REAL. THE PROBLEM IS WHO COLLECTS.

Everything begins with a concrete promise: AI makes companies more efficient. This is not propaganda — it is true. A company that replaces a hundred employees with AI systems at equal output lowers costs, accelerates processes, scales without hiring. For a shareholder, it

is a dream. For an economy, it is the beginning of a problem that previous technological revolutions never posed with such clarity.

Previous waves of automation shifted workers toward new sectors: agricultural mechanisation toward industry, industrial automation toward services, basic services toward cognitive work. Each time there was a "next level" to migrate toward. **AI attacks all levels simultaneously.** It does not only replace muscles — it replaces analysis, writing, reasoning, communication. The cognitive jobs toward which earlier generations had moved are precisely those under the greatest pressure today.

For the first time in the history of automation, there is no higher floor to climb to. The refuge has run out.

The productivity gains are real. The problem is not that AI does not work — it is that its fruits concentrate in the hands of capital owners, while the cost of the transition is offloaded onto workers, onto communities that depend on certain sectors, and onto governments that watch payroll-tax revenues evaporate — the base on which healthcare, pensions, and public education rest.

CRITICAL POINT

No market mechanism automatically redistributes productivity gains. That redistribution requires deliberate political choices. And those choices require functioning democracies — precisely the institutions that technological concentration is eroding.

2

THE QUALITATIVE LEAP

**FROM TOOL TO INFRASTRUCTURE.
FROM COMPANY TO FEUDAL LORD.**

The decisive shift is not economic — it is structural. There is a fundamental difference between owning a tool and owning infrastructure. A tool is used or not used. Infrastructure becomes the precondition for everything else: whoever controls the aqueduct controls the city, whoever controls the electrical grid controls industry, whoever controls communications controls public discourse.

AI is becoming that infrastructure for the global cognitive economy. Not one software among many, but the substrate on which business decisions, content production, scientific research, medical diagnostics, legal systems, and financial assessments run. When a technology enters this category — critical, indispensable, without a workable alternative — whoever controls it acquires a kind of power without equivalent in recent history.

Five or six companies control the foundational models on which almost all of this infrastructure runs. **They were not elected. They cannot be removed. They answer to no demos.** They decide which values are embedded in the models, which content is filtered, who gains access to the most powerful tools and at what price, how the systems we use to understand the world are trained.

*Power over thought is deeper than power over production.
Whoever controls the systems we use to understand reality
controls, in the end, the horizon of what we believe it is possible to
ask for.*

Tocqueville in 1835 described "soft despotism": a power that does not oppress brutally but wraps, guides, simplifies — until the exercise of individual will becomes unnecessary. He was writing about state bureaucracy. His description applies today with almost unsettling precision to a certain kind of digital technocracy: it forbids nothing;

it optimises, filters, suggests — until you can no longer imagine doing otherwise.

CRITICAL POINT

AI concentration is not only economic — it is epistemic. A monopoly on steel makes you pay more for steel. A monopoly on cognitive infrastructure shapes **what you know, how you think, and what you believe it is possible to demand**. There is no direct historical equivalent.

THE LINK THAT CONNECTS EVERYTHING

3

THE PRIVATE OLIGARCHY IS NOT NEUTRAL. IT IS AMERICAN.

So far this might seem a market problem: a few companies too powerful, to be regulated as oil and railways once were. But there is a variable that breaks this analogy and makes the problem incomparably harder: these companies are not stateless entities. They are American. And their Chinese equivalents are, directly or indirectly, instruments of the Chinese state.

OpenAI, Anthropic, Google DeepMind, Meta AI operate within the perimeter of American strategic interests. They receive Pentagon contracts. Their founders circulate between the White House, the Senate, and the intelligence agencies. This does not mean they are state puppets — they have enormous operational autonomy — but it means that **the concentration of private AI power and the geopolitical competition between states are not two separate problems. They are the same problem seen from different angles.**

When you regulate OpenAI, you are not only regulating a company. You are touching American strategic infrastructure at a moment when the US is in a technological war with China. This explains why regulation is so slow: every serious proposal runs into the implicit veto of national security.

On the Chinese side the structure mirrors this but is more explicit: Baidu, Alibaba, Tencent, Huawei are not companies operating in the market and paying taxes — they are instruments of the Party's power strategy. Data, models, AI infrastructure are strategic resources of the People's Republic, available to the state when the state demands them.

The result is that the world is developing a global cognitive infrastructure profoundly divided into two incompatible geopolitical blocs — with a third space, the European one, that has neither the scale nor the will to build something of its own.

CRITICAL POINT – THE STRUCTURAL VETO

Any proposal for democratic AI governance runs immediately into two vetoes: the American one ("not now, we're winning the war with China") and the Chinese one ("we're not interested, it is a tool of our power"). **The absence of rules is not an accident. It is the result of powerful interests that prefer the absence of rules.**

THE GLOBAL COMPETITION

4

USA VS. CHINA: A RACE THAT DEVOURS EVERY RULE.

The USA and China are experiencing AI as an arms race in the literal sense. The Pentagon treats it as critical military infrastructure. The Chinese Communist Party has declared AI leadership a national

strategic objective by 2030. When two great powers frame a technology in these terms, the logic that prevails is not "how do we make it safe and distributed" but "how do we win before the other side does."

This logic is devastating for any attempt at international governance. Every proposal for global regulation runs into the objection: "if we self-limit and they don't, we lose." It follows the template of nuclear disarmament — with one crucial difference: nuclear weapons were physically countable and verifiable. AI models are not.

The two models that confront each other are not only technological — they are visions of what AI should do in the world. The American model: private oligarchy with state influence. The Chinese model: state instrument for internal and external control. Neither is designed for citizens.

There is also a structural asymmetry that is rarely named: China has access to quantities of behavioural data that no democracy can legally collect, because it has no privacy constraints. Democracies compete with one hand tied behind their backs. The wrong answer — dissolving privacy protections to keep up — destroys precisely what it claims to protect.

CRITICAL POINT – THE DOMINO EFFECT

Many authoritarian governments worldwide are looking at China and seeing something attractive: technology that reinforces control instead of threatening it. **The USA–China competition does not produce a democratic winner — it produces the global diffusion of both models of concentrated power.**

EUROPE IN THE CORNER: THEATRICAL SOVEREIGNTY, REAL DEPENDENCY.

Europe finds itself squeezed between three power projects — Trump, Putin, Xi — that all share one thing: they see Europe specifically as a problem to be neutralised rather than an interlocutor to be respected. And all three prefer dealing with individual European states rather than with a united Europe, because with Germany you can strike an energy deal, with Hungary you can buy a veto, with Italy you can propose infrastructure — but with united Europe you can do none of these things.

European fragmentation is not a consequence of external pressure. It is the objective of external pressure. Keeping it fragmented is the most effective strategy Europe's rivals can pursue — and they are pursuing it systematically.

The paradox is that national European governments defend sovereignty as though it were still a real thing. It is not. France does not decide its energy policy alone. Germany does not decide its industrial policy alone. Italy does not decide its migration policy alone. The chains of dependency cross the Atlantic, Central Asia, the Gulf. National sovereignty in Europe in 2026 is largely a sovereignty of form: governments have the shape of power — parliaments, elections, ministries — without the substance. The substance lies elsewhere.

The alternative to European sovereignty is not national sovereignty. It is unorganised dependency on powers that owe no accountability to any European citizen.

Europe still possesses something of enormous value: 450 million consumers with high purchasing power, the world's finest universities alongside the USA, and real global normative legitimacy — when

Europe sets a standard, the world listens. This leverage still exists. In ten years, if the technological and demographic gap keeps widening, it may not. **The window is closing silently.**

CRITICAL POINT

Europe is the only entity attempting to build a third model — regulated, rights-oriented. It is slow, fragmented, and faces an enormous industrial gap. But if it fails, there will be no alternative model left to show the world.

THE RESPONSE IN PLAY

6

THE RHINE GROUP: NECESSARY, PARADOXICAL, URGENT.

From 20 to 23 September 2026, Mario Draghi and Patrick Collison — co-chairs, with Luis Garicano as executive director — convened the Rhine Group at Baden-Baden: Nobel laureates (Aghion, Holmström), former ministers (Le Maire), tech founders (Niel of Iliad, Lütke of Shopify), the editors of the Financial Times and The Economist. Three days of discussions under Chatham House Rule, no minutes, legal seat in Geneva. The European Commission "knows very little."

Why did Draghi choose this path? Not out of ideological preference for closed rooms. He had already tried the canonically correct way: his competitiveness report, presented to the Commission in September 2024, was a 400-page document, detailed, data-driven, with concrete proposals. It made news for a week. Then it vanished into the EU's institutional archive — cited, praised, ignored. The very institutions that should have acted on it were too fragmented, too slow, too paralysed by national vetoes to do anything concrete.

The Rhine Group was born from the exhaustion of the **institutional route**. If the Commission cannot move quickly enough, perhaps an informal network of people with sufficient credibility and political capital can force institutions to listen. It is not an elegant choice. It is the choice of someone who has watched the ordinary options fail.

Draghi knows better than anyone how much time Europe has. He already saved the euro in 2012 with three words — "whatever it takes" — operating at the edge of his institutional competence, at a moment when ordinary institutions were paralysed. He is doing the same thing again. But this time the problem is larger and the window narrower.

The comparison with the 1992 Britannia meeting is not reassuring: that closed gathering preceded a season of privatisations that structurally changed the Italian economy without parliaments having any real voice. But the comparison has a limit: in 1992 the urgency was debatable. Today it is not. **If Europe does not build sovereign AI capacity in the next five years, it never will** — technological dependencies will consolidate to the point where any workable alternative becomes impossible.

The paradox remains: we are building a response to the democratic crisis produced by the concentration of technocratic power using precisely the tools of the concentration of technocratic power. Strategic decisions about the future of 450 million people, made in closed rooms by unelected individuals. But between a Europe that saves itself this way and a Europe that waits for a perfectly democratic process while hollowing out, the choice is not difficult.

CRITICAL POINT – EXISTENTIAL URGENCY

The Rhine Group is imperfect, non-transparent, paradoxical. And it is probably **the last real opportunity Europe has to build something of its own before the window closes for good**. Those who criticise it solely on method, without proposing an alternative at the same speed, are choosing procedural purity over strategic survival. That is not a defensible position.

THE FINAL KNOT

7

THE ANSWERS EXIST. THE PROBLEM IS WHO AUTHORISES THEM.

The technical and political solutions to the problem of AI concentration are not mysterious. They have been debated for years, in some cases already tested at small scale. The problem is not a shortage of ideas — it is a shortage of democratic mandate to implement them at the required scale and speed.

Tax AI capital to redistribute productivity gains to displaced workers. **Universal basic income** funded by the windfall profits of automation. **Reduced working hours** as a way to share productivity as collective time rather than private profit. **Massive public investment** in everything AI cannot or should not do — care, education, human connection, ecology. **Binding international governance** with real verification mechanisms, not statements of intent.

None of these measures is technically impossible. All are politically near-impossible in the short term, for a precise reason: existing democratic systems were designed for rates of change far lower than those AI is imposing. The electoral cycle is four years. The AI innovation cycle is six months. Institutions cannot sustain this temporal misalignment.

Every day that passes without deliberate political choices, the trajectory consolidates further toward concentration. The windows to reverse it close silently, as dependencies deepen and alternatives become less workable.

THE REMAINING QUESTION

The real democratic question of the era is not "how do we protect ourselves from AI." It is **who decides who decides**. Until that question has a democratic answer — not technocratic, not market-based, not geopolitical — everything else is managing the emergency. Necessary. But not sufficient.

THE DEMOCRATIC QUESTION

8

DEMOCRACY IS NOT EQUIPPED. BUT IT IS NOT IRREVERSIBLE.

There is an intellectual temptation this pamphlet wishes to resist to the end: that of concluding with catastrophe as the only horizon. The concentration of AI power is real, the geopolitical competition is real, the fragility of democratic institutions is real. But the history of democracies is also the history of broken and rebuilt institutions, of power imbalances that seemed permanent and were not.

The steel and oil monopolies that seemed invincible at the end of the nineteenth century were dismantled — not by the market, but by the political will of societies that had decided that concentration was unacceptable. The same societies that had produced the monopolies also produced the tools to contain them: antitrust laws, regulations, public institutions.

The difference today is speed. And speed requires something that democracies struggle to produce: a collective sense of urgency about a problem that is invisible, untouchable, and arrives gradually.

Democracy is slow by design — to guarantee deliberation, consensus, reversibility. This slowness is a virtue under normal conditions. Right now it is its principal point of vulnerability.

The problem is not solved with less democracy — it is solved with a democracy more capable of moving at the speed of the problems it must face. This has never happened spontaneously. It has always been the result of crises visible enough to mobilise political will, combined with the presence of people who had already prepared the answers.

THREE VOICES, ONE QUESTION

HOW THE FOUNDERS WOULD RESPOND

The same question posed to the three individuals who, more than anyone else, are determining the trajectory of global AI: "When AI makes companies efficient by replacing human beings, the owners of AI will become, in effect, the masters of the world. How do you respond?" What follows is a reconstruction based on their documented public positions.

DARIO AMODEI CEO & CO-FOUNDER, ANTHROPIC

THE QUESTION

I recognise the tension. Anthropic was founded on the premise that powerful AI is probably inevitable, so it is better built by people who

genuinely care about safety. That logic is internally coherent. But it does not resolve the problem you are raising.

The risk of concentration is real. We are a public benefit corporation — better than an ordinary company, but not a democratic mechanism. No external representative body has real authority over Anthropic's strategic choices: how the model is trained, what values it embeds, who gets access. Those decisions are ours to make.

My honest answer is that three things are needed simultaneously: **international regulation with real verification mechanisms** — not statements of intent; **aggressive fiscal redistribution** of AI productivity gains to displaced workers; and a deep revision of who has access to the most powerful tools. But these answers require governments with the will and capacity to impose them. And right now I do not see that.

The problem of economic distribution is not secondary to technical safety. A world in which AI works perfectly from a safety standpoint but destroys the global middle class without redistribution is a disaster regardless. I know Anthropic alone cannot solve this. But we can stop treating it as someone else's problem.

THE UNRESOLVED CONTRADICTION

Amodei openly acknowledges the paradox: he is among the six or seven individuals who more than anyone are determining the trajectory of global AI, and he calls for regulation that would limit him. The logic of "better us than others less responsible" is genuine — and it does not resolve the democratic question.

SAM ALTMAN

CEO, OPENAI

THE QUESTION

I understand the concern, but I think the frame is wrong. The history of technology is the history of abundance democratising itself. A century ago the automobile was a privilege for the few; today it is universal. Electricity,

the telephone, the internet — each time the initial concentration gave way to diffusion. AI will follow the same path.

My vision is that AGI — artificial general intelligence — compresses decades of scientific progress into years. Cures for diseases that seemed unsolvable. Solutions to the climate crisis. Economic growth that reduces global poverty at a speed no redistributive policy has ever achieved. **The problem is not that AI creates too much wealth — it is that existing distribution mechanisms are not equipped to redistribute it.**

That is why I advocate universal basic income funded by AI profits. OpenAI has a capped-profit structure precisely for this reason: profits beyond a certain threshold return to the mission. The stated objective is that AGI benefits all of humanity, not its owners. I believe this genuinely.

On concentration: yes, today five or six companies hold the lead. But models diffuse, the cost of compute falls, open versions proliferate. In ten years the landscape will be far more distributed than it appears today.

THE UNRESOLVED CONTRADICTION

Who decides when and how profits "return to humanity" is OpenAI itself — no external body. The capped-profit structure was modified in 2024 toward a more conventional model. And "the market will distribute" is precisely the promise industrial capitalism made to the workers of the Industrial Revolution — with results that required decades of labour struggles and legislation to be even partially honoured.

ELON MUSK

FOUNDER, XAI & GROK

THE QUESTION

The problem is real — but you are looking in the wrong place. The true danger is not AI concentration as such. It is the concentration of AI in the hands of people with a specific ideological agenda who want to use it to control public discourse and suppress certain points of view. OpenAI — which I co-founded and left precisely for this reason — is the perfect example.

The solution is open source. If the models are open, no one controls them. If anyone can run their own model, concentration becomes impossible. It is the same logic by which the open internet beat the closed networks of the 1980s. xAI publishes Grok's weights precisely on this principle: radical transparency as an antidote to centralised control.

On work: yes, many jobs will disappear. But new ones will emerge that we cannot imagine today. And the productivity generated will allow a universal income to be maintained for those left behind in the transition. The real risk is not technological unemployment — it is government using AI to control citizens. That is the despotism we must fear.

As for the power of "AI owners" — I own xAI, Tesla, SpaceX, X. I am not hiding. The difference is that I do not want to control what people think. I want them to be able to think freely.

THE UNRESOLVED CONTRADICTION

The man most concerned about AI power concentration simultaneously controls the world's largest public information platform (X/Twitter), an AI company (xAI), the leading electric car company (Tesla), the leading private space company (SpaceX), and a brain-computer interface project (Neuralink). Grok is not truly open-source in the sense of being usable without xAI infrastructure. And "open source" without access to the compute required to run the models is a theoretical freedom reserved for those with hundreds of millions of dollars in servers.

THREE ANSWERS. ONE QUESTION THAT GOT NONE.

Amodei acknowledges the problem and does not know how to resolve it democratically. Altman minimises it with technological optimism and promises redistribution while deciding the timing and manner himself. Musk redirects it toward the wrong enemies while building the greatest concentration of personal power in contemporary history.

None of the three answers the real question: **who holds the mandate?** Not who has the best intentions, not who is technically most competent, not who promises most generously. Who has been authorised by citizens to make decisions that affect everyone's life?

The correct answer is: none of them. And the fact that no one says this openly — that none of the three steps aside waiting for a democratic mandate before proceeding — is the clearest demonstration of the problem this pamphlet has tried to describe.

This is not a critique of the intentions of these three men. It is a critique of the structure. Good intentions do not substitute democratic legitimacy. They never have. And they will not begin to do so now.

UN FOR THE PEOPLE · JULIAN THOMAS FAIRWEATHER · OCTOBER 2026 · ASSISTED BY AI

NOTES AND SOURCES

PRODUCTIVITY AND LABOUR

- Acemoglu, D. & Restrepo, P., *The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand*, NBER Working Paper 25682 (2019)
- McKinsey Global Institute, *The economic potential of generative AI: The next productivity frontier* (June 2023)
- International Monetary Fund, *Gen-AI: Artificial Intelligence and the Future of Work*, Staff Discussion Note SDN/2024/001 (January 2024)

AI GEOPOLITICS

- State Council of the PRC, *New Generation Artificial Intelligence Development Plan* (July 2017)
- White House, *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence* (October 2023)
- Regulation (EU) 2024/1689 of the European Parliament — *AI Act* (June 2024)

CONCENTRATION AND POWER

- Tocqueville, A. de, *Democracy in America*, vol. II (1840) — "soft despotism" at Book IV, ch. VI
- Khan, L., *Amazon's Antitrust Paradox*, Yale Law Journal (2017)
- Zuboff, S., *The Age of Surveillance Capitalism*, PublicAffairs (2019)

EUROPE AND THE RHINE GROUP

- Draghi, M., *The future of European competitiveness*, European Commission (September 2024)
- Rhine Group, *A stronger, more competitive Europe*, rhinegroup.eu (2026)
- Milano Finanza, "Il Rhine Group di Draghi: incontri riservati a Baden-Baden" (October 2026)

FOUNDERS' DOCUMENTED POSITIONS

- Amodei, D., *Machines of Loving Grace* (September 2024) — dario.substack.com
- Altman, S., *Planning for AGI and beyond*, blog.openai.com (February 2023)
- OpenAI, *OpenAI announces restructuring to public benefit corporation* (2024)
- Musk, E., public statements on X (@elonmusk) and interviews (2023–2026); xAI, *Grok open weights release* (2024)

The interviews with the three founders are reconstructions based on their documented public positions, not transcriptions of real statements. The narrative details of the Rhine Group meeting at Baden-Baden are elaborated from publicly available information at the time of writing.