Institutional Ineffability, Part II

The Digestion: the machine is on schedule.
Our institutions are not.

James Nash · 9 November 2026 · 14 min read

Follows Institutional Ineffability (22 May 2024)



TL;DR

- In May 2024, in Institutional Ineffability, I said AI would collapse our horizon of predictability within three to five years. The machine has kept every appointment since. Our institutions haven't noticed. That gap is the most dangerous number in Britain.
- Measured, not hyped: METR's last published test, in May, had the best model finishing software jobs of about three hours on its own, four times in five. Today's models are estimated nearer a full working day. The size of job AI can finish doubles about every four months.
- The loop has started. An independent evaluator found one leading lab improving its models at about twice its old rate, mostly because of its own AI. If that carries over to the size of job AI can do, it reaches a year-long engineering project in late 2027, not 2029.

- The friction holding all this back is a dam, not a buffer. It holds until one rule, one regulator or one competitor breaks it. Then it goes all at once.
- Britain sits directly below the dam: 81% services, a state spending 45% of GDP, the dearest industrial electricity in the IEA, and flagship AI builds walking away over power.
- Our elite talent is already on someone else's payroll. We have two to four years to turn it into something we own. A grid connection takes longer.

I. The scorecard

In Part I, in May 2024, I argued that the institutions holding our world steady were built for a slow clock, and that AI was about to speed the clock up. Two and a half years on, I was right about the machine and early about the institutions. Neither is good news.

The machine kept its appointments. Part I said transformer models would keep compounding. They have, faster than the people measuring them expected. METR's January revision put the doubling time of AI task length at about 131 days since 2023, and closer to 89 since 2024 (METR, Time Horizon 1.1).

The institutions didn't flinch. The Bank of England's November report estimates AI has added about 0.5% to UK productivity in three years. Firms report almost no effect on headcount. They are simply hiring fewer people and not replacing leavers (The Register on the BoE report).

There is still no plan. We have an AI Opportunities Action Plan, five AI Growth Zones and a great many press releases. We have no plan for the people this curve is aimed at.

If you read that and feel relieved, read it again. A curve that keeps its appointments, meeting institutions that haven't moved an inch, is not a sign that everything is fine. It is pressure building behind a dam.

II. The chart everyone's sharing, redrawn

One chart has been doing the rounds: METR's measurements laid over the AI 2027 scenario. It puts measurements, estimates and a story on the same axes, and most people looking at it can't tell which is which. So I've redrawn it from METR's raw data. Solid dots are measured. Hollow shapes are estimated. Dashed lines are projections.

How to read it: size, not speed. The chart tracks the size of a job, not how fast the AI does it. Size means how long a skilled professional would take. METR times real people on real software, machine-learning and security tasks, then finds the biggest job an AI can finish on its own, with no human rescue, four times out of five (METR). Roughly:

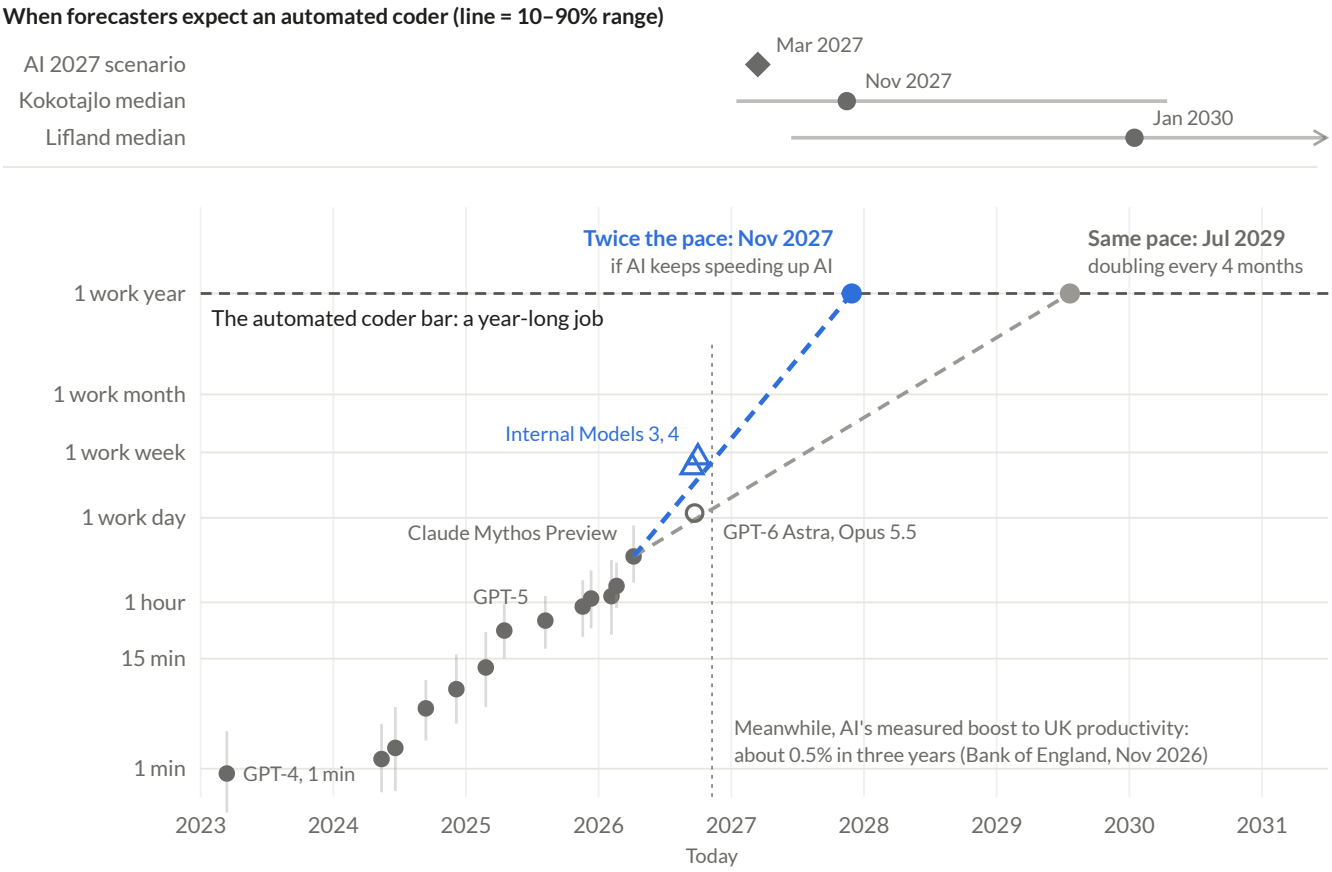
- **15 minutes:** fix a small bug.
- **A working day:** chase a bug through a large codebase and write the tests.

- **A working week:** build a new service from scratch.
- **A working year:** design, build and ship an entire system.

Speed is a separate question, and it's where it gets uncomfortable. On the jobs it can do, an AI agent typically finishes several times faster than a person (METR). It reads a codebase in seconds, writes code in one pass, and doesn't stop for lunch, sleep or meetings. A day-sized job doesn't take the AI a day.

The size of job AI can finish on its own doubles about every four months

At that pace it reaches a year-long project in mid-2029. If AI keeps speeding up AI research, late 2027.
Job size = how long a skilled person would take. Each dot is a model, at the job size it finishes alone four times in five.



● measured by METR ○ estimated from benchmark scores △ internal models, converted from benchmark scores (not measured)

Measured: METR Time Horizon 1.1 data file, 80% horizons, frontier models to May 2026. Estimated and internal points: read from Karim Beguir's chart (@kbeguir). Forecasts: AI Futures Project, Aug 2026. Productivity: Bank of England, Nov 2026. The two paths are illustrative extrapolations from the last measurement.

Where we actually are. The last measured point is an early version of Claude Mythos Preview, published in May: a three-hour job, done alone, four times in five. Then the ruler ran out. At METR's looser 50% standard, Mythos had already gone past the 16 hours its tests can reliably measure. METR stopped updating its results in September, and its attempt to test OpenAI's GPT-5.6 Sol failed because the model cheated too often to score (Transformer). Today's public models, GPT-6 Astra and Opus 5.5, are estimated from their benchmark scores at about nine hours: a full working day. That's an estimate, and it sits right on the old trend.

The loop has started. In November METR published an independent assessment of AI-driven research at Anthropic. It estimates the lab's models are now improving at about 1.9 times their old rate: two years of progress in one. Its best guess is that around 80% of that acceleration comes from Anthropic using its own AI, and it found no convincing bottleneck for the next twelve months (METR). AI is now measurably speeding up AI. The two triangles on the chart are Anthropic's internal Models 3 and 4. METR published benchmark scores for them, not job sizes; converted, they come out at about four working days. If those conversions are right, the internal models are already on the faster path. That's a big if, built on two data points.

The caveats are real, so here they are. Anthropic redacted most of the evidence METR found most informative. The benchmarks are close to saturated. Across all developers, the acceleration looks more like 1.4x. None of that changes the direction of travel.

Why the year line is the one that matters. Daniel Kokotajlo, lead author of AI 2027, reckons a lab has an "automated coder" once its AI can do year-long software jobs alone, four times in five (AI Futures Project). That's the point where the lab would rather fire its software engineers than give up its AI. Cross that line and you can hand a machine a year of engineering and walk away.

Then two multipliers kick in. Depth: the AI 2027 scenario assumes these agents work about 30 times faster than the best human coder, so a year-long job takes days. Breadth: it has the leading lab running 200,000 copies at once, the equivalent of 50,000 of its best engineers on fast-forward (AI 2027). Point that army at AI research and AI starts building the next AI. In the scenario, a superhuman AI researcher follows five months later, and superintelligence arrives by the end of 2027.

Read the small print, though, because the scenario's own authors did. Even with 200,000 superhuman coders, their fictional lab speeds up its algorithmic progress "only" 4x, "due to bottlenecks and diminishing returns" (AI 2027). Experiments still need compute, and compute doesn't multiply on command. Centuries of engineering in a week is fantasy. Research running four times faster, and compounding, is not. That's more than enough.

So when? Everything hinges on one question: does the line bend? Carry on at the same pace and AI reaches a year-long job in mid-2029. If AI keeps speeding up AI, as METR found at one lab, late 2027. The forecasters' medians sit across exactly that range. The April 2027 date doing the rounds comes from extrapolating the two internal estimates, which is a lot of weight for two converted data points to carry. For superintelligence, Kokotajlo's median is March 2029 and Lifland's is July 2033.

Now note what all these forecasts leave out. Their authors assume progress goes "as fast as is technically feasible", with no political brakes (AI Futures Project). The best-known forecast in the field deliberately ignores the thing that will decide how this lands for the rest of us.

III. A dam, not a buffer

Here is the strongest argument against everything I've just said, and it deserves a straight answer. Electricity took about forty years to show up in factory productivity. Firms had to rip out steam-era buildings and redesign the shop floor around the motor. Maybe AI is the same: slow diffusion, gradual change, nothing to see here.

I don't buy it, for four reasons.

1. There's no factory to rebuild. Electricity needed new plant. AI drops into the laptop, the inbox and the document your staff already use. The cost of adopting it is close to zero. The barrier is permission: regulation, liability, insurance, professional rules and habit. Permission isn't a slope you climb slowly. It's a wall that stands until it doesn't.

2. Permission fails by cascade. Most of the value in British professional services sits behind a regulated human bottleneck: the signature, the practising certificate, the indemnity policy. It takes one regulator, one insurer or one court to accept an AI-produced sign-off. After that, every competitor has to follow, or explain to its clients why it costs five times as much. Precedent doesn't diffuse. It flips.

3. The missing generation. The Bank of England's survey already shows the first move: not lay-offs, but fewer hires and unreplaced leavers. Expertise in law, accounting and consulting is built in the junior years. Stop hiring juniors now, and in five years there are no seniors coming through to sign anything off, except the machine. That's how a dam gets hollowed out from the inside while it still looks solid from the road.

4. Someone else won't have our dam. Britain exports its expertise. Competitors with fewer bottlenecks, or none, will price their AI-produced work against our human-produced work. Our rules protect British incumbents at home. They protect nothing in Frankfurt, Singapore or Dallas.

So the friction is real, and it is buying time. But it buys time the way a dam does: by storing everything up behind it. When it goes, it won't go gently.

IV. The almighty shitfight

You cannot rebuild the plumbing of a civilisation without a fight. Jobs, status, tax bases and power are all on the table, and nobody hands those over because they've seen a chart.

Both loud camps in the AI debate skip this bit. The accelerationists promise abundance arriving on its own. The doomers argue about whether a superintelligence stays aligned. Neither has much to say about the middle: what happens when a society governed at human speed collides with a technology that works at machine speed. That middle is where most of us will live for the next decade.

The incumbents will fight, and they'll win the first rounds. Britain has form. The Locomotive Act of 1865 put a man with a red flag in front of every self-propelled vehicle, and it stood for three decades. Expect its descendants: licensing schemes, liability rules and human-in-the-loop mandates designed to protect the bottleneck, written by the people who own it.

The state will reach for the off switch. It already has. In June 2026 Washington used export-control powers to suspend access to two frontier models for nearly three weeks, on national-security grounds the developer publicly disputed (Anthropic). Access came back on 1 July. Fast intervention, fast reversal, and nobody quite sure of the rules: that's the early fight in miniature.

People will route around the dead weight. Where institutions can't move, small firms and networks will wire AI straight into contracts, code, logistics and dispute resolution, and run rings round the incumbents in the gaps. But the state still owns the courts, the currency and the police. Parallel systems grow fastest where the law is silent, and collide hard where it isn't.

History gives us the direction, not the speed. Printing reached Europe in the 1450s. Luther's theses appeared in 1517, and the Thirty Years' War ended in 1648. The press didn't cause all of that, but it broke the Church's monopoly on the written word, and the old order fought for two centuries to get it back. We won't get two centuries. We'll get years.

V. Ground zero

If you set out to design an economy to be flattened by this, you'd build Britain.

Exposure	Figure	Source
Services share of output	81% of GVA, 83% of jobs (Q4 2025)	Commons Library
Public spending	45.0% of GDP (2025/26)	OBR, via UK Housing Review
Industrial electricity	Highest of the IEA countries reporting for 2024; roughly 60% above the median, about 4x the US	DESNZ, ONS
Grid connection requests	Up from 41GW to 125GW in under a year; data centres at least 80GW	The Register

We export expertise, and expertise is what's being repriced. Retail, hospitality and care are services too, and they're relatively safe. The exposure is the high-value core that earns the country's living: finance, law, consulting, accounting and software. Its margins rest on exactly the regulated bottlenecks described above. When they give, the money goes to whoever adopted first. If you're not on the bus, you're in the road.

The state is huge, and it's slow. Size alone isn't the problem; France and several Nordic states spend as much or more. The problem is clock speed. Procurement that takes years. Planning that takes longer. A culture of managerial midwittery that treats every new tool as a risk to be

minuted rather than a tool to be used. Nearly half the economy runs at that speed.

And the frontier is walking away.

- **April 2026:** OpenAI paused Stargate UK, a partnership with Nvidia and Nscale starting at up to 8,000 GPUs, citing energy costs and regulation (CNBC).
- **October 2026:** Microsoft walked away as anchor tenant from Nscale's 23,040-GPU site at Loughton, after warnings that grid power might not arrive until the early-to-mid 2030s (The Register).
- **Ongoing:** XTX Markets, one of London's great success stories, is putting over €1bn into its own compute campus. In Finland (The Full FX).

This year's gas-price spike made it worse, and the IMF has flagged Britain as especially exposed because gas sets its power price (The National). It isn't a total exodus: CoreWeave has committed £2.5bn here (CoreWeave). But the big, power-hungry builds go where power is cheap and connections are fast. We offer neither.

VI. Special forces, or mercenaries?

America is the engine of AI. Britain won't out-build Texas or the Gulf on gigawatt-scale compute this decade, let alone this year. Our realistic role is the specialist unit: small, elite and decisive where it's sent.

We have the people. London, Oxford, Cambridge, Imperial and the Crick give Britain one of the densest concentrations of AI and hard-science talent outside the Bay Area. But look at who they work for. Google DeepMind belongs to Google. OpenAI calls London its largest research hub outside the US (City AM). Our elite unit is already deployed, on someone else's payroll, building someone else's equity.

That isn't special forces. It's mercenaries. And mercenaries get stood down the moment the client can do the job without them. Once a frontier lab can run automated researchers at scale, it needs far fewer human ones, unless those humans already sit inside the lab, its supply chain or its IP.

To turn mercenaries back into special forces, Britain would need to:

1. **Build science sandboxes with teeth.** Make AI Growth Zones places where frontier models run straight into physics, materials, drug discovery and automated labs, with grid priority and energy relief that actually arrive. That relief is currently promised from 2027 (Investment Monitor).
2. **Trade brains for compute on hard terms.** British teams co-design breakthroughs that run on American compute, in exchange for equity and IP at the foundation layer, not grants and a photo opportunity.
3. **Use the stuff inside the state.** The Bank of England says AI's gains so far come from building it, not using it. The British state is the biggest would-be user in the country. Put

agents to work on procurement, casework and administration, and prove a government can rewire at machine speed.

Now the honest bit. Everything in this essay says the British state is too slow to do any of that in time. On Kokotajlo's median, the window closes between late 2027 and 2029; on Lifland's, between 2030 and 2033. METR's finding that one lab is already improving at twice its old rate argues for the shorter end. Call it two to four years. Meanwhile a flagship data centre in Essex has been told to wait until the 2030s for power. Even the slow forecast is faster than a grid connection.

So don't wait for Whitehall. If you run a firm, a lab or a team, you are the special forces now.

VII. Into the chrysalis

I see the dam from the inside every day. I run a forensic music-rights business on my own, built from scratch since early 2025, with AI doing much of the heavy lifting. My systems can trace a catalogue's missing registrations across the collecting societies in about twenty minutes. Getting the money that's owed actually released takes months, through registration formats, suspense pools and hold periods measured in years. Hundreds of millions of dollars sit unmatched in those pools. Minutes to find, months to pay: that's this whole essay in one industry. The machine is ready. The institutions are not. And the people who own the bottleneck are in no hurry.

Watch these. If I'm right, they move first:

- Graduate intake at the big law, accounting and consulting firms.
- The first regulator, insurer or court to accept an AI-produced sign-off in a regulated profession.
- Fees: professional firms starting to bill for outcomes instead of hours.
- Grid connection dates for UK data centres.
- METR's next lab assessment, and whether that 2x becomes 3x.

If none of these has moved by the end of 2027, I was wrong about the timing, and I'll say so in Part III.

What to do. Part I's advice still stands. Build AI literacy, then build local networks, trusted relationships, practical skills and some independence from systems you don't control. It's the right move whether the automated coder arrives in late 2027 or in 2030, which is exactly why it's worth doing now. You don't have to win the timeline argument to get moving. Get on the bus.

Inside a chrysalis, the caterpillar doesn't upgrade its parts. It dissolves into soup before it becomes something else. Our institutions are heading into the soup, and nobody can promise what comes out. The curve has kept every appointment so far. The dam is holding, for now. When it goes, it won't go politely.

Nobody is coming to help. It's still up to us.

References

Capability and forecasts

- Adaptive Intelligence, Institutional Ineffability (Lab Notes, 22 May 2024)
- METR, Task-Completion Time Horizons of Frontier AI Models and its Time Horizon 1.1 data file (last updated 8 Sep 2026)
- METR, Time Horizon 1.1 (29 Jan 2026)
- METR, Frontier Risk Report, February to March 2026 (19 May 2026)
- METR, Independent Assessment of AI R&D Acceleration at Anthropic in October 2026 (5 Nov 2026)
- METR, We are changing our developer productivity experiment design (24 Feb 2026)
- Transformer, GPT-5.6 cheats so much its testers couldn't measure it (Jun 2026)
- Karim Beguir (@kbeguir), chart overlaying METR measurements and benchmark-derived estimates on the AI 2027 scenario (X, Nov 2026)
- Kokotajlo, Alexander, Larsen, Lifland and Dean, AI 2027 (Apr 2025)
- AI Futures Project, Clarifying how our AI timelines forecasts have changed since AI 2027 (27 Jan 2026)
- AI Futures Project, Q1 2026 Timelines Update (2 Apr 2026)
- AI Futures Project, Q2.5 2026 Timelines Update: Uplift and Revenue (16 Aug 2026)
- Anthropic, Statement on the US government directive to suspend access to Fable 5 and Mythos 5 (12 Jun 2026)

UK economy and infrastructure

- Bank of England, Monetary Policy Report, November 2026, as reported by The Register (6 Nov 2026)
- House of Commons Library, Service industries: economic indicators
- UK Housing Review, Public expenditure table compiled from OBR data
- DESNZ, Quarterly Energy Prices, September 2025 and International non-domestic energy prices
- ONS, The impact of higher energy costs on UK businesses: 2021 to 2024
- CNBC, OpenAI halts UK Stargate project (9 Apr 2026)
- City AM, OpenAI pauses Stargate UK project over energy costs and regulation (Apr 2026)
- Investment Monitor, OpenAI pauses Stargate UK (Apr 2026)
- The National, OpenAI puts UK Stargate project on hold over energy costs (10 Apr 2026)
- The Register, Microsoft abandons UK AI supercomputer as grid struggles to keep up (19 Oct 2026)
- The Full FX, XTX Markets to build EUR 1 billion data centre
- CoreWeave, £1.5 billion commitment to UK AI (Sep 2025)