ON TIMING, CAPITAL, AND GENERAL PURPOSE TECHNOLOGY

The Meh Case for AI

Why the most likely outcome is neither transformation nor catastrophe, why it will still be expensive, and why the gains arrive a decade after the money does.

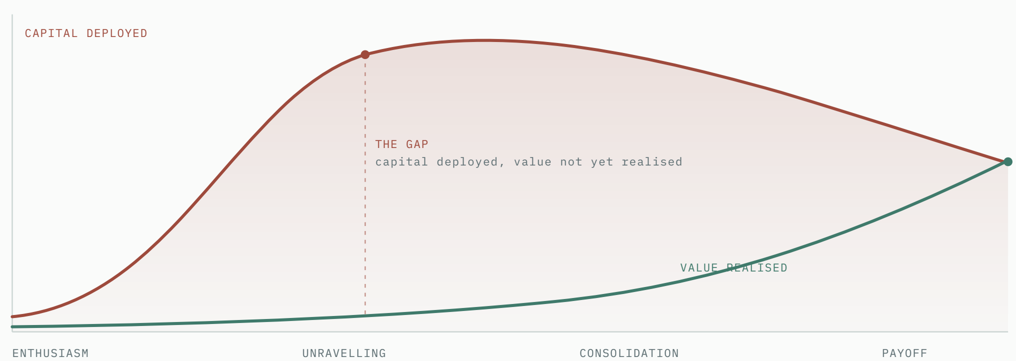

The shape the argument describes, and the shape fibre optic cable, British railways and the canals all took before it. Capital arrives on the schedule of enthusiasm. Value arrives on the schedule of institutions. The people who are right about the technology and wrong about the interval are, financially, indistinguishable from the people who are simply wrong.

I The scenario nobody argues for

Two stories about artificial intelligence get told. In one, the technology compresses a century of scientific progress into a decade, and the only serious question is what happens to human purpose afterwards. In the other, it slips control and the question is whether anyone is left to ask questions.

Both are interesting. Neither is likely.

There is a third story. AI is a real and useful general purpose technology whose economic benefits arrive roughly a decade after the capital that funded them. The gap in between produces a financial unravelling of some significance. The industry consolidates into a handful of firms large enough to raise the question of public rescue. And by the mid-2030s the world is meaningfully better off, in ways that will seem in retrospect to have taken a surprisingly long time to show up.

This is the meh case. It is under-argued for an obvious reason: nobody's incentives point at it. Labs raising capital cannot sell it, critics building audiences cannot dramatise it, and journalists cannot headline it. It has no constituency. It is also, on the evidence, the modal outcome.

What follows is the argument, with the numbers.

II The pattern is not new, and it has a shape

Start with the analogy that fits best.

Between 1996 and 2001, following the Telecommunications Act, American telecom firms deployed somewhere between $500 billion and $1 trillion in capital. Bond issuance alone exceeded $500 billion. They laid more than 80 million miles of fibre optic cable, which represented over three quarters of all digital wiring installed in the United States up to that point.

Large portions of it were never lit. Contemporary industry estimates commonly placed the unused share above eighty percent years after construction, though those figures came from industry rather than audited measurement and the range across sources is wide.

The firms that laid it were destroyed. Global Crossing filed for bankruptcy in January 2002 listing $22.4 billion in assets against $12.4 billion in debt, having spent about $15 billion in five years and never turned a profit. Its peak market capitalisation had briefly exceeded General Motors. WorldCom followed in July, the largest bankruptcy and the largest accounting fraud in American history at the time. By mid-2002 some twenty-three telecom firms had gone under. The industry owed about a trillion dollars. Roughly $2 trillion of the era's $7 trillion market capitalisation decline was telecom alone.

And then the fibre got used.

Distressed funds bought the bonds at pennies on the dollar. Level 3 acquired Genuity's network assets for about $60 million, WilTel for $486 million, ICG for $163 million, and eventually Global Crossing itself for $3 billion. The demand that justified the buildout arrived between 2005 and 2010: YouTube, the iPhone, Netflix streaming, Amazon Web Services. All of it ran on cable laid by companies that no longer existed, bought by their successors for a fraction of construction cost.

The technology thesis was correct. The timing was wrong by roughly five years. Everyone who was right too early lost everything, and the asset survived them.

This pattern is old enough to be boring. In the British railway mania of the 1840s, Parliament authorised over 260 railway acts in 1846 alone, covering some 9,500 miles. Investment reached about seven percent of GDP. The railway share index peaked in August 1845 and fell roughly two thirds by 1850. About a third of authorised lines were never built, and investors who had bought partly paid shares found themselves facing capital calls on a collapsing asset. George Hudson, the Railway King, was found to have been paying dividends out of capital.

The railways got built anyway. Eight thousand miles of them became the physical substrate of Victorian industry.

The canal mania of the 1790s ran the same course. So did the dot-com bust, which destroyed roughly $5 trillion in market value. It also left behind the infrastructure, and, according to research at the University of Maryland, about half the startups still alive five years later.

The recurring structure is worth stating plainly, because it is the spine of this essay. **Capital-intensive general purpose technologies are systematically overbuilt during the enthusiasm phase. The first owners are wiped out. The asset is acquired cheaply by successors. And the productive use arrives late enough that the original thesis looks foolish for years before it looks obvious.**

III Why the gains come late

The overbuild pattern describes what happens to capital. It does not explain why the benefits lag. For that, the best evidence comes from a paper written in 1990 about a technology commercialised in 1880.

Paul David's "The Dynamo and the Computer" asked why electrification, obviously transformative, took roughly forty years to appear in productivity statistics. Edison's lighting system dates from 1880. By 1900, electric lighting reached about three percent of American residences and electric motors accounted for under five percent of factory mechanical drive. Productivity did not accelerate until the 1920s.

The reason was not the technology. It was the factory.

Nineteenth century factories were built around a central steam engine driving overhead shafts, belts and pulleys. Machines had to be positioned by their proximity to power, and buildings were multi-storey to shorten the drive shafts. Dropping electric motors into that arrangement changed almost nothing, because the constraint was architectural. Productivity jumped only when factories were physically rebuilt around distributed unit drive, with each machine independently powered, laid out by the logic of the work rather than the geometry of the shaft. That required tearing down buildings and reorganising production, which required a generation.

Erik Brynjolfsson, Daniel Rock and Chad Syverson formalised this in 2021 as the Productivity J-Curve. General purpose technologies require complementary intangible investment: process redesign, new business models, retraining, organisational restructuring. That investment is expensive and poorly measured, so measured productivity is understated early, while the intangibles are being built, and overstated later, when they are being harvested. Their adjusted total factor productivity was about sixteen percent higher than official measures for software and hardware intangibles by the end of 2017. For AI specifically they find the effects "small but growing."

Jovanovic and Rousseau date electrification's diffusion from about 1894 to a plateau around 1929. Bresnahan and Trajtenberg's original formulation of general purpose technology, with its criteria of pervasiveness, innovation-spawning and continuous improvement, describes something that diffuses across an economy over decades rather than quarters.

The relevant question for AI is not whether the technology works. It is how much of the value requires an organisation to be rebuilt around it, and that fraction appears to be large. A firm that installs a language model without changing how work is allocated, reviewed, priced or staffed captures the equivalent of an electric motor bolted to a shaft.

IV Why the value does not become revenue

There is a second reason to expect the gap between capability and return to be wide, and it is more specific to this cycle.

Value created is not value captured, and the two can diverge enormously. William Nordhaus estimated that innovators historically capture about 2.2 percent of the social value they create. Broader work in the appropriability literature puts producer capture somewhere between two and fifteen percent. The gap goes to consumers, and that is generally a good thing for everyone except the party that has to service the debt.

For AI the leakage runs through four channels, and three of them are structurally closed to a vendor.

Competition passes it to consumers. In competitive markets, a cost reduction becomes a price reduction. This is the ordinary case and it is why highway travel time savings, which are the principal justification for sixty years of American infrastructure investment and are officially valued by the Department of Transportation at around nineteen dollars per person-hour, have never been paid for by a single road user.

Administered pricing confiscates it by formula. In healthcare, Medicare's budget neutrality requirement means that when volume or cost grows anywhere, the conversion factor is cut to hold total spending flat. Inflation-adjusted reimbursement to radiologists per beneficiary fell by roughly a quarter to nearly a third between 2004 and 2021, depending on the inflation measure and procedure basket used, while work relative value units per beneficiary rose about thirteen percent. The productivity showed up and the payment fell to match. This is not a market failure. It is the system working as designed, and it covers a large share of a sector that is roughly a fifth of the American economy.

Cost-measured sectors have no mechanism to pay at all. Public education output is measured in the national accounts at input cost, principally salaries. A district that becomes more productive does not generate a surplus it can spend on software. It generates a smaller number in the output statistics. The same is true of most government activity.

Only concentrated markets with pricing power let the gain reach the firm, and therefore let a vendor price against it.

Layered on top of this is a supply-side problem specific to AI: the product is commoditising faster than the capital can be recovered. Epoch AI finds the best open-weight models trailing the closed frontier by about three to four months. Inference for GPT-4 class capability has fallen by something between sixty and a thousand times in three years depending on the quality bar. A firm spending billions on a training run is selling into a market where a materially similar capability is available at a fraction of the price within two quarters.

The result shows up in the aggregate, though the aggregate is contested. There is no agreed definition of AI revenue: analysts variously count model providers only, AI software, cloud AI services, infrastructure, or enterprise deployment, and the totals differ by more than an order of magnitude depending on the choice. Contemporary industry estimates put AI-related revenue at roughly one tenth of annual infrastructure spending. David Cahn of Sequoia estimated a roughly $600 billion annual gap between what the infrastructure requires and what the ecosystem earns. Goldman Sachs estimated the buildout would need $1 trillion or more in annual profit against a $450 billion consensus. These are estimates rather than accounting identities, and they should be read as such.

V The financing is more fragile than the technology

The interesting question is not whether that gap closes eventually. It is what happens while it is open, and the answer depends on how the buildout was financed.

Hyperscaler bond issuance ran about $121 billion in 2025 against a $28 billion average across the preceding four years. That is the number that matters, because it marks the point at which the buildout stopped being funded from cash flow. Alongside it sits a larger pool of private credit and off-balance-sheet structures, plausibly several hundred billion dollars, with materially lower disclosure and weaker creditor protections. AI is now the single largest driver of American corporate issuance.

Three financing structures sit underneath that, and they fail at different points.

Hardware-collateralised lending fails first, and is already stressed. The reference case is a $7.5 billion facility at roughly eleven percent variable, secured on GPUs and customer contracts, with repayment beginning in January 2026 as collateral values were falling. The structural flaw is simple: collateral depreciates at roughly the rate at which compute cost per unit improves, which is about a doubling every eighteen to twenty-four months, while principal amortises on a fixed schedule. The two curves cross. On plausible assumptions, loan-to-value breaches one within two years. The interest rate is not the trigger. The depreciation curve is.

Convertible debt fails on a calendar. A $2.25 billion convertible priced at a 1.75 percent coupon is not a financing achievement. It is a sold equity option, where upside substitutes for yield. If the stock never reaches conversion, the issuer refinances at market. A 1.75 percent coupon refinancing at seven percent is a fourfold increase in interest burden, and typical five to seven year tenors put that wall around 2029 to 2031. This is the only tipping point with a date attached, which makes it the one that can be prepared for.

Investment-grade issuers fail last and differently. Firms funding substantially from cash flow are not constrained by the coupon. They are constrained by the discount rate applied to revenue that arrives years out. On a stylised buildout with revenue beginning in year five, net present value turns negative somewhere between a ten and twelve percent required return. They stop building because projects stop clearing hurdle, not because credit closes.

There is also a reflexive loop worth naming. At roughly a quarter of projected investment-grade issuance, the buildout is large enough to affect the rates it borrows at. More issuance means more duration supply, which pushes long yields up, which raises the hurdle for the next tranche. This is self-limiting rather than explosive, and it is the mechanism that most resembles telecom in 1999.

The leading indicator is already moving, and it is not rates. Bond coverage on hyperscaler issues fell from about five times in February 2026 to below two times by July, at spreads sitting near their tightest since 1997. Demand depth collapsing while price holds is an unstable arrangement. It is the market repricing risk before any central bank does.

To see why this matters beyond the sector, consider a figure from Jason Furman: investment in information processing equipment and software is about four percent of American GDP, but accounted for roughly ninety-two percent of GDP growth in the first half of 2025. Excluding those categories, the economy grew at an annualised rate of 0.1 percent.

An unravelling in AI capital expenditure would not be a sector event.

VI What consolidation looks like, and the bailout question

If the gap persists long enough, the industry consolidates. This is not speculation but the observed pattern in capital-intensive shakeouts, formalised by Steven Klepper and visible in telecom's collapse into a handful of long-haul owners.

The mechanics are predictable. Firms without cash flow fail first: the neoclouds, the compute resellers, the model companies whose unit economics depend on subsidised inference. Their assets are acquired at a fraction of construction cost by firms with balance sheets. The survivors emerge larger, more vertically integrated and considerably more concentrated than the sector that entered the cycle.

Whether that ends in public rescue is a genuine question rather than a rhetorical one, and it deserves care.

The framework for systemic designation turns on size, interconnectedness, substitutability, complexity and cross-border activity, not size alone. That framework is legally specific to finance. If an AI rescue happens it will look less like the bank bailouts and more like the automotive one. That rescue was justified on industrial rather than financial grounds: supplier networks, regional employment concentration, and an estimated $39 to $105 billion in lost tax revenue and unemployment costs had General Motors and Chrysler been allowed to fail.

The fiscal record there is worth stating precisely, because it cuts both ways. The Troubled Asset Relief Programme overall closed with a headline profit of roughly $15.3 billion. The automotive component specifically lost money, somewhere between $9.2 billion on Treasury figures and $12 billion on Congressional Budget Office accounting, with higher estimates from critics. A rescue is not automatically a fiscal disaster. It is also not free, and the auto piece is the relevant precedent.

The conditions under which AI infrastructure would qualify are identifiable in advance. AI-linked debt would have to propagate widely enough through private credit and pension allocations that failure transmits into unrelated balance sheets. A small number of firms would have to become the operational backbone for cloud, payments and government services, such that substitutability is genuinely low. And regional employment and tax bases would have to concentrate around datacentre construction. None of those is currently established. All are worth monitoring, and the monitoring is cheap.

VII The uncomfortable asymmetry: crime adopts first

There is a category of AI adoption that faces none of the friction described above, and it is worth confronting because it cuts against the comfortable reading of a slow transition.

Legitimate adoption is slowed by regulatory headcount floors, occupational licensing, procurement cycles, safety review, liability exposure, board approval and integration cost. Criminal enterprises face none of these. They have no compliance function, no auditor, no scope-of-practice statute and no insurer. **One would expect the least regulated sectors to adopt fastest, and crime is among the least regulated sectors there is.**

The mechanism is the one this essay has applied throughout. Fraud at scale is labour-constrained. Romance and investment scams require human operators sustaining conversations over weeks or months. Business email compromise requires someone who can write convincingly in a particular register. Those are exactly the constraints a language model relieves, and relieving them collapses the marginal cost of an attempt while expanding the addressable victim pool.

The FBI's Internet Crime Complaint Center recorded $20.877 billion in losses in 2025 across 1,008,597 complaints, a twenty-six percent increase year over year and nearly triple the 2021 figure. About eighty-five cents of every dollar lost came from cyber-enabled fraud that exploited a human decision rather than a technical vulnerability. Business email compromise alone accounted for $3.05 billion. Phishing losses grew 208 percent. Elder fraud reached $7.75 billion, up fifty-nine percent.

For the first time, the 2025 report tracked artificial intelligence as a distinct descriptor: 22,364 complaints and $893 million in losses, concentrated in investment fraud at $632 million.

Here the honest reading supports the meh case rather than undermining it. **AI-attributed losses are about 4.3 percent of the total.** Even in the sector with no institutional friction whatsoever, no compliance burden and every incentive to adopt immediately, artificial intelligence accounted for a twentieth of the damage in its first year of measurement. The FBI notes the figure understates reality, because victims frequently cannot tell whether AI was involved. But the direction of the finding is clear enough: AI is not changing what crimes are committed. It is scaling them, and the scaling is real but early.

This has two implications.

The first concerns open weights. Capable models with published weights are available to anyone, and the marginal uplift they provide to a fraud operation is real. There is a serious argument that frontier capability should remain closed for this reason. The difficulty is that the position is close to unenforceable in practice: open-weight models trail the closed frontier by about four months, multiple jurisdictions release them, and a capability gap measured in a single quarter does not constitute a control regime. A policy that depends on capability remaining scarce is a policy that expires quarterly.

The more defensible position accepts diffusion and shifts to the defensive side of the ledger. If frontier capability is going to be broadly available regardless, the argument for directing efficiency gains toward hardening software systems, verification infrastructure, authentication and consumer protection is straightforward. Voice-based verification is already compromised by cloning. Detection of awkward phrasing was never a complete defence and is now no defence at all. The controls that survive are the ones that do not depend on a human noticing something wrong.

The second implication is economic, and it belongs in the productivity accounting. Security expenditure is a cost of protection, not an addition to output. If AI raises attack volume and defenders must spend more to hold position, that spending is resources diverted from production. Some fraction of AI's measured output will be consumed by an arms race rather than appearing as welfare. In a period when the productivity gains are already arriving late, that is not a rounding error.

VIII The deepest objection: discovery is not the bottleneck

The strongest case for AI transforming growth rather than merely improving it runs through science. Aghion, Jones and Jones showed formally that automating the idea production function can raise the growth rate itself, and in the limit produce something like a singularity. If AI makes discovery cheap, the argument goes, the constraint that has bound growth since the Enlightenment is released.

This is a real mechanism and the essay should concede it. It is also, I think, aimed at the wrong constraint.

For many commercially important domains, ideas increasingly appear not to be the binding input. Validation does.

Consider what has actually happened where AI-accelerated discovery is furthest advanced. In pharmaceuticals there are between 173 and 200 AI-originated programmes in clinical development, up from roughly two dozen in late 2023. As of this writing, none has yet received regulatory approval. Phase I success rates for AI-discovered molecules run at eighty to ninety percent against forty to fifty-two percent traditionally, which is genuinely impressive. Phase II success runs at about forty percent against a traditional range of twenty-nine to forty percent, which is no advantage at all.

The interpretation that fits is uncomfortable for the transformative case. AI appears to be improving the odds of finding a molecule that is safe. It has not yet demonstrated any improvement in finding one that works. And the failure mode that dominates drug development is efficacy, not safety. The most credible sceptical reading is that early AI programmes selected easier targets, a pattern well documented in the early days of monoclonal antibodies, where initial success rates reflected careful target choice rather than technical superiority.

Materials science tells the same story more starkly. DeepMind's GNoME predicted 2.2 million crystal structures, of which about 380,000 were assessed as stable, and improved stability prediction hit rates from roughly one percent to over eighty percent. Berkeley's autonomous A-Lab synthesised forty-one of fifty-eight targets in seventeen days. Set those numbers side by side: 380,000 predicted stable structures against forty-one syntheses in seventeen days. Even at that pace, validating the predictions would take centuries. As of early 2026 no AI-originated material had entered commercial production, one showcase compound turned out to have been known since 1972, and a Fritz Haber Institute study found that over eighty percent of AI-recommended candidates would exhibit crystallographic disorder in practice.

This is not a failure of the technology. It is the discovery constraint being relieved and revealing that it was never the binding one.

The broader literature points the same way, and it long predates AI. Park, Leahey and Funk analysed 45 million papers and 3.9 million patents across six decades and found both becoming systematically less disruptive, a pattern holding nearly universally across fields. The finding is contested, with critics attributing part of the decline to dataset artefacts, and that caveat is worth carrying. But it sits alongside documented declines in research productivity in semiconductors and pharmaceuticals, and Bloom, Jones, Van Reenen and Webb's finding that sustaining Moore's Law required roughly eighteen times more researchers in 2014 than in 1971. Meanwhile university technology transfer offices hold large patent portfolios of which only a small fraction is licensed and a smaller fraction reaches market.

More ideas have not been producing proportionally more progress for some time. There is already a substantial backlog of invented, published, patented technology sitting unvalidated, and it is difficult to see how adding to the front of that queue accelerates the exit.

The bottleneck increasingly appears to lie downstream of invention: in trials, scale-up, regulatory throughput, capital, distribution and adoption. Those are constrained by patient recruitment, physical synthesis capacity, statutory review periods, factory construction and human willingness to change practice. A model that generates a thousand candidate molecules in an afternoon does not shorten a Phase III trial by a day.

The strongest counterargument is not that AI discovers more. It is that AI may automate experimentation itself. Self-driving laboratories, autonomous synthesis, robotic assay pipelines and closed-loop design cycles do not merely propose hypotheses. They test them, without waiting for a graduate student. If the experiment automates, the validation bottleneck opens.

This deserves a serious answer, and I think there are three.

The first is arithmetic, and it is already in the numbers above. Berkeley's A-Lab is automated experimentation, running at full tilt, and it produced forty-one syntheses in seventeen days. Against 380,000 predictions that is roughly nine hundred years of work. Automating the experiment raised throughput by a large multiple and left the backlog essentially untouched, because the discovery side scaled by a larger multiple still. Relieving a constraint faster than you relieve the constraint behind it does not help.

The second is that the binding validation is not laboratory work. It is clinical and regulatory. A Phase III trial requires human subjects, recruitment across sites, institutional review, and statutory observation periods measured in years. No robot administers a drug faster than a patient can take it. The same holds for materials, where the gate is qualification and scale-up in a production environment rather than synthesis at bench scale.

The third is historical, and it is the most uncomfortable for the counterargument. Pharmaceutical research automated experimentation once already. High-throughput screening and combinatorial chemistry were adopted across the industry in the 1990s, raising the number of compounds that could be tested by orders of magnitude. Drug development productivity continued to fall throughout, on the trend that came to be called Eroom's Law. We have run this experiment. Automating the experiment did not produce more approved drugs.

None of that says autonomous laboratories are unimportant. It says that they relieve a constraint which was not the binding one, and that we have direct evidence of what happens when they do.

IX What arrives, and when

None of this argues that AI does not work. It argues that the schedule is wrong by about a decade, and that a decade is long enough for the financing to matter.

The distribution I would defend looks roughly like this.

Three to five years out, the gains are concentrated in narrow, well-specified tasks with fast feedback: code generation, customer support, document processing, translation, first-draft anything. These are real and measurable at the firm level. Call centre studies show productivity gains of fourteen to fifteen percent on average and thirty-four percent for novices. They are also small at the level of the economy, because the affected task hours are a small share of total hours and because the sectors where they land are frequently the ones where the gain cannot be captured or measured. This is the period in which the financing gap is widest and the unravelling most likely.

Five to ten years out, the complementary investment starts to pay. Organisations that rebuilt around the technology, rather than bolting it onto existing process, begin to separate from those that did not. This is the J-curve inflection, and it is where the intangible investment made during the disappointing years finally shows in output. Consolidation is complete by this point. The surviving firms own infrastructure acquired at a discount from the firms that built it.

Ten to fifteen years out, the discovery pipeline delivers. The molecules found in 2026 clear Phase III around 2034 if they clear it at all. The materials predicted in 2024 reach commercial production in the mid-2030s if the synthesis and scale-up work. New occupations that do not currently have names are staffed and growing, in the pattern Autor and colleagues documented when they found that roughly sixty percent of employment in 2018 sat in job titles that did not exist in 1940. Living standards are measurably higher and it is genuinely difficult to attribute the improvement, because by then the technology is infrastructure rather than news.

One acknowledgement belongs here, because this essay is overwhelmingly about software. If general purpose robotics improves faster than expected, physical task automation arrives sooner and this timeline compresses. That is a real possibility and I do not want to argue it away.

I would note two things that make it less of an exception than it first appears. Robotics is capital-gated in exactly the way software is not, requiring manufacturing capacity, supply chains for actuators and high-torque motors, and physical installation. It is therefore subject to the same overbuild-and-lag dynamic this essay describes, only more so.

And the economics turn on utilisation rather than capability. A hundred-thousand-dollar robot amortised over five years costs about three dollars an hour at three-shift utilisation and about twenty-eight dollars an hour at household utilisation, against a loaded human wage nearer seventeen. The viable zone is therefore structured commercial environments running continuously, which is precisely where purpose-built automation already works better. A fixed arm beats a humanoid on cost, speed and reliability for a known task in a known position. The general purpose form factor's advantage is generality, and the economically viable zone is where generality is worth least. Faster robotics would compress the timeline for physical work. It would not obviously compress it for the reasons people expect.

It is worth being concrete about what that world looks like, because an essay called the meh case should describe the meh, and readers otherwise leave remembering only the crash.

By the mid-2030s, software engineering is perhaps twice as productive per engineer, and there are more engineers rather than fewer, because the cost of writing software fell and the quantity demanded rose to meet it. Customer service is mostly automated at the first tier and staffed by a smaller, better-paid group handling everything the system escalates. Legal discovery, document review and compliance monitoring are dramatically cheaper, which means small businesses and individuals receive legal review they were previously priced out of entirely. Personalised tutoring is ordinary rather than remarkable, and its effect on outcomes is real, measurable, and considerably smaller than its advocates promised in 2025.

Drug discovery is faster but not miraculous. The pipeline is fuller, the failure rate at the efficacy stage is roughly what it always was, and the approvals arriving are genuinely useful without constituting a compressed century. Scientific software is ubiquitous and invisible, in the way statistical software became invisible. Logistics, scheduling, forecasting and back-office operations across the economy are quietly more efficient in ways nobody writes about, which is generally how the largest aggregate gains arrive.

Nobody calls it a revolution. Total factor productivity growth is perhaps half a point higher than it would otherwise have been, sustained, which compounds into something substantial over a generation and is nearly invisible in any given year. That is not utopia. It is also not nothing, and it is a great deal better than the decade that preceded it.

The convergence point deserves emphasis, because it is how these transitions always end. Nobody experiences electrification as an event. They experience a factory that works differently and a house that is lit, and the forty years of disappointing productivity statistics disappear from memory entirely.

---

x Why this is the modal case, and what would falsify it

The meh case should be the default not because it is moderate but because it requires the fewest things to go unusually right or unusually wrong.

The transformative case needs four things at once. Productivity gains must arrive without the organisational rebuild that every prior general purpose technology demanded. Validation capacity must expand alongside discovery capacity. Capture must hold up against fast commoditisation. And revenue must arrive before the financing structures break. Each is possible. All four together is a demanding conjunction.

The catastrophic case requires capability to advance faster than every institutional and physical constraint documented here, including grid interconnection queues where the median request-to-operation interval exceeds five years, transformer lead times reaching five years, and regulatory processes with statutory minimum durations.

The meh case requires only that the historical pattern holds: that a capital-intensive general purpose technology is overbuilt, that its benefits require reorganisation to realise, that the first owners of the infrastructure do not capture its value, and that the gains arrive on the timescale of institutions rather than the timescale of engineering.

I would rather be falsifiable than persuasive, so here are the conditions under which I would abandon this argument.

A sustained economy-wide total factor productivity acceleration above about 1.5 percent annually attributable to AI would break the timing claim. Acemoglu's estimate is around 0.07 percent per year over a decade; anything an order of magnitude above that is a different world.

The first AI-originated drug approval accompanied by a demonstrated Phase II efficacy advantage would break the validation-bottleneck argument, because it would show the constraint being relieved rather than merely relocated.

The open-weight gap widening back toward twelve months or more would restore pricing power and break the capture-failure engine.

Hyperscaler AI revenue converging toward capital expenditure, closing the roughly $600 billion gap, would remove the compute-glut premise entirely and with it most of section five.

Any of those would be good news. I do not expect them soon.

XI Coda

The uncomfortable thing about the meh case is that it is simultaneously the least alarming scenario and the one most likely to produce a genuine financial event. A technology that fails outright destroys the capital and stops. A technology that transforms immediately justifies the capital. A technology that works, slowly, on a schedule mismatched to its financing, destroys the capital and then delivers the benefits to whoever buys the assets afterwards.

That is what happened to fibre, to the railways, and to the canals. In each case the people who were right about the technology and wrong about the timing were indistinguishable, financially, from the people who were simply wrong.

A decade later, everyone uses the thing without remembering it was ever in doubt.

The asset survives. The investor frequently does not. History suggests that this is not the failure mode of transformative technologies. It is often how they arrive.

you@example.com

Subscribe

No tracking, no list rental, unsubscribe in one click.