PAUL GIBBONS ADVISORY

# THE PARADOXES OF AI

Nine contradictions in enterprise AI with guidance for leaders

DISCUSSION DOCUMENT — BETA

Paul Gibbons | Adaptive Adoption™
paulgibbonsadvisory.com | October 2026

2026 ENTERPRISE AI ADOPTION SERIES
One of seven discussion documents

2026 ENTERPRISE AI ADOPTION SERIES

# Other whitepapers in this series

The Paradoxes of AI is part of a seven-part 2026 whitepaper series on enterprise AI adoption.

## Adaptive Adoption

Executive Briefing

## Change Management Was a Breakthrough

Was it Enough?

## Master the Craft

Why AI Capability Is a Craft

## Solve for Trust

The RIST Framework

## The Case for a Chief AI Officer

## The Leadership Delta

Why AI Leadership Is Different and What To Do About It

Paul Gibbons · Paul Gibbons Advisory

paulgibbonsadvisory.com

# The Paradoxes of AI

## Nine contradictions in enterprise AI with guidance for leaders

**Paul Gibbons**

### THE ARGUMENT

Since 2023, hundreds of billions have been spent on enterprise AI. The returns are harder to find.

AI has made intelligence dramatically cheaper, faster, and more widely available. Yet more intelligence has not reliably produced more value.

Technology adoption has always produced paradoxes that often become intelligible only in retrospect. Those once obscure ideas, for example Jevons Paradox and Goodhart's Law, are now part of the AI strategy canon. But they are from another era.

AI adds relevance to these older puzzles, but has surfaced new, unresolved conundrums.

This paper introduces **nine new paradoxes** from my research and practice.

Some extend familiar economic arguments; others arise from the peculiar conditions of abundant machine intelligence. Individual capability expands faster than organisations can absorb it, creating more work to verify, more options to prioritise, and more dependencies to coordinate—while organisations are still redesigning work, developing judgement, and establishing what deserves trust.

The critical challenge is to redesign systems and develop human capability while the technology keeps changing. Leaders, many of whom are only passingly familiar with AI's potential, have to make decisions that will likely affect whether their business survives in the next decade, let alone prospers.

This paper offers practical guidance for those dilemmas, grounded in emerging evidence and the history of technology adoption.

# Business paradoxes from the pre-AI era

*“What made it possible for me to exercise freedom of choice also made it impossible for me to know the future.” — Ted Chiang, Story of Your Life, in film as Arrival*



In 1937, while the modern corporation was taking shape, Ronald Coase asked why firms exist at all. His answer turned on the relative costs of coordinating work inside an organisation and arranging it through the market. The differential, between internal and external frictions, determines the firm's boundaries.

AI reopens that question.

A small AI-powered team can quickly complete work that takes entire departments working for months in a large organisation. **Where AI lowers the cost of buying and coordinating expertise outside the firm, the case for scaling internally weakens.** A fluid network of microenterprises is the vision – the behemoth firm replaced by a swarm.

The counter-argument goes that platforms from providers such as Microsoft, Google, or AWS could, in theory if not yet in reality, make coordination easier within firms also, allowing them to leverage capital and capability in novel ways.

Whether firms grow or shrink will depend on which coordination costs, internal or external, fall furthest. (My money is on the swarm: after 35 years in tech adoption, I've seen more sluggishness, turf wars, internal frictions, and sclerotic cultures than I've seen agility. But as Niels Bohr, Danish Nobel winner, said: "It is difficult to make predictions, especially about the future.")

Alongside Coase, the social sciences have produced other paradoxical insights:

- Jevons, in 1865, noticed that more efficient steam engines could mean **more** coal burned, not less. The same argument applies to debates about AI's energy demand and the volume of work that cheaper intelligence makes economic.
- Polanyi, in 1966, observed "we know more than we can tell," which is why the knowledge most worth automating, **tacit knowledge**, is the hardest to hand over.
- Baumol, in 1967, showed that when machines make some work cheap, the work they cannot do gets dear. (Consider the collapse of entry level software development jobs and the soaring, sometimes eight-figure salaries, for senior architects as the canary in this coal mine.)
- Harvey, in 1974, described a group agreeing to a plan not one of them wanted — the ***Abilene Paradox*** — which should be read aloud before every AI steering committee or board meeting.
- Goodhart, in 1975, warned that a measure that becomes a target stops being a good measure. (If AI usage becomes an incentivized target, the quality of that use will fall.)
- Collingridge, in 1980, described a dilemma: early in a technology's development, its effects are hard to foresee; once it is entrenched, changing its course becomes difficult. (A quiver, to be sure, in the Doomer's bow.)
- Bainbridge, in 1983, identified an irony of automation: people lose practice operating a system, yet must take over when it encounters problems it cannot handle. (Who today remembers addresses and phone numbers, or could follow complex driving instructions or navigate their way around new territory with a paper map?)

- Solow, in 1987, observed that the computer age was visible everywhere except in the productivity statistics. (As is with enterprise AI: a tsunami of spending, and a dribble of ROI.)
- David, in 1990, examined the long delay between electrification and its broad productivity benefits. Factories needed new layouts and ways of working to make full use of electric motors.

These arguments have long bedeviled strategists but were written before anyone had spent eight figures on model licences. They help explain why technical advances can coexist with disappointing returns.

# THE NINE NEW PARADOXES OF AI

*"The more powerful technology becomes, paradoxically, the more important the human becomes." (From Impact, 2019, Gibbons)*

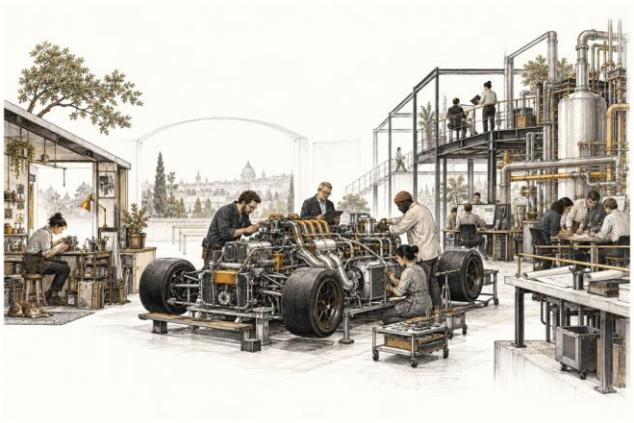
The nine new paradoxes in enterprise AI draw on that intellectual history, social science research, and on emerging evidence. These are signals from the frontier: propositions to test against the work organisations are doing now. Some extend familiar economic arguments; others connect observations that are usually considered separately.

As we will see, AI makes some kinds of intelligence abundant while the capacity to use it well remains scarce. A founder can acquire capabilities once reserved for large firms; an enterprise can spend heavily and struggle to turn individual gains into useful work. These tensions have consequences for what leaders buy, how they redesign work, and whose capabilities they develop.

None of the new paradoxes is yet definitive; the frontier is moving too swiftly. **However, leaders do not have the luxury of awaiting clarity:** our decisions won't wait for tidy resolution of these paradoxes. So, beneath each paradox there is a strategic heuristic or guide: something a chief financial officer, a chief information officer, or a chief people officer can act upon.

# The nine new paradoxes of AI

Nine tensions for leaders at the frontier



01
## Racecar
The better the machine, the more the driver matters.

02
## Illusion
The first summit is not the mountain.

03
## Cheap Intelligence
Cognition got cheap. Getting value from it did not.

04
## Plug-and-Play
The product is standard. The work is particular.

05
## Participation
Participation is part of the build, not an afterthought.

06
## Temporal Shear
Keeping up can leave you further behind.

07
## Endowment
The richest clients have the least to gain.

08
## Outsider
The next significant contribution may come from someone without the expected credentials.

09
## Big Science
Universities helped create the field; they now rent access to its frontier.

THE PARADOXES OF AI · PAUL GIBBONS

# 1. The Racecar Paradox

The better the machine, the more the driver matters

*Enterprises are upgrading the engine faster than they are redesigning the system around it and developing the people to drive it.*




**EXPLANATION**

Put your author in an F1 Ferrari and it will be neither safe nor fast. Yet enterprises keep upgrading the engine while giving too little attention to the drivers. Make my Ferrari engine bigger, and it will get less safe, and no faster.

AI-enabled expertise now includes verifying outputs, supervising delegated work, and recognising failure. Those abilities need practice, even among people with deep knowledge of their field.

For many of us, building agents, designing multi-agent control planes, and interpreting their telemetry are newly acquired skills. We are learning as we go, while models and harnesses evolve faster than many organisations can make ethical and productive use of them.

The University of Melbourne and KPMG's 2025 study included about 32,000 employees across 47 countries. **Fewer than half** reported receiving AI training. Among employees using AI at work, two-thirds reported

relying on output without evaluating it, while more than half reported mistakes in their work due to AI. They identify a practical exposure: access to powerful tools can arrive well before the habits and skills needed to use them dependably.

The Racecar also needs a racetrack. Dirt roads and big engines don't mix. AI amplifies the GIGO rule – data provenance and security might be thought of as the new track. To risk overworking the analogy, unscripted pathways take users (the drivers) places where they should not go. These side-of-track crash bars are the upgraded governance, safety, and policy features of today's AI landscape.

The Cheap Intelligence Paradox (*vid infra*) explains the economic logic of investing in what remains scarce. Here the question is what people must learn and how work must be redesigned as the machine becomes more capable.

### WHAT IT COSTS

A firm can buy models, infrastructure, licences, and agentic tools in a quarter. It cannot acquire sound judgement, redesigned work, supervision capability, and trustworthy use at the same speed. The visible cost is the licence bill. The larger cost is confident but unverified work: rework, errors, escalation, and decisions made on output nobody checked.

### WHAT TO DO

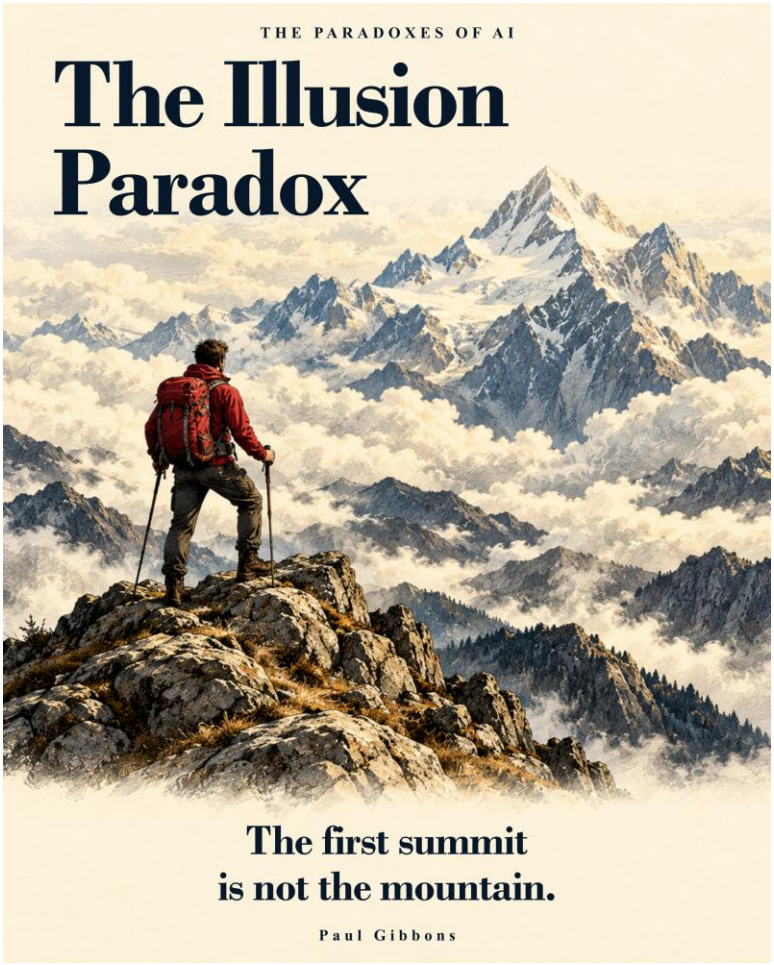
Do not approve model and infrastructure spend without a matching commitment to systems redesign and role-specific human capability. Fund practice, supervision standards, clear policies, operational measures, and time for the people who do the work to redesign it. Treat this as the deployment, not change-management overhead. Rewrite the definition of expertise used in hiring and promotion: in an AI-enabled organisation, expertise includes knowing how to supervise the machine.

**Invest in drivers; invest in tracks.**

# 2. The Illusion Paradox

The first summit is not the mountain.

*Early, real gains can make the larger frontier hard to see.*




**EXPLANATION**

*"Pick unreasonable goals and tell the model that they're very reasonable." (Felix Rieseberg, Anthropic)*

At four, my son told me, "I know how to play chess, Daddy." In one sense, he did. He could move the pieces and enjoy a game. How could he know the depth of mastery still beyond his sight? He could not imagine what a high-school chess champion thinks about, let along Magnus Carlson.

AI creates a similar predicament. Even in late 2026, 95 percent of ChatGPT subscribers were on the unpaid tiers. On free subscriptions, you can accomplish very little – either not very ambitious, or not for very long. But such users think, "I use AI."

At work, where the vast majority use AI at least a little, someone can get real value from Claude through quicker drafts, useful advice, and conclude that they **"know AI."** What remains invisible is how much further skilled use can go through better context, verification, workflow design, and delegation. Like my four-year-old son, they know nothing of what is possible at the frontier.

My senior executive friends who do rudimentary research and drafting think, “I use AI.” My friends who ask GPT for travel advice “use AI.” However, none have had the visceral wonder I had six months ago when I handed Claude Code a book I wrote and asked it to build a website for it. By the time I had grabbed a cup of coffee, a gorgeous website that I would have paid $15k for was waiting. Two dollars in tokens. When I asked a frontier AI (Astra) to review 100 academic sources and 20 hours of interview transcripts and saw it spin up 15 agents to complete the task, my jaw dropped. The first time my Chief of Staff dashboard went live, with email, scheduling, meeting prep-research, news curation, and a punch list of calls to make – I felt as if I had crossed a frontier. For non-technical individuals, agentic work was not possible in January 2026; among AI power users, my jaw-drops from January are, today, quite modest uses.

Most workplaces, for complex reasons, are not set up for that “what if I tried this” kind of experimentation. People are busy. Work habits have been grooved over decades. I’m well aware that using voice, not typing, to command AI is much faster and much more creative, but fifty years of typing into computers (after I gave up punch cards) is a hard habit to break! Even with the best will in the world, to change habits, organisational systems may push back, for good and bad reasons.

Self-report surveys supply **some** useful signals about trust, enthusiasm, friction, and willingness to experiment. But people saying they are willing to experiment does not mean that they do. Leaders should take those signals seriously and build on them. Establishing organisational value requires additional behavioural evidence: changes in cycle time, quality, rework, customer resolution, or the economics of the work. Measuring those outcomes should help users improve their practice, without treating their enthusiasm as something to be debunked.

In contrast, my journey with AI has sometimes made me want to throw my PC out of the window. Setting up OpenClaw in February took a day. When I do coding and work inside my computer operating system, I sometimes spend hours pasting red error messages from the terminal into an LLM to figure out what broke. When I built a RAG system for my own 35 years of consulting knowledge, by an order of magnitude, the hardest thing was peering into nearly 4 decades of Windows folders and finding value amidst the piles of detritus.

## WHAT IT COSTS

This is why I insist that leaders build. Until the power of the technology is viscerally experienced, and leaders have likewise experienced the enormous frictions of trying to reorchestrate even simple workflows, they just don’t “get it.” They are asking followers to do something that they (barely) understand in the abstract, and not at all practically.

Leaders can mistake early fluency for mature capability, or enthusiasm for return. In both cases, they lose the chance to build on what users have discovered and to make the next summit visible.

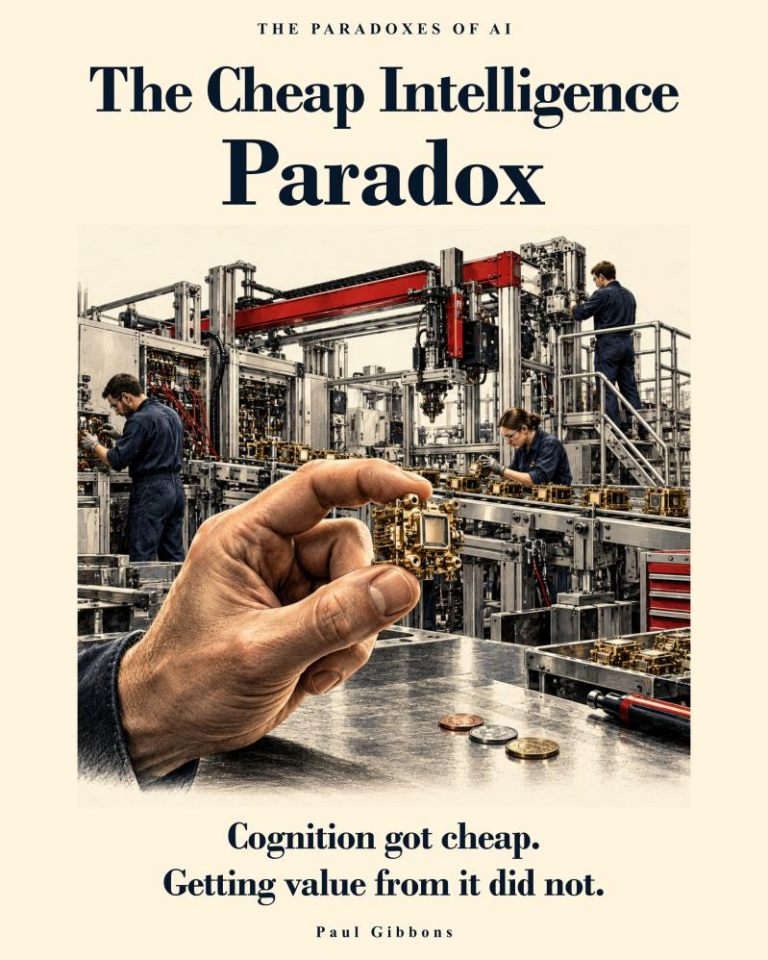
## WHAT TO DO

Don’t rely exclusively on self-reports of use and value. Use self-report to find enthusiasm, friction, and emerging practice; then pair it with behavioural measures that show what changed in the work. Give people ways to see the next mountain: shared practice, examples of excellent use, peer comparison, deliberate skill levels, and operational evidence. Have them share their “wows” and frustrations. Elevate the people in your business, who get their “use-cases” from YouTube, or Substack (where the best “builder” content is,) and give them a sandbox to play in.

# 3. The Cheap Intelligence Paradox

Cognition got cheap. Getting value from it did not.

*The price of a unit of machine intelligence has collapsed. Turning it into valuable work still requires substantial investment in systems and people.*




**EXPLANATION**

Stanford's AI Index shows the price of GPT-3.5-level intelligence falling from twenty dollars to seven cents per million tokens between November 2022 and October 2024 - a 280-fold fall over that period. Imagine, if you will, gasoline at a few cents a gallon.

Lower unit prices do not make intelligence free. They make a greater number of uses economic, so demand for compute and total spending can still rise: the Jevons effect introduced earlier. **A cheaper token does not bring an equivalent reduction in the cost of redesigning a process, developing a workforce, or clarifying who makes a decision.**

Economists call inputs that increase one another's usefulness *complements*. AI and the human capability to deploy it can work this way: cheaper intelligence increases the opportunity, while systems redesign and skilled use determine how much of it becomes valuable. This helps answer the leader's question: after all this

investment, why are we not seeing more value? The organisation around the model can remain the limiting factor.

A modern formulation of this longstanding pattern is Brynjolfsson, Rock, and Syverson's productivity J-curve: businesses invest in new processes and human capital before the resulting gains become visible. Weak early returns can therefore reflect investment still taking effect, as well as poor deployment. Leaders need to distinguish those explanations before buying more technology or abandoning the effort. As the poker saying goes, "know when to hold 'em, and know when to fold 'em."

**Tokens are cheap inputs within an expensive operating system.** Inference at scale, integration, governance, systems redesign, and human capability all carry real cost. With today's frontier models so powerful, the rule of thumb I offer leaders is value is about 10% model power, and about 90% "the rest." (Moreover, as the Racecar Paradox shows, human capability development may be a substantial portion of that 90%.)

Paul David's account of electrification shows why cheaper power alone was insufficient. Factories also had to change how production was organised. AI can help people redesign work; responsibility for choosing, testing, and governing the new system remains human.

The cheap intelligence paradox is an old economic story, made more urgent by AI.

## WHAT IT COSTS

The false economy lies in buying what has become cheap and visible while underinvesting in what still limits value. The easy purchase crowds out the harder work: systems redesign, human capability, and leadership.

That underinvestment means the rest of the system changes too slowly. The bill reappears as rework, weak adoption, poorly supervised output, and a growing gap between technical capability and economic return.

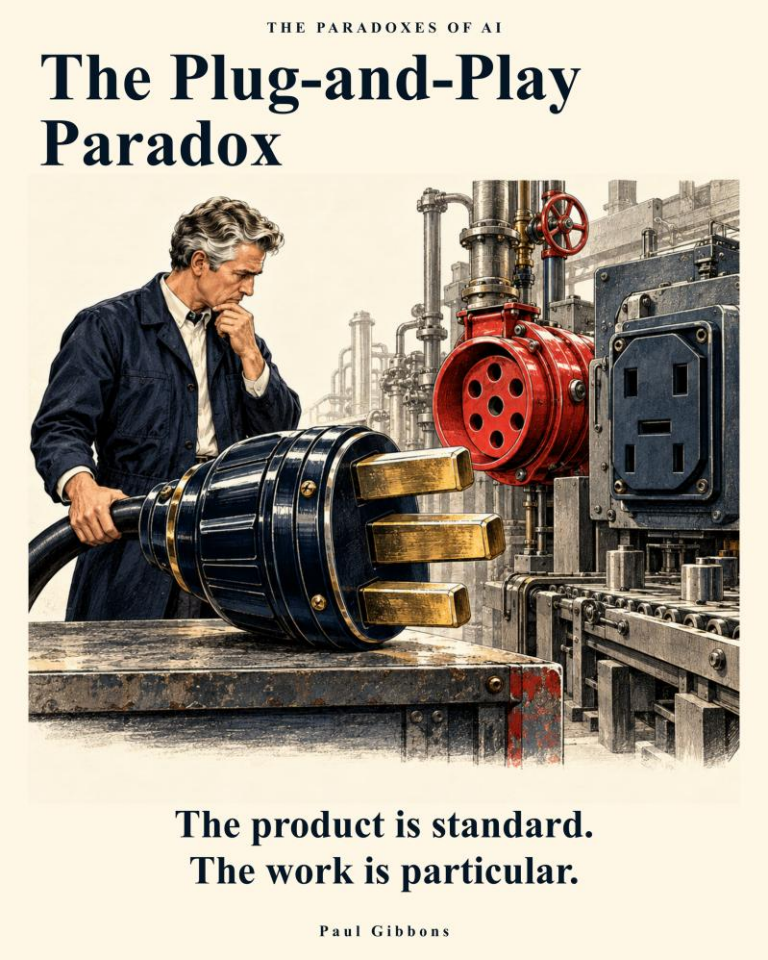
## WHAT TO DO

Spend on what limits value: systems redesign, human capability, and leadership. Approve an AI proposal when it changes the system in which cheaper thinking becomes useful work.

# 4. The Plug-and-Play Paradox

The product is standard. The work is particular.

*The same model can serve thousands of organisations. A dependable workflow must fit the organisation using it.*




**EXPLANATION**

My consulting brethren and startup agencies advertise plug-and-play workflow solutions. The promise confuses a product they can supply with an organisational transformation they cannot deliver for you. The benchmarking, best-practice, and road maps forget the most important thing.

Take two companies in the same industry, of the same size, in the same country. Look inside: different technology stacks, staffing, capabilities, objectives, decision rights, and histories of what has been tried and failed. Even the exceptions differ. Who can approve a discount? Which record is authoritative? What happens when a customer's account contradicts it?

**Nobody can sell you the finished workflow, because that idiosyncratic workflow is the part that's yours.**

This is an expensive lesson to relearn and enterprises have been relearning it since the 1990s. Some stuff you can buy, the important stuff you have to build.

In his account of electrification, economist Paul David explained why installing electric motors was insufficient to realise their productive potential. Factories organised around shafts and belts needed layouts and working arrangements that exploited distributed power. Buying a new engine left much of the old production system intact.

Packaged AI tools can be immediately useful to individuals: writing assistance, presentation generation, and working with sources. An agentic process that acts across company systems, handles exceptions, respects permissions, and leaves someone accountable is a different undertaking. The model's capabilities may be available out of the box. Reliable operation requires design and testing *in situ*.

McKinsey's 2026 State of AI survey found that nearly three-quarters of respondents at AI high-performing organisations reported fundamentally redesigning workflows, compared with one-quarter of other respondents. High performers reported significant value from AI and attributed at least 5 per cent of their organisation's earnings before interest and tax to its use. The survey cannot establish causation, but it gives weight to the distinction between buying access and changing how work gets done.

There is a further implication for orthodox change management. **With packaged enterprise software, a predefined process could go live even where worker participation was perfunctory.** Where an AI-enabled workflow has yet to be designed, participation becomes part of creating the thing to be adopted, not something added on as an afterthought, to get "buy-in." The people doing the work know the exceptions, dependencies, and informal judgements that a demo rarely reveals. You cannot roll out a workflow that hasn't yet been imagined, let alone built.

## WHAT IT COSTS

A licence buys access. The deployment budget must also cover integration, local design, testing, supervision, and maintenance. Omit those costs and they return as repeated pilots, consultant dependency, manual workarounds, and staff repairing a process that never fitted. The next software purchase can look attractive precisely because the difficult work remains unfunded.

## WHAT TO DO

Buy useful software and reusable components, but budget explicitly and sufficiently for making workflows work here, IRL. Choose a consequential workflow and involve the people who perform it, including those receiving its outputs. Define exceptions, permissions, quality measures, and responsibility before extending the system.

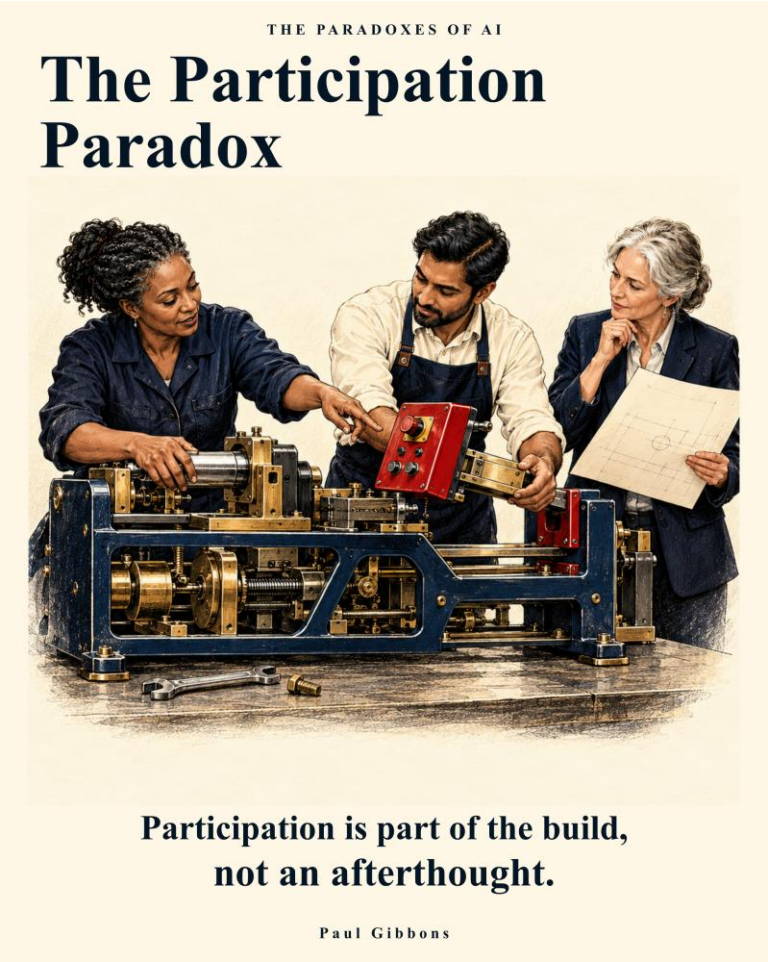
The Participation Paradox takes up participation in that design; the Racecar Paradox examines the capability people need to operate it. Better platforms can reduce the local effort. They cannot settle the organisation's decisions about responsibility, acceptable error, and coordination.

**You can buy the components. You still have to design how the work gets done.**

# 5. The Participation Paradox

Participation is part of the build, not an afterthought

*Participation moves from a recommendation that leaders could sideline to a practical requirement for designing work that people can use, trust, and improve.*




**EXPLANATION**

Involving users in technology strategy and workflow design has long been recommended, with roots in organisation development and sociotechnical research. However, it was always more discussed by academics than done in practice: Organisations treated it as optional. A workshop held, a town hall delivered, a communications plan completed: **participation theatre** could tick the “user engagement” box on the change plan without seriously consulting anyone on anything.

And businesses, for a long time, got away with such performative engagement rituals. A largely prescribed SAP or Salesforce process could be mandated even when users disliked it. Those systems also required local design, but compliance with an established workflow could keep the operation running. When a team is redesigning judgement-heavy work around AI, important questions remain open: who decides, which exceptions matter, what needs checking, and how the work reaches the next department.

Practitioners bring knowledge that helps answer those questions. Their involvement becomes part of creating the system they are expected to adopt.

I coined People-first AI ™ as a humanistic and ethical commitment. I now see it as a practical change tactic too. Trust, upskilling, and engagement belong inside the build. People develop skills while solving real problems, discover where the system deserves their trust, and have a say in decisions that affect their work. This requires time and influence, including the ability to challenge a proposed design. It begins before leaders "cut a cheque for the tech."

The IKEA effect, from cognitive science, suggests that an eighty-dollar table can acquire an unreasonable place in your affections because you have assembled it yourself. Successful effort can increase how much we value what we have made. In a workflow, participation can also improve the fit: practitioners know the workarounds, handovers, and awkward cases that a process diagram misses. Their attachment is useful; whether the design works still needs testing.

### WHAT IT COSTS

**When participation is postponed, the organisation pays for design decisions before discovering whether they fit the work**. Staff then invent workarounds, engineers revisit requirements, and a change team is asked to repair trust around choices people had no part in making. Money saved by excluding practitioners from the build can return as rework and a system people use reluctantly or avoid. Participation needs a budget while consequential decisions are still open.

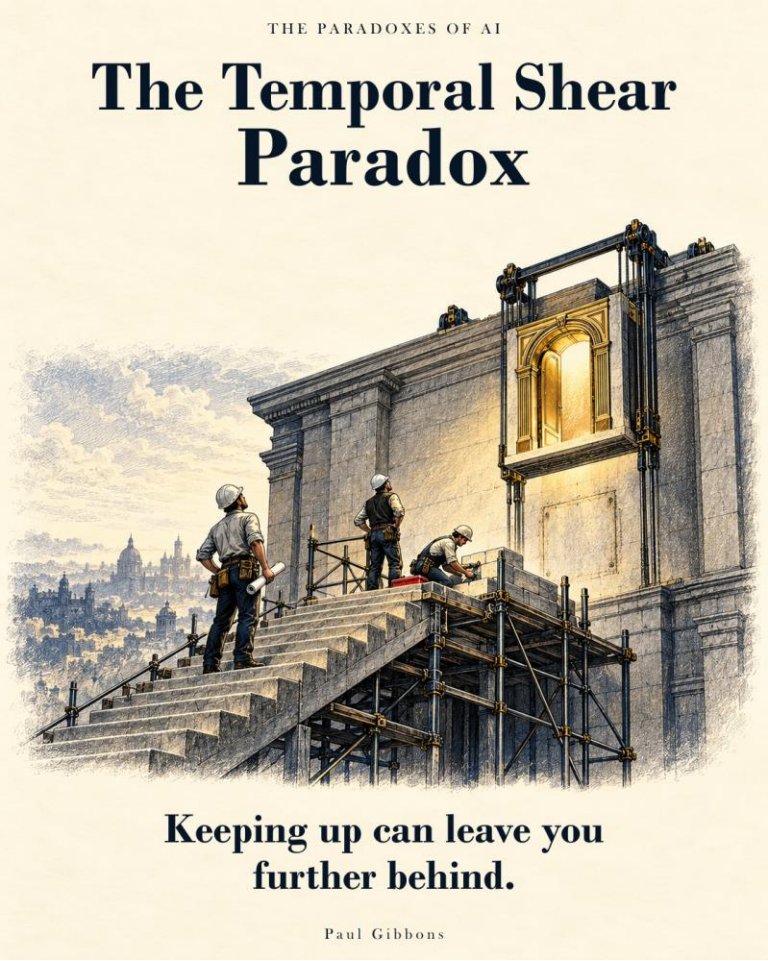
### WHAT TO DO

Build Adoption Pods ™ around specific workflows. Bring together the practitioners doing the work, the person accountable for its outcome, a forward-deployed engineer, and a forward-deployed change practitioner. Give the pod time and authority to design, test, and improve the process. Change expertise stays inside the team throughout the build, helping people develop capability, address concerns, and establish dependable working practices. Include those who receive the workflow's outputs. Assess the pod through improvements in the work and sustained, appropriate use. Participation belongs in the delivery budget from the outset.

# 6. The Temporal Shear Paradox

Keeping up can leave you further behind

*The work needed to make today's capability useful can become an obstacle to using tomorrow's.*




**EXPLANATION**

I can subscribe to a new AI tool with two clicks. This means my home "AI-lab" for knowledge work is within weeks of the frontier: OpenClaw in February; new models upon arrival (open source and otherwise); "lab models" (Gemini, OpenAI, and Anthropic) that upgrade to the latest versions automatically; and new GitHub plug-ins as they arrive (daily).

But the subscriptions accumulate faster than the new working practices. Making those useful in my working life takes considerably longer – **even at a boutique advisory firm, integrating new tools is a full-time job**. I can keep acquiring better technology without acceleration and augmenting my outputs as I had hoped. When commercial demand spikes, investigating new tools and upgrading the "home-lab" falls behind.

This is acute in big enterprises.

Think of five clocks: model releases measured in weeks, agentic advances in months, budgeting in quarters, infrastructure replacement in years, and institutional change sometimes in decades. These are illustrative

timescales, but the difficulty is concrete: decisions made on one clock must work with commitments made on another.

Enterprise budgeting and architecture review restrain indiscriminate buying. They cannot guarantee that the assumptions behind an approved investment will remain useful throughout implementation. A firm may integrate a capability, validate its behaviour, and train people to use it, only to find that a new model makes some of that work unnecessary. The original decision need not have been foolish.

Anthropic's engineers describe adding **context resets** to compensate for one model's tendency to finish tasks prematurely as its context filled. With a stronger model, that behaviour disappeared, leaving the workaround as overhead. Technical progress had changed what the surrounding system needed to do.

In my Leadership Delta ™ framework, I call the stress between these clocks temporal shear: connected organisational plates moving at different velocities. Here, the difficulty is that the work needed to integrate one generation of capability may still be under way when the next changes what is worth integrating.

Modular architecture can reduce the cost of change. Replacing a component, however, does not establish that the whole workflow will behave reliably. If each new use case also waits for the same data team to change shared pipelines, the pace of the model matters less than that dependency. A stable system can remain valuable for years. A newly purchased one can already be poorly suited to the work.

### WHAT IT COSTS

Repeated switching can consume the time needed to develop useful practice. Integration work is abandoned, training is revised, and subscriptions survive the uses that justified them. Conversely, holding every decision until the slowest system is ready leaves useful improvements waiting. The organisation can pay both for moving too quickly and for committing too rigidly.

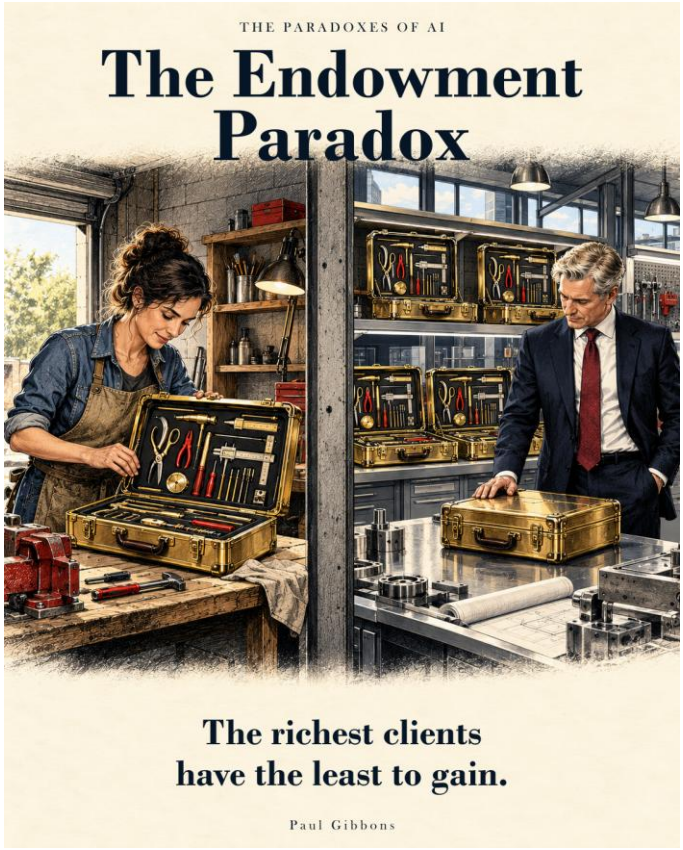
### WHAT TO DO

Give different decisions appropriate lifetimes. Keep technical components replaceable, and test changes against the work they must perform. Preserve useful operating practices long enough for people to become capable, while revisiting assumptions that newer models may have overtaken. Review subscriptions and workarounds as deliberately as new purchases. Approve a migration when its expected benefit justifies the learning, testing, and integration it requires. Coordination across these clocks is a continuing leadership responsibility.

# 7. The Endowment Paradox

The richest clients have the least to gain

*A modest AI budget can add more capability to a resource-poor founder than a large one adds to an already well-equipped firm.*




**EXPLANATION**

Connections, capital, and expertise have long shaped who could turn an idea into a business. AI lowers some of those barriers. A curious tinkerer can spin up working code they could not previously write. A founder can draft a marketing campaign, explore financial scenarios, or organise company knowledge without employing a specialist in each field. The gain includes being able to attempt work that was previously out of reach. (In AI speak, this is capability AI versus efficiency AI.)

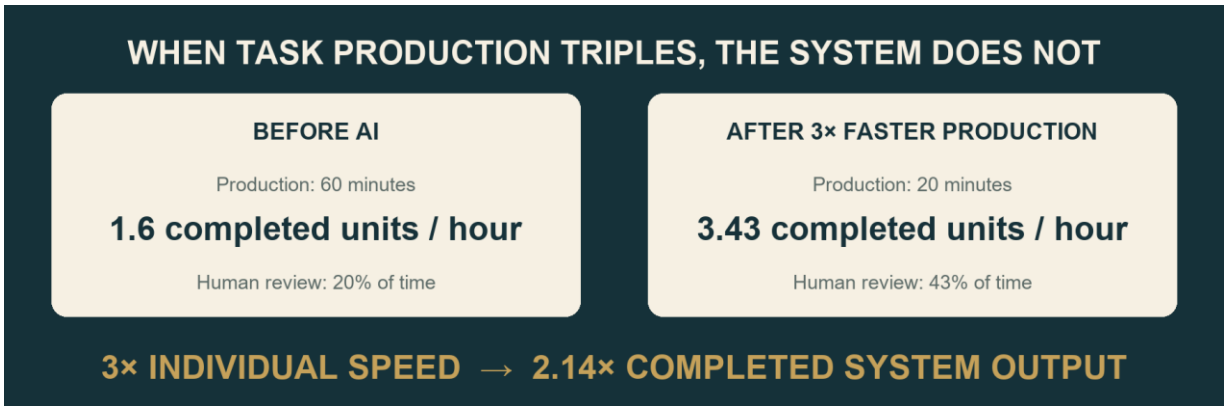
There are two effects to distinguish. An individual can improve within a role; a small business can acquire a range of capabilities it previously lacked. Supplying eighty percent of a missing capability can be worth vastly more than adding another source of one already abundant; or, as an economist would put it, diminishing marginal returns may favour filling a capability gap over buying a close substitute for expertise already available. That depends on the task. AI can also complement experts, and large firms retain advantages in proprietary data, distribution, and the scale over which they can spread costs.

Complex systems introduce another limit**. Producing more candidate drug molecules does not remove the need for testing and regulatory approval**. Producing more content does not create more audience attention or accelerate a company's approval process. An individual may work much faster while the wider system continues at its former pace. A founder who controls the whole process can often act on a useful

output immediately; a practitioner in a large firm may have several departments between the output and its use.

Consider a simple two-person system. Each person creates work and checks a quarter of the other's output. At human speed, a completed unit entails one hour of production and fifteen minutes of review: together, they complete 1.6 units an hour.

Now make each work three times faster, while serious human checking still takes fifteen minutes per hour of output. The team completes 3.43 units an hour, not 4.8: a threefold improvement in individual production becomes only a 2.14-fold improvement in system output.




The constraint has moved. Review, judgement, exception-handling, and coordination become the limiting work. **Before the speed-up, checking consumes a fifth of each person's time; afterwards it consumes almost half.** (I don't know about you, dear reader, but I do not enjoy that part of work at all.)

The people are not merely working faster: they are supervising a much faster flow of consequential work, with less slack to notice what has gone wrong. The verification tax rises as the machine produces more work than a human can responsibly supervise. AI also creates an infinite backlog: more projects, analyses, and improvements become plausible than an organisation can prioritise, coordinate, and carry through.

**Frontier AI does not remove the bottleneck; it shifts it elsewhere in the production system.**

### WHAT IT COSTS

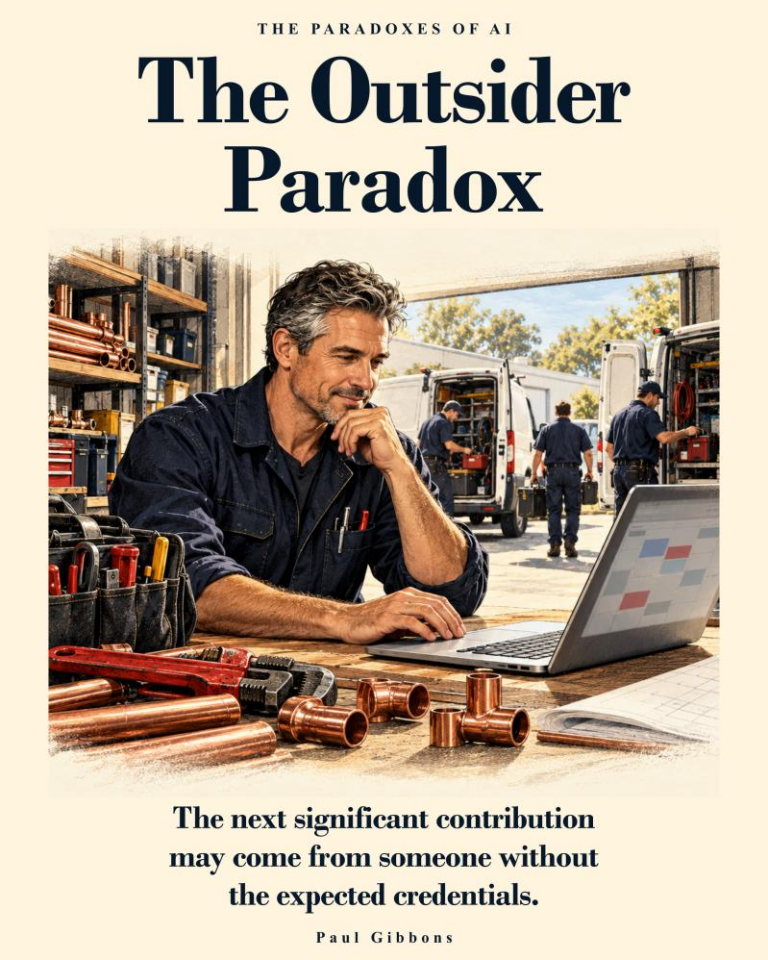
A low-cost subscription can make a striking difference to a founder while a large enterprise programme produces modest bang per buck. This is a proposition about missing capabilities and coordination costs, not an established ranking of firms by size. The expensive mistake is assuming that faster individual tasks will add up to enterprise returns without changing the system through which the work must pass.

### WHAT TO DO

Identify where AI adds a capability you lack and where it improves one you already have. Give a small team responsibility for an outcome from beginning to end, with access to the people and information it needs. Measure the full cost and value of the work, including review and coordination. **The question for a large firm is how much of the freedom and immediacy enjoyed by a capable founder it can reproduce safely inside its own operations.**

# 8. The Outsider Paradox

The next significant contribution may come from someone without the expected credentials




**EXPLANATION**

In 2024, Demis Hassabis shared the Nobel Prize in Chemistry for a breakthrough in a fifty-year-old problem: predicting a protein's three-dimensional structure from its amino-acid sequence. A chess prodigy and games developer who trained in neuroscience had solved a decades-old problem in chemistry with AlphaFold2. Hassabis and his team showed how AI might power research breakthroughs by people outside the core of the field and solve one of its hardest problems.

I was deep in the protein folding project as an undergraduate in the 1970s; it seemed impossible even for dipeptides (two amino acids, not two hundred.)  Hassibis' win strikes me as astonishing, that a non-chemist could win the Nobel. But for non-chemists it might seem abstract – like AI solving Go. Whatever. Interesting to be sure, but non-Go players and non-chemists might struggle to find the "so-what."

The following example might be much more practical and understandable. I gave a keynote in Denver, and following me was a plumber who had dropped out of high school at fifteen. After a stint in the military, he had opened an HVAC business that, in under eighteen months, had grown to 44 technicians. He had built his entire operational infrastructure in Claude Code.

My jaw dropped. Until that conversation, I had scarcely considered the complexity beneath a plumbing business. Thousands of possible parts. Different equipment models that require the correct component. Jobs to schedule, customers to coordinate, invoices to issue, and 44 people whose work had to fit together. Sending someone with the wrong part wastes both time and money. By his account, software he had built with Claude Code handled all of this coordination, allowing the business to operate with a remarkably lean administrative team.

AI had helped someone who had never studied algebra, never mind coding, build a complex operations layer to his business.

These examples raise an exhilarating question. What could a curious 21-year-old in Kigali contribute to economics, physics, or mathematics with access to powerful tools and ways to test their ideas? What could an aspiring entrepreneur in Kuala Lumpur build without an MBA or the money to assemble a specialist team? Citizen science and entrepreneurship could draw on a much wider pool of talent.

This is another dimension of People-first AI ™: opportunity at the bottom of the economic pyramid. People with limited resources may gain the means to attempt difficult work, create value, and build wealth for their families and communities. The possibility is that fewer good ideas will fail simply because their authors lack the expected route into the field.

### WHAT IT COSTS

Institutions can overlook valuable contributions when credentials become a substitute for examining the work. The loss includes businesses never attempted, useful ideas never tested, and capable people never considered. For those people, the cost is opportunity denied. For an organisation or economy, it is productive talent left unused.

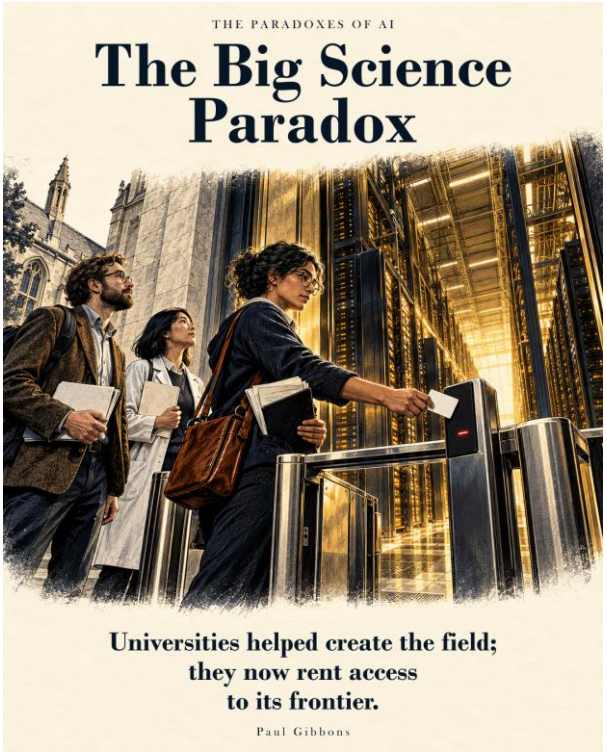
### WHAT TO DO

Create routes through which people can demonstrate capability: working prototypes, practical assessments, and research contributions that others can examine and test. Pair access to AI with domain mentorship and rigorous verification. Give under-resourced people opportunities to build, alongside the support needed to recognise mistakes. Keep the standards required for safe and dependable work, while widening the routes through which people can meet them.

# 9. The Big Science Paradox

Universities helped create the field; they now rent access to its frontier

*Frontier AI has outgrown university-scale resources. Private firms control much of the infrastructure on which academic researchers now depend.*




**EXPLANATION**

Universities helped create the field; they now rent access to the infrastructure needed to pursue its frontier. Modern AI developed in universities, but the resources required to build the largest models have shifted the balance towards industry. While academic researchers can still make important discoveries, driving frontier model development requires a scale of compute, data, and specialist staffing that individual universities cannot readily assemble.

Nur Ahmed, Muntasir Wahed, and Neil Thompson documented this shift in *Science* in 2023. Industry's growing command of compute, data, and talent was translating into greater influence over research outputs. Stanford's 2025 AI Index shows the distinction: industry produced nearly 90 per cent of notable models in 2024, while academia remained the leading institutional producer of highly cited research. Contributing to a field and commanding the resources to drive its frontier are different things.

Big science has outgrown university budgets before. CERN and the James Webb Space Telescope rest on public institutions pooling resources for research beyond the means of an individual university. Frontier AI has followed a different institutional path, with commercial firms taking the lead. In my view, their ability to mobilise capital quickly, encouraged by enormous prospective commercial returns, helps explain that outcome. A comparable public effort requires governments to assemble both the funding and the political will. It is hard to see a public consortium, let alone a single university, spending $10 billion on infrastructure and hiring researchers at eight-figure salaries.

Universities have, as a result, become customers of the firms driving the frontier, purchasing cloud compute or securing access through corporate grants. Harvard offers consolidated AWS billing for research groups; Stanford HAI distributes Google Cloud research credits. These arrangements enable valuable work, but they also reveal the dependency. The companies supplying the research infrastructure may also own the models under investigation. Their decisions about access influence which independent experiments are possible.

## WHAT IT COSTS

This has consequences beyond university budgets. **Control of research infrastructure influences whose questions receive sustained attention and resources.** Commercial demand rewards problems with paying customers; social value and purchasing power do not reliably coincide. A problem in economic development may matter enormously to the people affected while offering little prospect of commercial return.

The frontier labs have made substantial commitments to science and education through grants, subsidised access, and research partnerships. These contributions deserve recognition. But access offered on terms set by a company is different from research capacity governed by a public mandate. Eligibility, resources, and continuity remain dependent on the provider's decisions.

The concern is therefore institutional, rather than a judgement about the generosity of particular leaders. How should societies secure sustained AI capacity for problems whose public value exceeds their commercial appeal? Public research infrastructure and independently governed funding can give those purposes a durable claim on resources. Without them, important parts of the public research agenda remain contingent on corporate priorities.

The cost can take the form of experiments forgone, questions left untested, and research priorities shaped by commercial access. For an enterprise buyer, the related risk is treating a provider's roadmap as sufficient independent advice. Academic affiliation alone cannot establish access to frontier systems; access alone cannot establish independence.

## WHAT TO DO

Assess technical capability and independent judgment separately. Examine a research partner's access to relevant systems and its commercial dependencies. Support arrangements that let independent researchers test consequential claims, including public research compute and partnerships with protected publication rights. For enterprise decisions, combine hands-on technical knowledge with scrutiny that is independent of the supplier.

## POST-SCRIPT: THE NATIONAL IMPLICATION

A country need not produce the leading AI model to benefit substantially from AI. Building domestic capability may serve commercial, scientific, and security purposes. Turning available intelligence into better businesses and public services requires another programme of work. Governments should assess both ambitions explicitly. A successful technology industry cannot stand in for productive use across the economy.

For leaders, the common question is where additional intelligence changes the work, and what people need in order to make that change dependable. The more powerful technology becomes, the more consequential those human choices become.

# What to do Monday

**Before approving the next AI spend, insist on these nine answers. A proposal that cannot answer them is not ready.**

## NINE RULES FOR THE NEXT INVESTMENT DECISION

**01 NO CAPABILITY PLAN, NO MODEL SPEND**

Fund the people, controls, and operating system needed to supervise the machine.

**02 MEASURE OUTCOMES, NOT ENTHUSIASM**

Track cycle time, quality, rework, customer outcomes, and unit economics.

**03 FUND THE BOTTLENECK, NOT THE MODEL**

Cheap intelligence is useless inside an unchanged operating system.

**04 NO WORKFLOW OWNER, NO AUTOMATION**

Define responsibility, exceptions, testing, and maintenance before buying the licence.

**05 PUT PRACTITIONERS INSIDE THE BUILD**

Give the people who know the real work authority in its design.

**06 MAKE COMPONENTS REPLACEABLE**

Do not hard-wire today's model into a system that will outlive it.

**07 MEASURE THE WHOLE OUTCOME**

Faster tasks do not count if work still sits in queues or creates a verification tax.

**08 TEST CAPABILITY, NOT CREDENTIALS**

Look for demonstrated work, practical judgement, and rigorous verification.

**09 NEVER LET A SUPPLIER GRADE ITS OWN SYSTEM**

Separate technical access from independent judgement before placing the bet.

# Selected sources


### FOUNDATIONS

Ronald Coase, The Nature of the Firm (1937); W. S. Jevons, The Coal Question (1865); Michael Polanyi, The Tacit Dimension (1966); William Baumol, Macroeconomics of Unbalanced Growth (1967).

Jerry Harvey, The Abilene Paradox (1974); Charles Goodhart, Problems of Monetary Management (1975); David Collingridge, The Social Control of Technology (1980); Robert Solow, We'd Better Watch Out (1987).

Ted Chiang (1998), *Story of Your Life*, in *Starlight 2*, edited by Patrick Nielsen Hayden. Tor Books. Opening epigraph. Story

### BOOKS AND RELATED WORK

Paul Gibbons (2019a), *Impact: 21st Century Change Management, Behavioral Science, Digital Transformation, and the Future of Work*. Phronesis Media. Book details

Paul Gibbons (2019b), *The Science of Organizational Change: How Leaders Set Strategy, Change Behavior, and Create an Agile Culture*. Phronesis Media. Book details

Paul Gibbons and James Healy (2025), *Adopting AI: The People-First Approach*. Phronesis Media. Book details

Paul Gibbons and Tricia Kennedy (2023), *Change Myths: The Professional's Guide to Separating Sense from Nonsense*. Phronesis Media. Book details

Paul Gibbons (2026), *The Leadership Delta: Why AI Leadership Is Different and What To Do About It*. Discussion document, Paul Gibbons Advisory. Temporal Shear Paradox.

### RESEARCH AND EVIDENCE

Stanford HAI (2025), AI Index Report, chapter 1. Inference prices and the institutional distribution of model development. Cheap Intelligence and Big Science paradoxes. Report

Erik Brynjolfsson, Daniel Rock, and Chad Syverson (2021), The Productivity J-Curve: How Intangibles Complement General Purpose Technologies, American Economic Journal: Macroeconomics, 13(1), 333–372. Paul David (1990), The Dynamo and the Computer, American Economic Review, 80(2), 355–361. Cheap Intelligence Paradox. J-curve working paper

Erik Brynjolfsson, Danielle Li, and Lindsey Raymond (2025), Generative AI at Work, Quarterly Journal of Economics, 140(2), 889–942. Endowment Paradox. Published study

METR (2025), Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity; experiment-design update (24 February 2026). Illusion Paradox. Trial · Update

McKinsey & Company (25 August 2026), The State of AI in 2026: On the Road to ROI. Plug-and-Play Paradox. Report

Michael Norton, Daniel Mochon, and Dan Ariely (2012), The IKEA Effect: When Labor Leads to Love, Journal of Consumer Psychology, 22(3), 453–460; Tavistock Institute, Socio-Technical Systems. Participation Paradox. IKEA effect · Sociotechnical history

University of Melbourne and KPMG (2025), Trust, Attitudes and Use of AI: A Global Study; Lisanne Bainbridge (1983), Ironies of Automation, Automatica, 19(6), 775–779. Racecar Paradox. Global study · Automation

Nobel Foundation (2024), Chemistry prize announcement; UCL (9 October 2024), biography of Demis Hassabis; University of Chicago, John Jumper's chemistry training. Outsider Paradox. Nobel · UCL · Chicago

Nur Ahmed, Muntasir Wahed, and Neil Thompson (2023), The Growing Influence of Industry in AI Research, Science. Big Science Paradox. Author-institution account

Stanford HAI, Google Cloud Credit Grants; Harvard SEAS, Harvard Cloud Program; CERN, Member States; NASA, Webb Team. Institutional examples in the Big Science Paradox. Stanford · Harvard · CERN · Webb

Stewart Brand (2018), Pace Layering: How Complex Systems Learn and Keep Learning, adapted from The Clock of the Long Now (1999), with antecedents in How Buildings Learn (1994). Temporal Shear Paradox. Essay

Anthropic (8 April 2026), Scaling Managed Agents: Decoupling the Brain from the Hands. Temporal Shear Paradox. Engineering account

Martin Fowler (2015), Microservice Trade-Offs; Zhamak Dehghani (2019), How to Move Beyond a Monolithic Data Lake to a Distributed Data Mesh. Temporal Shear Paradox. Microservices · Data platforms

Microsoft (2025), The Year the Frontier Firm Is Born. Introduction. Philip Lane, ECB (23 March 2026), AI and the Euro Area Economy. Conclusion. Microsoft · ECB

Felix Rieseberg (13 October 2026), *How to Be More Ambitious: A Tutorial for Humans with Models*. Illusion Paradox. Essay