

BEYOND THE CAPABILITY REFLEX

ChatGPT Changed the Question — Why AI Strategy Cannot Start With Technology

Why the organisations demanding an AI strategy in 2023 are repeating the mistake they made with digital — and what starting from the decision would change

Giovanni Leonardi

April 2023



ChatGPT Changed the Question — Why AI Strategy Cannot Start With Technology

Transformation · No. 378

Originally written April 2023 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“When a tool can do almost anything, the scarce discipline is deciding what it should do here.”

1. The meeting that keeps happening

It is the spring of 2023, and the same meeting is taking place in boardrooms everywhere. Someone — often the chief executive, sometimes a non-executive who tried the thing over the weekend — puts a laptop on the table and shows the room a demonstration. A few typed sentences, and a machine writes a passable draft of a customer email, a policy summary, a fragment of working code. There is a short silence. Then the question arrives, and it is always the same question: what is our AI strategy?

By the following meeting there is a workstream. By the one after that, a budget line, a shortlist of vendors, and a senior sponsor with "AI" newly appended to their title. Motion arrives quickly. What has not arrived — and what, many months later, will still be conspicuously absent — is an answer to a smaller and far more awkward question: which decision, exactly, are we trying to make better?

I want to argue something that, in the enthusiasm of this moment, will sound almost perverse: an AI strategy that begins with the technology is not a strategy at all, and the organisations now moving fastest to acquire the capability are quietly repeating a mistake they have made at least twice before.

2. We have seen this shape before

Strip away the novelty and the pattern is familiar. A general-purpose technology becomes suddenly visible. It is genuinely impressive. Leaders feel, correctly, that it matters — and conclude, incorrectly, that mattering is the same as knowing what to do about it. And so "strategy" quietly collapses into "adoption": the strategic question — where should we compete, and how should we win — is displaced by the logistical one — how do we get this technology into the building.

A decade ago we called it digital. The digital strategies of that era read, more often than not, like technology shopping lists: a portal here, a mobile app there, a data lake because everyone else was digging one. A few years later we did it again with cloud, where a great many "cloud strategies" turned out, on inspection, to be data-centre exit plans wearing a more fashionable name. In each case the technology was real and the value was real.

What was missing was the connective tissue between them — a clear account of which business problem the capability was meant to move, and how anyone would know if it had.

The reliable tell of a technology-led strategy is that you can describe everything you are going to buy and build, and nothing you are going to decide differently.

Generative AI is a more dramatic case, but it is the same case. That is not a reason to be cynical about it. It is a reason to bring to it the one discipline the previous rounds lacked.

3. Why the reflex is so hard to resist

If starting with the technology were obviously foolish, we would have stopped by now. We keep doing it because the technology-first move is, in the moment, deeply rational for the people making it.

- It is legible. A capability programme produces things you can point to — a platform stood up, licences procured, a centre of excellence launched. Problem-definition produces, at first, only harder conversations.
- It is safe. When every competitor is announcing an initiative, being seen to move is its own protection. The board did not ask whether you had improved a decision; it asked what your AI strategy was, and a workstream is a serviceable answer.
- It is sold to you that way. The entire vendor ecosystem is organised around capability, not around your problems. Its natural sentence is "here is what our technology can do," and it is a short step from there to a roadmap that is really a purchase order.
- It flatters our fear of being left behind. When a technology moves this fast, waiting to define the problem feels like standing still while others sprint. Speed becomes a virtue in itself, detached from any account of where we are meant to be going.

None of these forces is stupid. Together they produce a predictable result: a great deal of activity, arranged around the technology, waiting for a problem to justify it.

4. What the reflex actually costs

Here is the mechanism, and it is worth being precise about, because the cost is not obvious until you look for it.

A capability that is not attached to a specific decision has no owner accountable for an outcome, no baseline it is meant to improve, and therefore no honest way to succeed or to fail. It cannot fail, because no one agreed in advance what success would look like. It cannot really succeed, because there is nothing to measure it against. So it does the only thing left to it: it persists. It becomes the pilot that never quite ends, the proof-of-concept that is forever nearly ready, the impressive demonstration that never crosses into the work.

I have watched this play out, and the details are always more mundane than the promise. A mid-sized insurer, a few years into the digital era, launches a "digital lab" with real fanfare — a backlog of a dozen ideas, a respectable budget, a wall of coloured notes. Two years and several million pounds later, exactly one thing has reached customers: an app that lets them photograph a document instead of posting it. Genuinely useful; it trims some handling time in the call centre. Also a rounding error against what was spent, and a quiet embarrassment against the ambition that launched it. The failure was not that the technology could not do more. It was that not one of those dozen ideas had begun life as a decision someone owned and wanted to make better. They had begun as things the technology could do — and a thing the technology can do, with no one on the hook for a result, drifts.

"A hundred ungoverned experiments with no thesis about where the value sits is not a strategy. It is an expense with good production values."

That story is being written again right now, in the vocabulary of generative AI rather than of mobile and apps. I would rather we did not need it a third time. The waste is seldom the licence fee. It is the year of senior attention, the credibility spent, and the strategic question that went unasked while everyone was busy adopting.

5. The question first, the technology second

The alternative is not more caution. It is a different order of operations.

Begin with the organisation's decisions, not its technology. Which decisions do we make badly, slowly, or inconsistently? Where does judgement pile up behind a small number of scarce people? Which pieces of work are expensive precisely because they are language — drafting, reviewing, reconciling, explaining — and therefore, for the first time, plausibly cheaper? Only once you can name the decision does the technology question become answerable at all: does this class of tool change the economics of this decision by enough to be worth the disruption of adopting it?

1. Name the decision or the work, and who owns it today.
2. Establish, however roughly, what it currently costs — in time, in error, in delay.
3. Then, and only then, ask whether the technology moves that number, and by enough to matter.

This is slower to start and far faster to finish, because everything downstream — the pilot, the metric, the sponsor, the go/no-go — now has something real to attach to. When a tool can do almost anything, the scarce discipline is deciding what it should do here. Problem-definition is not bureaucratic throat-clearing before the interesting part. It is the interesting part.

6. "But this time it really is different"

The strongest objection to all of this deserves to be stated at full strength, not waved away.

It runs like this. Previous technologies were special-purpose, so you could sensibly specify the problem first. This one is general-purpose and moving faster than any planning cycle. You cannot know in advance what it will be good for; the only way to find out is to put it into as many hands as possible and watch what happens. Insisting on a defined problem before you begin is not discipline — it is a way of guaranteeing that you learn last.

I think this is largely right, and it is the best reason to run wide, cheap experimentation early. But it argues against the wrong target. I am not arguing against experimentation; I am arguing against mistaking it for strategy. By all means put the tool in people's hands and let them surprise you — that is precisely how you generate candidate problems worth having. What you cannot do is let the experimentation substitute for the judgement about which of those candidates deserves real investment. And notice what the general-purpose nature of the technology actually implies. When a tool is narrow, the tool itself tells you what to do with it. When a tool can do almost anything, it tells you nothing — and the discipline of choosing where to point it becomes more valuable, not less. That this technology is genuinely different is the reason problem-definition matters more this time, not the reason it can be skipped.

7. The question worth asking

ChatGPT did change something. But it did not change the answer, and the answer was always going to be unglamorous: value comes from improving specific decisions and specific work, and technology is how you do that, not why you do it. What it changed was our willingness to ask the question at all. It took a demonstration this vivid to get boards to say the word strategy about their own operations with real urgency.

That urgency is a gift, if we spend it well. The organisations that will look prescient in a few years are not the ones that moved first in the spring of 2023. They are the ones that used the excitement to ask, at last, which of their decisions deserved to be better — and then reached for this remarkable technology to make them so. We are, once again, at risk of mistaking motion for direction. The correction costs nothing but the willingness to ask a smaller question before the large one, and to sit with the awkward pause while the room works out the honest answer.