

From Strategy to Roadmap

What Leaders Must Do, What They Must Stop, and What They Must Dare

The levels tell you what to aim for. This is the sequence that gets you there, the two ways organisations get it wrong, and the discipline that separates a dip from a failure.

Tommy Mars & Edwin Knoop

Business Strategy · AI Strategy · Digital Transformation · July 2026

Most organisations begin their AI journey at step five.

They begin with the tool. A platform is selected, a licence is bought, a pilot is scoped, and only afterwards does anyone ask which problem it was meant to solve, what would have to change for it to work, and who would be accountable if it did. **The sequence is inverted, and the inversion is not a detail.** It is the explanation for most of what follows.

The evidence is now unambiguous. Sixteen per cent of organisations report no measurable return at all on their largest AI deployment. Fewer than forty three per cent report substantial returns. And sixty one per cent believe they are creating more value than they can actually demonstrate, which is a polite way of saying they cannot prove it.

That last figure is the interesting one. It is not a technology failure. It is a measurement failure, and behind it sits a design failure: **nobody established the baseline before the tool arrived**, because measurement was treated as something you do at the end rather than something you decide at the beginning.

There is a sequence that works. It is not complicated, and it is not new. What is rare is the discipline to follow it in order.

THE EVIDENCE

16% report no measurable return on their largest AI deployment

<43% report substantial returns on their AI investment

61% believe they create more value than they can actually demonstrate

THE SEQUENCE, IN SHORT

01 Diagnosis 02 Ambition 03 Process redesign
04 Operating model 05 Technology 06 Measure & scale

Most organisations enter at 05. The first two steps are a decision, not a build.

The six steps are a dependency chain, not a project plan. Each step depends on the one before it, but that does not mean they run in single file.

The first two steps, diagnosis and ambition, are **a decision, not a build**. For a bounded Level 1 or Level 2 case they can be settled in days. For Level 3 they can be the hardest thinking a leadership team does: a clean sheet reconsideration of what the business would be if it were designed today, with today's technology, rather than inherited from decades of accumulated process. What would we build, what would we never build again, and what would we stop doing entirely.

Once those two are settled, the rest advance together. A Level 1 win runs alongside deeper Level 2 redesign. Waiting for one to finish before starting the other wastes the only genuinely scarce resource, which is organisational attention. And measurement begins at step one, not at step six. **The baseline you fail to capture in week one is the baseline you can never reconstruct in month twelve.**



FIGURE 1 · The sequence. Steps 01-02 are a decision, not a build phase. Steps 03-06 advance together. Most organisations enter at 05.

A note on step 05. Technology sits late in the sequence because it is derived, not because it is small. A tool chosen before the decision becomes technical debt with a licence attached. When the choice does arrive it is an architecture question rather than a purchase: what the data layer has to support, how the capability integrates with the systems already running the company, what it commits you to, and what it would cost to leave. It belongs inside the IT strategy, not beside it.

There are two ways to get AI wrong. One is to wait, because you believe this passes. The other is to move fast, because you believe you are already late. They end in the same place.

The most useful question we ask a leadership team is not how healthy the business is. It is how the business reads the change around it. But before that question can be answered it has to be aimed, because leadership teams are reading three different changes at once and rarely separate them. Is the world permanently more turbulent? Is my industry being structurally rearranged, and how soon? Is the technology a genuine capability shift? These are nested, and the outer lens filters the inner ones.

The third question no longer discriminates. Almost everyone now believes AI is real and will change everything. That belief is close to universal, and it predicts almost nothing. Most organisations that hold it have changed very little: no material shift in investment, no change to the governance model, no redesign of how the company runs. **Enthusiasm is not a strategy, and it is not a plan.**

So the disbelief has moved. It sits in the industry reading. Few executives deny that AI matters; many quietly deny that it reaches them, at least not soon, and not in a business like this one, which runs on relationships, judgement and craft. That is why debating benchmarks with a sceptical board so rarely works. You are arguing one level away from the doubt.

There is something honest in the hesitation. Where AI does the core technical work, the change is dramatic: molecules, materials, perception, code. Where AI has to work through an organisation, the change has so far been modest. **That is not evidence that AI is overrated. It is evidence that the bottleneck has moved.** In a laboratory the organisation is not the constraint. In a company, it is.

WHY WAITING IS EXPENSIVE

AI has not yet changed how companies are run because companies have not changed how they run. Which is exactly why waiting costs. The constraint is your operating model, and an operating model is not something you can buy in six months when the pressure finally arrives.

THE TIMING TRAP

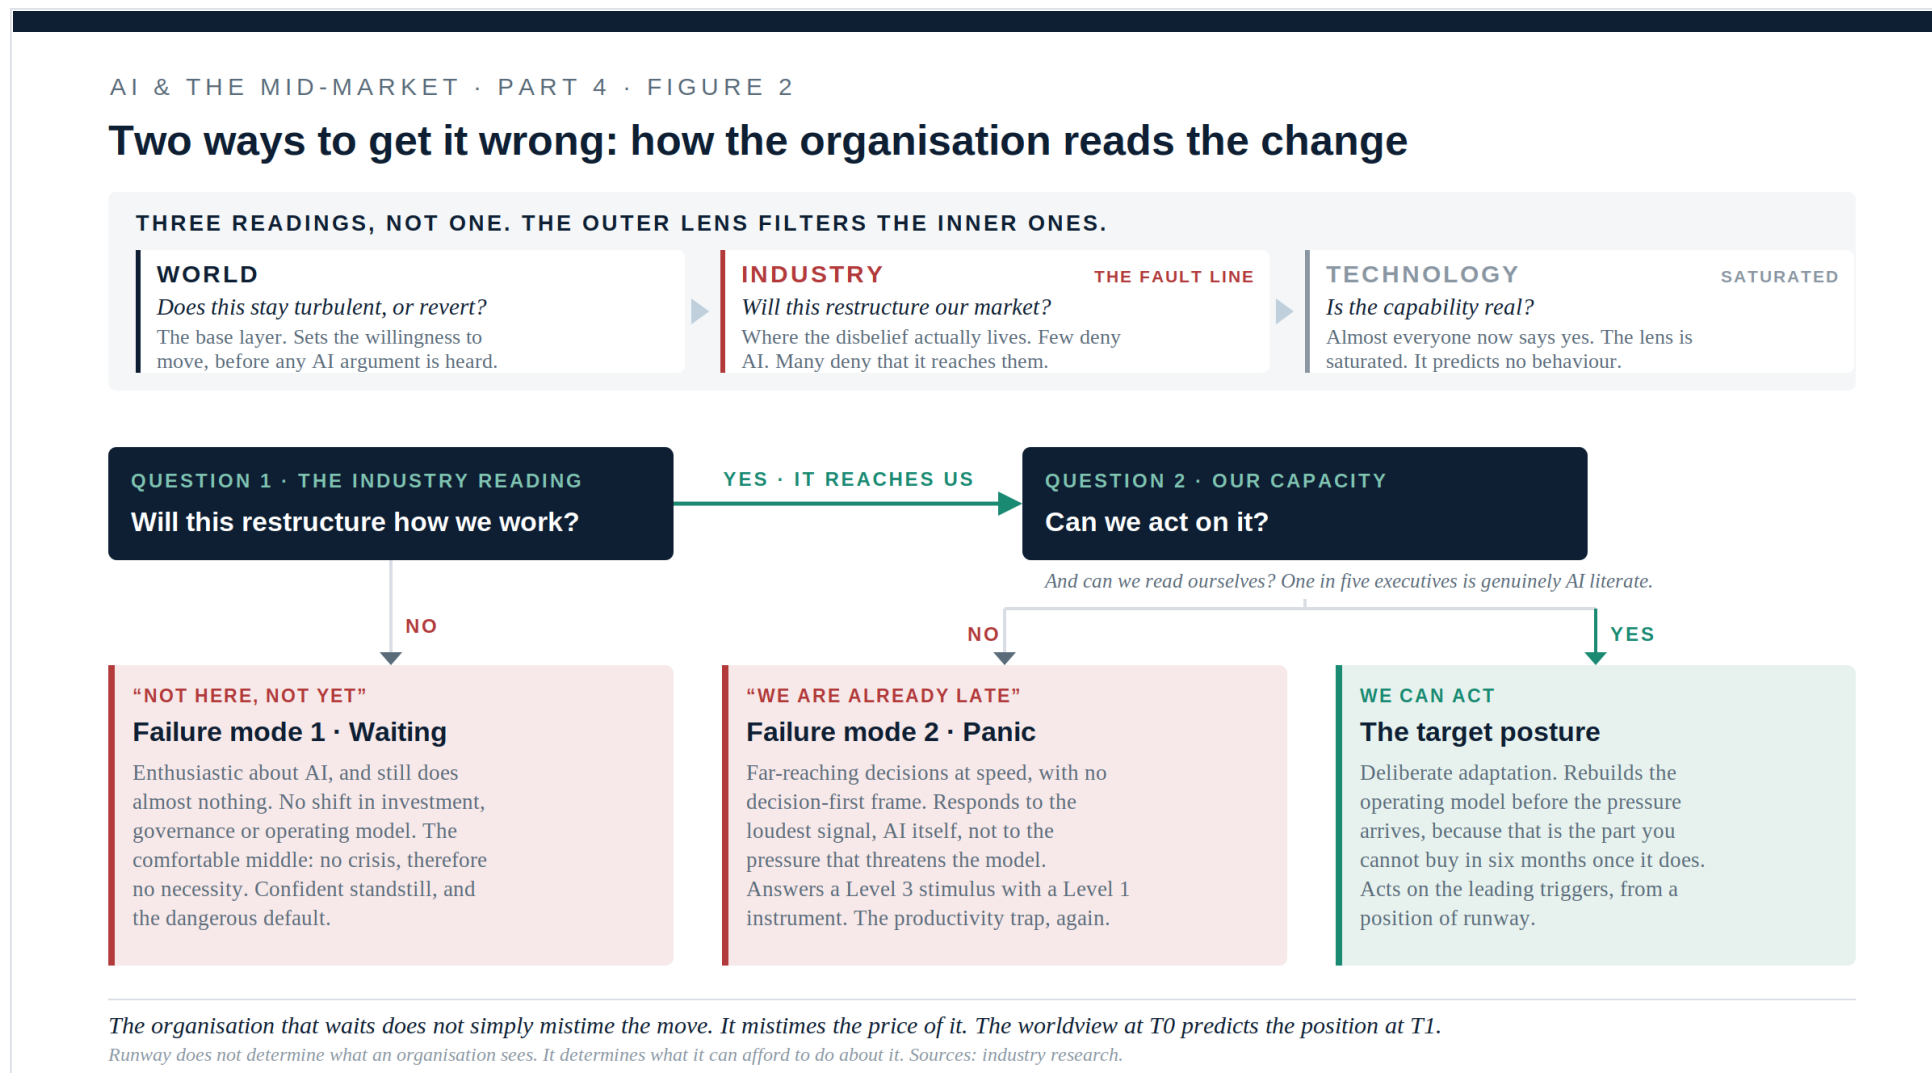
Industry research into business model reinvention finds the leading triggers are technological disruption and shifting customer behaviour. Competitive pressure is not a leading trigger. It follows, once the first two have opened the market.

The organisation that waits rejects exactly the triggers that arrive first, and moves only when the last one is undeniable. By then the early movers have booked their transition, and the runway to respond from strength has gone.

The organisation that waits does not simply mistime the move. It mistimes the price of it.

The worldview at T0 predicts the position at T1.

FIGURE 2 · The technology lens is saturated and predicts no behaviour. The disbelief sits in the industry reading, and that is where the argument has to be made.



Friction is the fastest route to a starting point, and most of it is not where people look for it. Some sits inside a single team. Far more sits between them, in the handovers where work is put down by one function and picked up by another, where waiting time accumulates and nobody is accountable for the total.

Inside the organisation the seams are written down in service levels, response times, uptime commitments and the process model that surrounds them. Those documents describe how the parts of a business are supposed to meet. **The gap between what they describe and what actually happens is your first map.**

Outside the organisation the seams are written down somewhere else. In Part 3 we made the point that a contract portfolio is one of the places where a business model is written down. It is also where a business model meets everyone else's.

And an end-to-end process almost never stops at the company boundary.

“An end-to-end process almost never stops at the company boundary. It runs straight through it, into your customers, your suppliers, and on a project into a dozen parties at once. Everyone maps their own piece. Nobody owns the whole. These boundaries are harder than any internal silo, because there is no manager above them to escalate to. What changes now is that a shared record no longer requires everyone to agree on one system. That was the blocker for twenty years. It is not the blocker anymore.”

Alan Samuel · Rise-X

THE FRICTION NOBODY OWNS

This is the friction that has been structurally unaddressable. Internal silos can at least be reorganised. Cross-boundary friction could not, because it required parties with different systems, different incentives and different lawyers to agree on a single version of the truth.

That is the constraint that has actually changed, and it is why some of the largest Level 2 opportunities now sit in the places nobody has ever owned.

THE SECOND LENS

Friction is not the only lens. Some of the largest opportunities, and most of the existential threats, never present as friction at all. They present as an ecosystem rearranging itself while your own processes run perfectly.

That is the lens the organisation that waits never picks up.

Redesign means subtraction. If nothing has been switched off, no process has been redesigned. A tool has been added to an existing way of working and the result has been called transformation. This is the political part, and it is why it does not happen.

“Redesign is only real when something is switched off. That is the part organisations avoid, because switching something off means telling someone their work has changed. Until one person owns the outcome and carries the number, the new tool simply gets layered on top of the old process, and everyone calls it transformation.”

Sobat Khawaja · Founding Partner, Provectus Consulting

Which brings us to the single most consequential structural decision: **one outcome owner, with a profit and loss mandate**. Not a steering committee, not a centre of excellence, not a sponsor. One person who carries the number.

Accountability and cost visibility, not model quality, separate the organisations getting a return from those still explaining why they are not.

WHAT SEPARATES THE RETURNS

3x

more likely to report established returns, where accountability for AI-informed decisions is clearly defined

5x

more likely, where there is genuine visibility into what the AI actually costs to run

6%

of organisations have genuinely redesigned the process around the AI, rather than adding it to the one they already had

THE STOPPING RULE

Every organisation has a process for starting an AI initiative. Almost none have a process for stopping one.

So the initiatives that are not working are never killed. They are quietly starved of attention, and the organisation learns nothing from them. Decide in advance who is permitted to stop something, and on what evidence. This matters more as systems become agentic, where costs scale with use rather than with value.

Not a governance question. A profit and loss question.

Value does not arrive on a straight line. It arrives on a J curve. Costs come first, capability accumulates invisibly, and the return lands later. Industry research into business model reinvention finds exactly this shape: a transition period in which the benefits stubbornly fail to appear, followed by materially higher revenue in the years after. The same research finds that genuine reinvention is rare. Out of tens of thousands of companies examined, only a few hundred actually did it.

Which produces the contradiction at the centre of this article. We have just argued that you must be willing to stop an initiative that is not working. We are now arguing that returns arrive late and that the dip is normal. Both are true, and holding both is the entire discipline.

So the question is not whether you are in a dip. It is whether you are in the right one.

The difference is visible, but only if you decided in advance what to look at. A dip on the curve shows movement in the leading indicators: adoption is rising, the redesigned process is being used, cycle times are falling, people are working differently. The financial return has not arrived, but the behaviour that produces it has. A dip that is a failure looks entirely different. Nothing has been switched off, the old process still runs alongside the new tool, adoption is flat, and the only thing accumulating is cost.

Judge an initiative on financial return in month three and you will kill every initiative that would have worked in month eighteen. Judge it on nothing at all and you will fund the failures for years.

WHAT REINVENTION ACTUALLY COSTS, AND RETURNS

<2%

of companies studied achieved genuine business model reinvention. A few hundred, out of tens of thousands. It is rare because it is hard, not because it is unrewarded.

+21%

revenue after the transition period, rising materially over the five years that follow. The benefit does not arrive during the dip. It arrives after it.

THE HONEST CLOSE

Much of what is reported as AI failure is measurement failure, not incompetence. Programmes that were working were stopped because the only thing being measured was a number that had not arrived yet.

The dip is an investment if you have runway, and a trap if you do not.

Reinvention pays, but only for those who can afford to wait for it. So the first question is not whether you should reinvent. It is whether you can still choose to.

The series ends here. The work does not.

Four articles: the productivity trap, the strategic case, the three levels, and the sequence that turns them into a roadmap. One argument runs through all of them. AI is an instrument. The decision comes first, and the decision is a business decision.

There is one part of this we are still working out, and we would rather say so than pretend otherwise. In Part 3, researchers at Eindhoven University of Technology identified [value capture, who earns what as the model changes](#), as the dimension both practice and research understand least. We agree. It is the hardest part of Level 3, and it is where the interesting questions now sit.

So we will end on a question rather than a conclusion. On value capture with AI, what would you most want to understand better?

- Who keeps the gain. When AI makes a chain more efficient, does the margin stay with you, move to your customer, or go to your supplier?
- What makes it stick. A gain every competitor can copy within a year is not captured value. It is a temporary discount to your customers.
- Contracts, partners and incentives. How the benefits of a redesigned chain get divided between the parties that made it possible.

ABOUT THE AUTHORS

Tommy Mars

Business strategy and digital transformation. He helps leadership teams connect AI investment to competitive advantage and measurable business outcomes.

Edwin Knoop

AI strategy and enterprise transformation. He works with organisations moving from AI experimentation to enterprise-wide operationalisation.

THE SERIES IN FOUR PARTS

1 The Productivity Trap

Why AI investment doesn't translate to strategic value.

2 Why You Want AI

The strategic case and the ecosystem argument.

3 Three Levels of Value

What AI delivers at each ambition level.

4 From Strategy to Roadmap ■ THIS ARTICLE

The sequence, the two failure modes, and the discipline.

CONTRIBUTORS

Alan Samuel · *Rise-X*

On a mission to improve how businesses collaborate. Rise-X builds trust across business ecosystems: transparency between parties, automated processes across company interfaces, and data sharing that stays secure where it must. He contributed on cross-boundary friction and the shared record.

Sobat Khawaja · *Founding Partner, Provectus Consulting*

She uses art, science and analytics to help clients adapt in a world of relentless change, and is a persistent critic of demo-led AI. Transformation, as she puts it, is messy, human, political and deeply contextual. She contributed on subtraction and single-owner accountability.

ONE LAST QUESTION, AND IT IS THE PRACTICAL ONE

So: what do the first ninety days look like in your business? Where would you start, what would you have to stop doing to make room for it, and what is standing in the way? The sequence does not change. What changes is what it costs you to execute it. Tell us, and the next thing we publish will be about that.

SOURCES · Industry research: EY/Oxford Economics · KPMG Global AI Pulse · McKinsey · BCG · Gartner · MIT Sloan · PwC · AlixPartners · TU Eindhoven · academic research. Figures are indicative, not guarantees. All views are those of the authors.

The problem before the tool.
The ambition before the investment.
The redesign before the deployment.
The owner before the scale.

*Two ways to get AI wrong: wait because you think it passes, or rush because you think you are late.
The organisations that get it right do neither. They decide.*

Follow Tommy Mars and Edwin Knoop for the full series.

#AIStrategy #MidMarket #DigitalTransformation #Leadership