



WHITEPAPER

2026

The AI Question Nobody is Asking.

Why the gap between AI ambition and AI value
starts with the wrong conversation.





Contents

Executive Summary	3
The AI adoption paradox	4
Why data strategy is the missing link	7
CASE STUDY: Strengthening data foundations to deliver against a regulatory committment	10
The cost of misalignment - what poor data strategy costs	12
CASE STUDY: Keeping data foundations in step with organisational ambitions	14
What good looks like: how AI-ready organisations operate	15
The conversation most organisations have not yet had	17
A practical approach to getting started	20
Where to go from here	23
How TXP approaches this	24
About the Author	25
TXP Point of View	26
Sources	27



Executive summary

Most organisations have started on their AI journey. The pilots are running, the demos have impressed, and somewhere a slide deck promises transformational outcomes, yet for the vast majority, meaningful business results remain elusive.

The gap between AI ambition and AI value is not primarily a technology problem. It is a data problem, and more specifically, a strategy problem. When data and AI investments are treated separately from business strategy, organisations end up with sophisticated tools doing the wrong things with the wrong information.

This whitepaper explores why AI without a coherent data strategy, consistently under-delivers. It examines the patterns that separate organisations extracting genuine AI value, from those still searching for a return on their experimentation.

Specifically, this whitepaper is written for senior leaders who are asking the right question... not 'how do we use more AI?' but 'how do we make AI actually work for us?' The answer, as this paper sets out, requires a different kind of conversation. One that connects data quality, governance, and business intent before technology decisions are made. Organisations that have had that conversation are already seeing results. Those that have not are at growing risk of being left behind, not by their competitors' AI tools, but by their own lack of strategic clarity.

88%

of organisations use AI, but only a small minority (~6%) report significant value, with most still struggling to translate adoption into meaningful enterprise impact.

McKinsey, State of AI 2025






The AI adoption paradox


Almost everyone is using it. Few are seeing real value.

There's no doubt that the adoption of AI is accelerating faster than almost any technology in history, and yet conversations with senior leaders across a breadth of industries reveal a persistent and uncomfortable truth: most organisations cannot point to AI investments that have materially changed their business.


McKinsey's State of AI report (2025) makes this hard to ignore. Almost every organisation (88%) is now using AI in some way. But far fewer are seeing real results, only 39% report any measurable impact on profitability, and when you look closer, only around 6% are seeing meaningful financial returns. In other words, while most are experimenting with AI, only a small number are turning it into real business value. That gap between widespread adoption and genuine impact is the real challenge organisations now face.



of organisations now use AI in at least one business function (McKinsey, 2025)



report AI impact significant enough to materially move enterprise profitability (McKinsey, 2025)



cite data quality as their top data integrity challenge, the leading barrier to AI value (Precisely, 2025)

The reasons behind this gap are well understood, even if they're not always acted on; AI is only as good as the data behind it. In many organisations, that data is incomplete, inconsistent, or siloed. As a result, outputs may look impressive, but they don't always translate into better business decisions.

This isn't a limitation of the technology itself. It's a question of approach. Too often, organisations move quickly to implement AI without first stepping back to ask a few fundamental questions; what decisions are we trying to improve? What data do we need to support them? Can we trust it? And, is it connected in a way that makes it usable?

These are not difficult questions, and they are rarely ignored entirely. But they are often asked too late, in the wrong sequence, or without the right stakeholders involved.

The experimentation trap

For multiple organisations AI investment has followed a recognisable pattern; a business unit identifies an opportunity or a problem, a technology team or external vendor proposes an AI solution, a pilot is commissioned, usually with controlled data, optimised conditions, and a narrow brief, and, the pilot works, it delivers 'outputs', the case study is written, and then somewhere between the pilot and enterprise-wide deployment, progress stalls!


This is 'the experimentation trap'. The conditions that make a pilot successful; clean data, a clearly defined problem and close management attention, rarely exist at scale across the organisation. The pilot proves the technology can work; it does not prove the organisation is ready to work with it. Gartner predicted that at least 30 per cent of generative AI projects would be abandoned after proof of concept by the end of 2025, not because the technology failed, but because of poor data quality, unclear business value, or costs that could not be justified without a credible return. More recently Gartner suggested this may have been conservative, with some estimates rising to 50% abandonment in practice.



The question most organisations are asking is 'which AI tool should we buy?'

The question they should be asking is '**what would our organisation look like if AI was genuinely embedded in how we make decisions, and what would we need to be true for that to happen?**'





“Nearly half of organisations cite data quality and governance as the biggest barrier to AI success.”

IDC, 2025

Why data strategy is the missing link

It is not the algorithm. It is what you feed it.

Data scientists are clear about what makes AI effective: it comes down to data. If the data is poor, incomplete, or inconsistent, the results will be too, regardless of how advanced the model is.

This isn't a new insight. It has always been true. Yet many organisations are still not acting on it. Research reflects this; data quality is consistently cited as the biggest challenge, and for many, it remains the primary barrier to making AI effective. Despite continued investment in AI, this hasn't materially shifted and that is in itself telling.

The reason is often structural. AI initiatives are typically led by innovation or technology teams, while data sits within IT or operations, and business strategy is defined elsewhere. When these areas are not aligned, organisations risk developing AI that is disconnected from operational reality, and data programmes that fail to deliver meaningful value.

What a data strategy really is (and what it isn't)

Data strategy is often overcomplicated. It's associated with governance frameworks, architecture diagrams, and long transformation programmes. That's understandable, but it misses the point.

At its simplest, a data strategy answers one question: what data do we need to make better decisions, and can we trust it? From there, the focus becomes practical; which decisions matter most? What data supports them today, and how reliable is it? Where are the gaps? And what would it take to fix them in a way that reflects the value at stake? Most importantly, how does this connect to AI? Because without that connection, AI remains an experiment, not a driver of outcomes.

The five dimensions of data readiness for AI

Through working with many organisations, we see a consistent pattern emerging in what separates those with genuine AI readiness, from those still preparing for it. It comes down to five dimensions:

- › Data quality: Is the data accurate, complete, and consistent enough to be trusted?
- › Data integration: Can the right data from across the organisation be brought together in a usable form?
- › Data governance: Are there clear rules about how data is managed, accessed, and maintained?
- › Data literacy: Do the people making decisions understand what the data means and what its limitations are?
- › Business alignment: Is the data programme connected to the decisions that drive business value?

Most organisations are reasonably strong on one or two of these dimensions. Very few are strong across all five. And it is only when all five are working together that AI investments can consistently deliver at scale.



How a leading pensions administration organisation strengthened its data foundations and delivered on a major regulatory commitment

A leading Pensions Administration organisation was working to deliver a complex pension valuation programme against a fixed regulatory timeline. As delivery progressed, it became clear that the data infrastructure underpinning it needed to be in better shape to carry the programme forward with confidence.

Rather than carry that risk forward, the organisation made a deliberate decision to address the foundations properly. TXP embedded within the organisation's teams, taking ownership of the highest-complexity components, designing and building the core calculation engine, delivering the management interface, and integrating the solution with a major third-party pensions platform. The organisation is set to meet its regulatory deadline.

Data accuracy improved materially, and the internal architecture is now in a shape that supports the next phase of its analytics ambitions, moving forward from a stable, trustworthy foundation.





Practical recommendations

What leaders can do now



Map your most valuable decisions against your current data.

Ask: 'what information do we actually use, and how confident are we in it'?



Commission a data quality audit focused on the use cases where AI investment is planned, not a full enterprise audit, but a targeted assessment of what you have versus what you need.



Bring data, technology, and business strategy into the same conversation.

If these three groups have not met to discuss AI readiness, that meeting is overdue.



Set a clear threshold for data quality before scaling any AI pilot.

'Good enough for a demo' and 'good enough for enterprise deployment' are very different standards.



Treat data governance as an enabler, not a constraint.

The organisations that move fastest with AI are those that have clear, trusted rules, not those that have the most flexible ones.



The cost of misalignment – what poor data strategy costs

This is not a theoretical risk, it's measurable.

One of the reasons data strategy is often overlooked is that the consequences of getting it wrong are not immediately visible. A failed AI initiative rarely appears as a clear line on a P&L. Instead, the impact accumulates over time through wasted technology investment, slower decision-making, reduced confidence in data, and missed opportunities where better decisions could have been made.

These costs are now being quantified more clearly. Gartner has long estimated that poor data quality costs organisations an average of \$12.9 million per year. In an AI context, the impact is amplified. When AI systems rely on poor-quality data, the issue does not remain isolated, it scales. A flawed dataset does not result in a single poor outcome; it introduces repeated, systemic errors across the decisions the model supports.


Further Gartner research highlights how widespread this challenge remains. Around 63% of organisations either lack, or are unsure whether they have the appropriate data management practices in place for AI. As a result, Gartner predicts that through 2026, 60% of AI projects will be abandoned due to a lack of AI-ready data. This reflects an ongoing pattern, as many organisations continue to struggle to move from experimentation to measurable value.

The impact is visible across sectors. In financial services, it can be seen in credit and fraud models that reinforce outdated patterns or produce unreliable results. In retail, forecasting models fail to reflect changing customer behaviour because underlying data has not kept pace. In the public sector, services can be directed based on data that no longer reflects the people or needs they are meant to serve. In each case, the limitation is not the model itself, but the data that underpins it.


There is also a less visible cost – delay. Continued investment in AI without addressing underlying data foundations reduces the return on that investment over time. Models become less effective, governance complexity increases, and the gap between potential and realised value widens. In contrast, organisations that strengthen their data foundations are better positioned to improve outcomes, adapt more quickly, and scale AI capabilities more effectively. Over time, this creates a cumulative advantage that becomes increasingly difficult to replicate. 12

\$12.9M

average annual cost of poor data
quality per organisation
(Gartner)



of AI projects without AI-ready data
will be abandoned through 2026
(Gartner, 2025)



of organisations lack the right
data management practices for AI
(Gartner, 2025)

Practical recommendations

Quantify your current AI investment and map it against measurable outcomes. If you can't articulate what a specific investment has changed about a specific decision, that is a signal worth acting on.

Calculate the decision-quality gap. For your three most important recurring business decisions, ask: 'how much better could that decision be with better data'? 'What would that improvement be worth'?

Set a data quality threshold as a deployment criteria for AI. If the underlying data does not meet a defined standard, the AI system should not be deployed, regardless of how impressive the demo was.

How a global healthcare organisation kept its data foundations in step with its ambitions

A global healthcare organisation was evolving rapidly with new products, new services, and expanding digital capabilities. As the organisation scaled, it needed its foundations to keep pace, so that decisions stayed well-supported and insight could grow with the business rather than lag behind it.

Working alongside the organisation's own teams, TXP modernised key elements of the data estate, improved the quality and flow of operational data, and created clearer integration pathways between systems.

The approach was deliberately sequenced, prioritising what was needed to support day-to-day operations first, while building the platform for more advanced analytics and AI capability in future phases.

The organisation came away with greater confidence in its core datasets, a more consistent view of information across teams, and a stronger foundation for the next stage of its digital and data transformation.





What good looks like: how AI-ready organisations operate

The **6%** are not doing anything extraordinary. They are doing something deliberate.

The organisations that consistently extract value from AI are doing the fundamentals deliberately and consistently. The difference isn't the technology, it's how they think about decisions, data, and ownership.

They start with the decision, not the technology

Instead of broad ambitions like improving customer experience, they focus on specific outcomes, what decision needs to be made better, faster, or more consistently. That level of clarity shapes everything that follows: the data required, the approach taken, and how success is measured.

They treat data as infrastructure, not a project

In these organisations, data quality and governance are ongoing responsibilities, not project-based fixes. There is clear ownership, clear standards, and a shared understanding that if the data isn't reliable, nothing built on top of it will be either.

They build capability, not dependency


External partners still play an important role, but they don't replace internal understanding. Organisations that get long-term value from AI invest in their own ability to interpret, adapt, and improve what they build. Over time, they become more self-sufficient, not more reliant.

They align three things that most organisations treat separately

Business strategy, data strategy, and technology decisions are not treated as independent conversations. They are addressed together; 'What are we trying to achieve'? 'What data do we need'? And 'how should we deliver it'? That alignment is what turns isolated AI experiments into outcomes that scale.



Practical recommendations



Run an AI readiness diagnostic before your next AI investment decision.

Map where you are across the four characteristics described above. Be honest about the gaps.



Identify one high-value decision in your organisation and build the data case around it.

Not a technology pilot, a data and decision audit.



Create a shared forum between your business strategy, data, and technology teams.

If this does not exist, create it. If it exists but meets rarely, change the cadence.



Establish a knowledge transfer requirement in any AI engagement with an external partner.

At the end of the engagement, your team should be able to explain, maintain, and develop what has been built.





The conversation most organisations have not yet had

Strategy before technology

At the centre of any successful AI programme is a clear alignment on intent. This is not primarily a technology discussion. It is a strategic one.

Organisations need to define what they are trying to achieve, what conditions must be in place to achieve it, and how current capabilities compare to those requirements. Without this clarity, technology decisions are made in isolation from business outcomes.

Establishing this alignment requires input from multiple functions and a shared understanding of both opportunities and constraints. It also requires discipline, prioritising foundational considerations over immediate technology adoption in order to enable more effective value realisation over time.

Organisations that consistently generate value from AI are not necessarily those with the largest budgets or the most advanced tools. They are those that connect business objectives to data capability before making significant technology investments.

“AI projects don't fail because the technology doesn't work. They fail because nobody agreed on what decision they were trying to improve before they started”

Ankur Gupta, Head of Data & AI Practice, TXP

What this alignment involves

In practice, this requires three elements to be considered together; First, a clearly defined business objective, expressed in measurable terms and linked to specific decisions or outcomes. Second, an assessment of current data capability, including availability, quality, and gaps relative to the objective. Third, a technology approach designed to address those gaps and enable the intended outcome.

These elements are interdependent. Business objectives shape data requirements, data constraints influence what can be delivered in the near term, and technology decisions determine where investment is required. Effective organisations manage these considerations as part of an ongoing, iterative process rather than a one-time sequence.

The role of external perspective

In many organisations, achieving this level of alignment internally can be challenging. Teams are often closely aligned to existing processes, data structures, and prior decisions. As a result, it can be difficult to objectively assess current limitations or identify where change is required. In addition, responsibilities for business strategy, data, and technology are frequently distributed across different functions.

External perspective can support this process by providing objectivity and structure. This includes facilitating alignment across stakeholders, challenging assumptions where appropriate, and helping organisations assess their current position against defined objectives.

The value is not solely in technical expertise, but in enabling a clearer, shared understanding of how strategy, data, and technology need to come together to deliver measurable outcomes.





Practical recommendations



Schedule the conversation.

Bring your business strategy, data, and technology leads into the same room with a single agenda item: 'are we building our AI programme on a foundation that can support it'?



Be specific about intent.

'We want to use AI to improve efficiency' is not sufficient.
'We want to reduce the time to process a new customer application from five days to same-day, without increasing risk' is a starting point.



Ask the data question first.

Before any technology conversation, establish: 'what data would we need to achieve this'? 'Do we have it'? 'Do we trust it'?



Build in external challenge.

If your organisation has not had an independent assessment of its data readiness in relation to its AI ambitions, that is a gap worth addressing.



A practical approach to getting started

Progress, in the right sequence

It's natural for organisations to pause at this point. If effective AI depends on strong data foundations, and those foundations take time to build, it can raise the question of whether to wait before moving forward.

In practice, most organisations don't need to choose between progress and preparation, what matters is sequencing.

Those making steady progress tend to focus first on areas where there is clear business value and a reasonable level of data readiness. This allows them to demonstrate early impact, while continuing to strengthen their data foundations in parallel.

It's a more balanced approach and one that enables progress without losing sight of what's needed to sustain it.



What effective organisations do differently

Organisations that translate AI into measurable value tend to follow a consistent set of principles;

They begin by establishing clarity on where AI can support the most important business decisions. This involves identifying a small number of high-value opportunities and assessing them in terms of both potential impact and current data capability. The objective is not to create a long list of use cases, but to focus on those that are both meaningful and achievable.

In parallel, they take a structured view of their data landscape. This includes understanding what data is available, how reliable it is, where gaps exist, and what would be required to address those gaps. This assessment is grounded in business need, rather than conducted as an abstract data exercise.

From there, attention shifts to enabling delivery. This typically involves improving data quality where it matters most, establishing clear ownership and governance for critical data assets, and ensuring that the necessary integration and infrastructure are in place to support deployment. Alongside this, organisations invest in building internal capability, so that what is delivered can be maintained, adapted, and extended over time.

As value is demonstrated, the approach is expanded. Additional use cases are prioritised based on business relevance and data readiness, with each iteration informed by what has already been learned. Over time, this creates a repeatable model for delivering AI outcomes, supported by stronger data foundations and increasing organisational confidence.



From isolated initiatives to repeatable outcomes

The key distinction is not speed, but direction.

Organisations that struggle often pursue multiple AI initiatives without a clear connection to data readiness or business value. Those that succeed take a more deliberate approach, linking decisions, data, and technology from the outset, and building capability as they go.

This creates momentum, early results build confidence, stronger data foundations enable broader application, and over time, AI shifts from a series of experiments to a capability that can be applied consistently across the organisation.



Practical recommendations

Do not wait for perfection before starting. Identify your highest-value, most data-ready opportunity and begin there.

Treat the first use case as a capability-building exercise, not just a technology deployment. The goal is to build the understanding to do it again.

Set explicit success criteria before you begin and measure them honestly at each phase gate.

Build the data foundation in parallel with the first use case, not sequentially after it. The organisations that move fastest are those that do not treat these as separate programmes.

Where to go from here

Turning intent into impact

The gap between AI adoption and measurable value is no longer a question of awareness. Most organisations recognise it. the challenge now is how it is addressed in practice.

Closing that gap does not require more experimentation or additional tooling alone. It requires alignment between business priorities, data capability, and the role technology plays in delivering outcomes. Without that alignment, even well-intentioned AI initiatives can struggle to translate into consistent value.

Organisations that are making progress tend to take a more deliberate approach. They define the decisions that matter, assess their data with a degree of realism, and apply technology in a way that is clearly connected to those priorities. Over time, this creates a more stable foundation for scaling AI beyond initial use cases.

This is also where many organisations benefit from stepping back. Not to slow progress, but to ensure that effort is being directed in the right areas, where there is both meaningful business value and the conditions required to support it.

The opportunity is not just to adopt AI, but to apply it in a way that is repeatable, scalable, and grounded in real outcomes. For most organisations, the question is not whether to act, but how to ensure that what comes next builds toward that outcome.

“AI delivers value when strategy, data, and technology are aligned, not when they are developed in isolation.”

Ankur Gupta, Head of Data & AI Practice, TXP

How TXP approaches this

Most organisations are not short of AI ambition but rather a clear line of sight between what matters commercially, what their data can realistically support, and what it would take to turn that into something repeatable.

That is the gap we can close, by starting with evidence. The decisions that would benefit from better information are identifiable, and the data to support them often exists in some form, but what is missing is the connection between them, and a realistic view of what it would take to make that connection work consistently.

We focus on establishing where value is genuinely available, combining what the organisation is trying to achieve with an honest assessment of what its data can currently support, identifying opportunities that are both meaningful and achievable, where improvement in decision quality would produce a measurable outcome and where the conditions to support are either in place or within reach.

From there, we making foundations reliable enough to build on, data that is trustworthy in the places that matter most, connected, governed, and understood by the people who depend on it. That shift, from data that exists to data that can be acted on with confidence, is where most of the practical work happens, and where the conditions for sustainable AI value are established.

What follows is the application of AI in a way that can be repeated and extended as part of how the organisation actually operates. Throughout, we treat delivery as capability-building. What is built should be explicable and maintainable by the people who will work with it long after the engagement ends. That is not a constraint on ambition. It is what makes the ambition sustainable.



About the Author



Ankur Gupta

Ankur is a data and AI strategist at TXP, focused on helping organisations close the gap between ambition and value. He works with mid-market enterprises to define how AI can support real business decisions, and what data and capability are required to make that possible. His approach centres on aligning business intent with data reality, ensuring that AI initiatives are grounded, measurable, and built to scale.

TXP Point of View

Most organisations don't struggle with ambition. They struggle with turning it into value, and that gap is rarely about technology alone.

Progress happens where strategy, data, and technology are aligned. Where there is clarity on what needs to change, and the capability to deliver it.

This goes beyond advice. It requires the ability to design, build, modernise, and run the systems that underpin change, supported by the right skills, applied in the right way. Consulting, engineering, and resourcing all play a role, but only when they are connected.

As a UK-based services integrator, TXP works with mid-market organisations to turn AI and digital ambition into measurable outcomes. The focus is not on increasing activity, but on improving alignment and connecting business priorities to data foundations and technology delivery.

Because progress is not defined by what is invested. It is defined by what works.





Sources

McKinsey & Company (2025)
The State of AI: How organisations are rewiring to capture value
Mckinsey.com

Gartner (2018; 2025)
Poor Data Quality Costs Organisations \$12.9M Per Year; AI-Ready Data Insights
Gartner.com

International Data Corporation (2025)
AI Adoption and Data Challenges Survey
idc.com

Precisely (2025)
Data Integrity Trends and Insights
Precisely.com

Forrester (2025)
Strategic AI Readiness: From Hype to Scalable Impact
Forrester.com



About TXP

TXP (Technology x People) is a UK leading digital transformation business, delivering consulting, development and technology resourcing solutions to a diversified blue-chip client base. With a UK base and global presence, the company currently works across public and private sectors, providing legacy modernisation services, custom software development and IT resource augmentation.

www.txp.co.uk
hello@txp.co.uk
[+44 \(0\) 121 788 4550](tel:+441217884550)