AI in Real Estate

What to believe, what to question and what to do

June 2026

Contents

Executive summary

Real estate firms are making consequential decisions about AI under conditions of genuine confusion. Vendors are overselling. Analysts are generalising. And the volume of commentary has become, paradoxically, an obstacle to clear thinking rather than an aid to it.

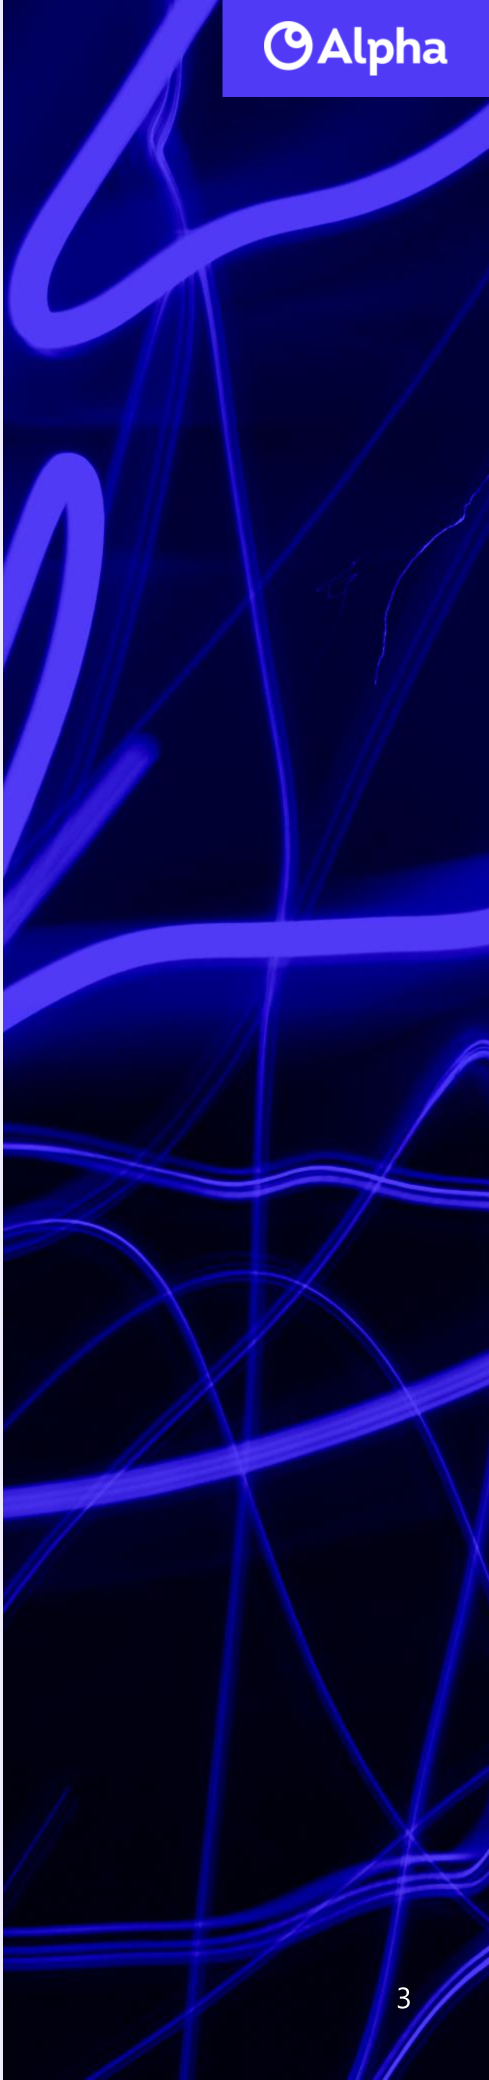

The Manufactured Confusion

The AI debate in real estate has produced a great deal of noise and very little that will help make better decisions. Vendor briefings, conference keynotes, analyst reports, and LinkedIn feeds have converged on a single narrative: AI is transforming everything, SaaS is dead, agents will run your business, and those who hesitate will be left behind.

The volume of AI commentary has become, paradoxically, an obstacle to clear thinking. Vendors have retrofitted AI branding onto features that existed under different names twelve months ago. Consultants have repackaged generic digital transformation advice with new terminology.

The consequences are real. Firms are launching pilots without the data foundations to support them. Investment committees are being presented with AI-generated analysis they cannot interrogate. Technology budgets are being reallocated toward AI initiatives with poorly understood cost structures and no governance framework.

This paper is not an argument against AI. AI is strategically important, the opportunity is genuine, and the firms that navigate this period well, will build durable competitive advantage. But getting it right requires cutting through the noise, including, as we will argue, some of the noise generated by your own instinct to be cautious.

The question this paper answers is straightforward: what is actually true, what is noise, and what should you do about it?

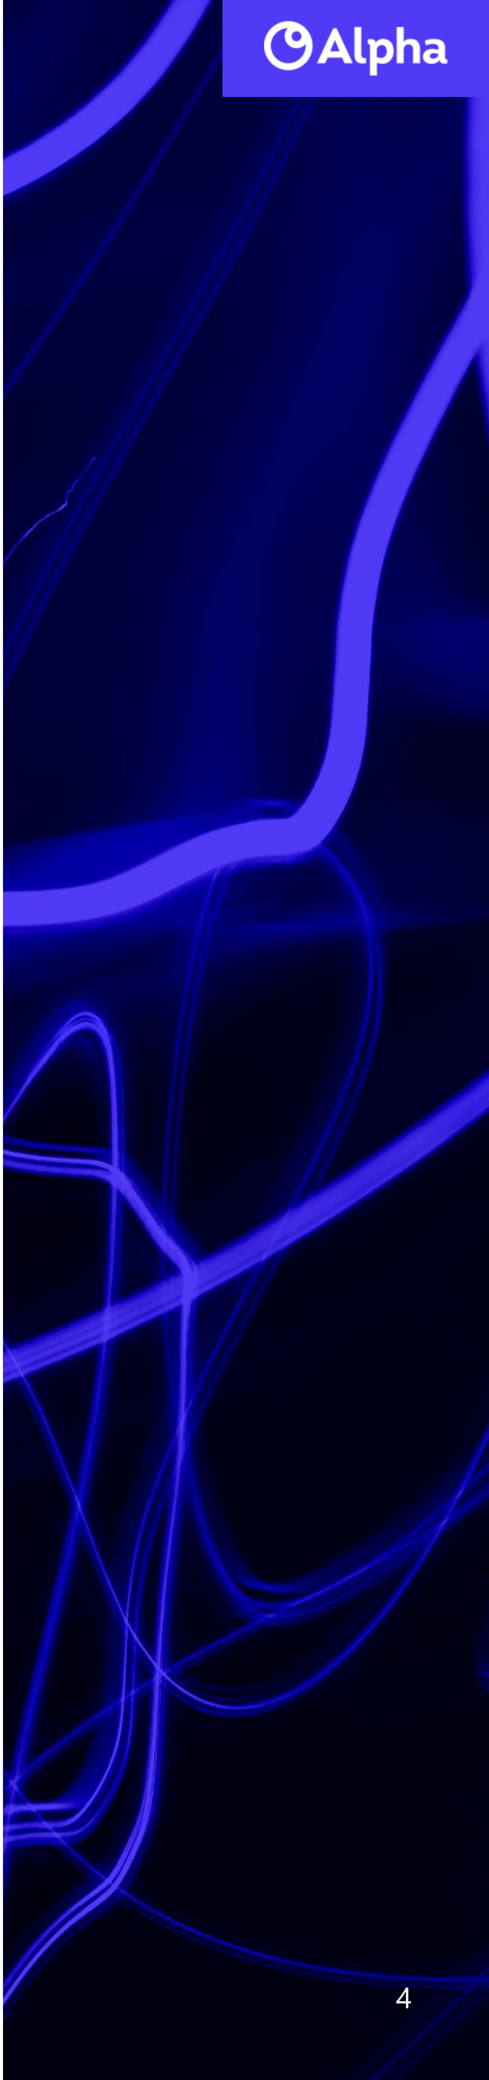

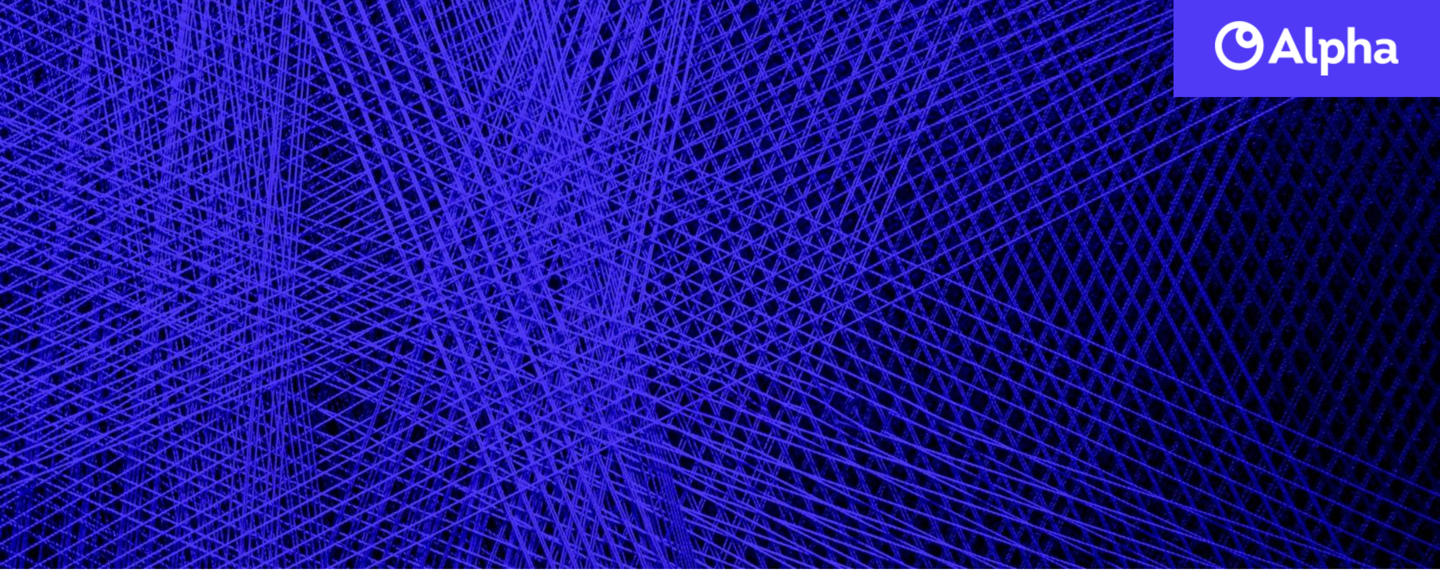

Your Scepticism Is Partly Right

Most senior real estate leaders we speak to are already sceptical of AI hype. That instinct is well-founded, but only in specific places. Blanket scepticism is as costly as blanket enthusiasm. The discipline is knowing which is which.

Where Caution Is Justified

The market is flooded with AI products that do not deserve the name. The most common pattern is a thin wrapper around a general-purpose large language model, presented as an industry-specific solution. For low-stakes productivity tasks, drafting emails, summarising notes, first-pass research, these tools can be useful. But they are being sold, and in some cases deployed, for purposes they cannot responsibly serve.

AI underwriting tools that cannot produce an auditable calculation trail are not underwriting tools. They are generators of plausible-sounding analysis without a defensible methodology. No investment professional should sign off on a multi-million-pound acquisition based on outputs they cannot trace, verify, or explain to their investment committee.

Vendor AI claims deserve similar scrutiny. Most established platforms have added AI features in the past eighteen months. Some represent genuine improvements. Many are existing features relabelled, or third-party LLM integrations that add a natural-language interface to an unchanged underlying system. Evaluating what has actually changed versus what has been rebranded requires harder questions than most vendor conversations allow.

And then there is the agentic AI wave. The concept is real and the long-term implications for real estate operations are significant. But agentic architecture is complex, expensive to run, and difficult to govern. Firms building agentic workflows without understanding token costs, auditability requirements, or process dependencies are not innovating. They are storing up expensive, hard-to-diagnose problems.

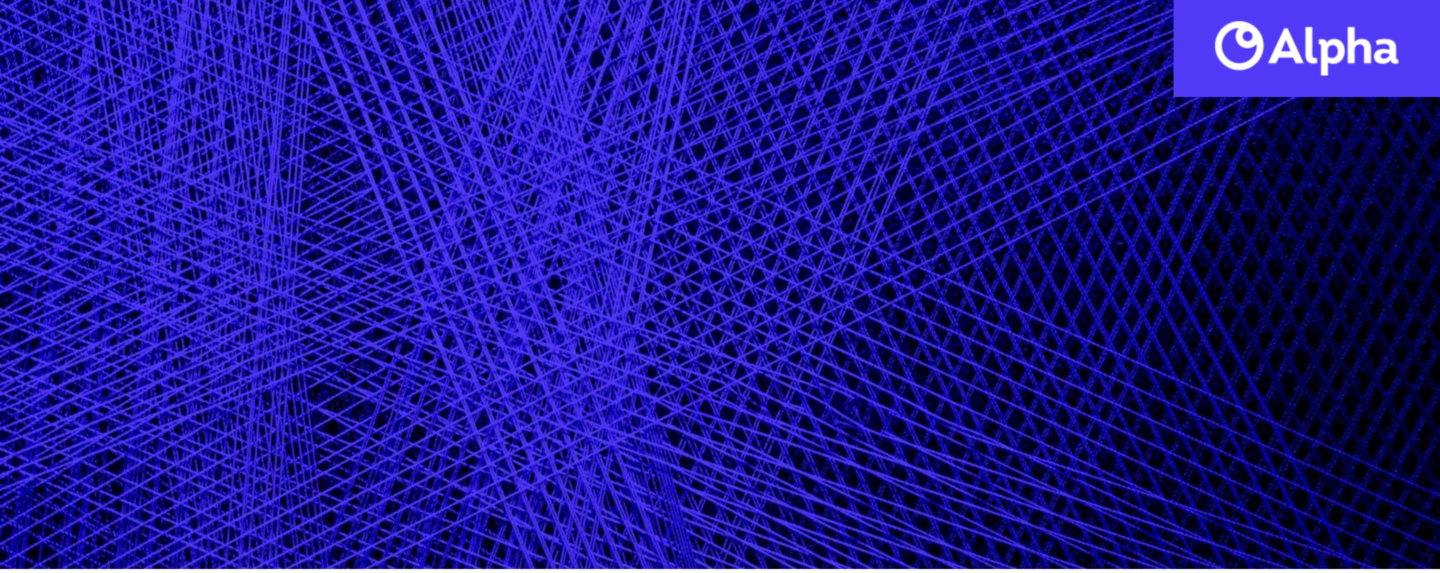

Where Scepticism Will Cost You

The harder conversation is this: some of what feels like prudent caution is actually strategic lag. In a sector where competitive differentiation is increasingly tied to operational efficiency and data quality, that lag compounds.

Dismissing AI entirely while peers automate high-volume, low-judgment workflows is not prudence. Lease abstraction, document interrogation, reporting generation, investor communication drafts, these are well-understood, production-proven use cases where productivity gains are real and measurable. Firms not exploring them are not being careful. They are falling behind.

There is also a category error embedded in much of the caution we hear: conflating 'AI cannot make our investment decisions' with 'AI cannot add value anywhere.' These are not the same claim. The first is correct and important. The second is wrong and costly. The space between those two positions, the well-defined workflows where AI genuinely accelerates competent human work is where the near-term value lies. Firms that cannot disaggregate these claims will systematically under-invest in the right places while remaining appropriately sceptical of the wrong ones.

The question is not whether to engage with AI. It is whether you are building the foundations that will make that engagement durable and defensible, or accumulating a problem that will become very expensive to unwind.

Three Things Nobody Is Telling You

The AI conversation in real estate has blind spots that are not commercially convenient for vendors and not instinctively prominent for technologists, but are directly relevant to the leaders who will be accountable for these decisions.

1 Cost Is Not Zero

AI is routinely presented as a cost reduction strategy. That is true at the output level, the right deployments do reduce operational costs. But the cost of building, running, and maintaining AI systems is itself significant and frequently underestimated. Token usage is not a fixed or predictable expense. Poorly designed agentic workflows, where AI models invoke other models unnecessarily or context windows are mismanaged, generate costs that scale non-linearly with usage. Firms building these systems without visibility into their cost architecture are not saving money. They are creating a liability.

Cost and governance must be designed in from the outset. If you cannot see what your AI is costing you, by workflow, by function, by agent interaction, you cannot manage it. And you cannot manage what you cannot see.

2 Auditability Is Non-Negotiable

Real estate investment management operates in an environment where decisions carry significant financial, regulatory, and fiduciary weight. Underwriting assumptions, valuation methodologies, investor reporting, and fund performance attribution all require a clear, complete demonstration of how a conclusion was reached and what data it was based on.

An AI system that cannot show its working is not a productivity tool in this context. It is a liability. Current market data reflects this: a significant proportion of investment committees report distrust of AI-generated analysis, not because AI is incapable, but because outputs cannot be traced to verified sources and auditable methodologies. That is a design problem, not an adoption problem. Governance and audit trails are not features to be added later. They are the product.

3 Not Every Workflow Needs an Agent

The enthusiasm for agentic AI has created a tendency to reach for the most sophisticated solution regardless of whether the problem requires it. Many of the highest-value automation opportunities in real estate do not require AI at all. They require well-defined processes, clean data, and workflow automation. AI should be reserved for workflows where its specific capabilities, reasoning over unstructured information, language generation, pattern recognition across large datasets, are genuinely required. Applying AI to problems that automation solves more cheaply and reliably is not innovation. It is waste.

SaaS Is Not Dead — It Is Being Misdiagnosed

The 'death of SaaS' narrative has gained significant traction. The argument, in its strongest form, is that AI agents will replace structured software platforms, that instead of buying and configuring systems, firms will simply instruct AI to build and run what they need, on demand. This is an intellectually interesting claim. It is also practically wrong for this sector, at this point in time, and for reasons that are not temporary.

Real estate investment management firms operate in regulated environments where auditability, traceability, and system-of-record integrity are non-negotiable. These requirements do not disappear because AI exists. The problem is not SaaS. The problem is firms deploying AI on top of broken processes and fragmented data, then attributing the failure to their platforms. A firm that cannot produce clean, consistent data from its existing systems will not solve that problem by replacing those systems with AI. It will reproduce it at greater speed and cost.

The durable architecture is a layered one: best-in-class core platforms providing system-of-record, data model, and auditability foundations; workflow automation, integrations, and selective AI built around that core to deliver operational efficiency and competitive differentiation. 'Build everything with AI' is as strategically naive as 'buy everything from one vendor.' Both are appealing in their simplicity. Both are expensive when they meet reality.

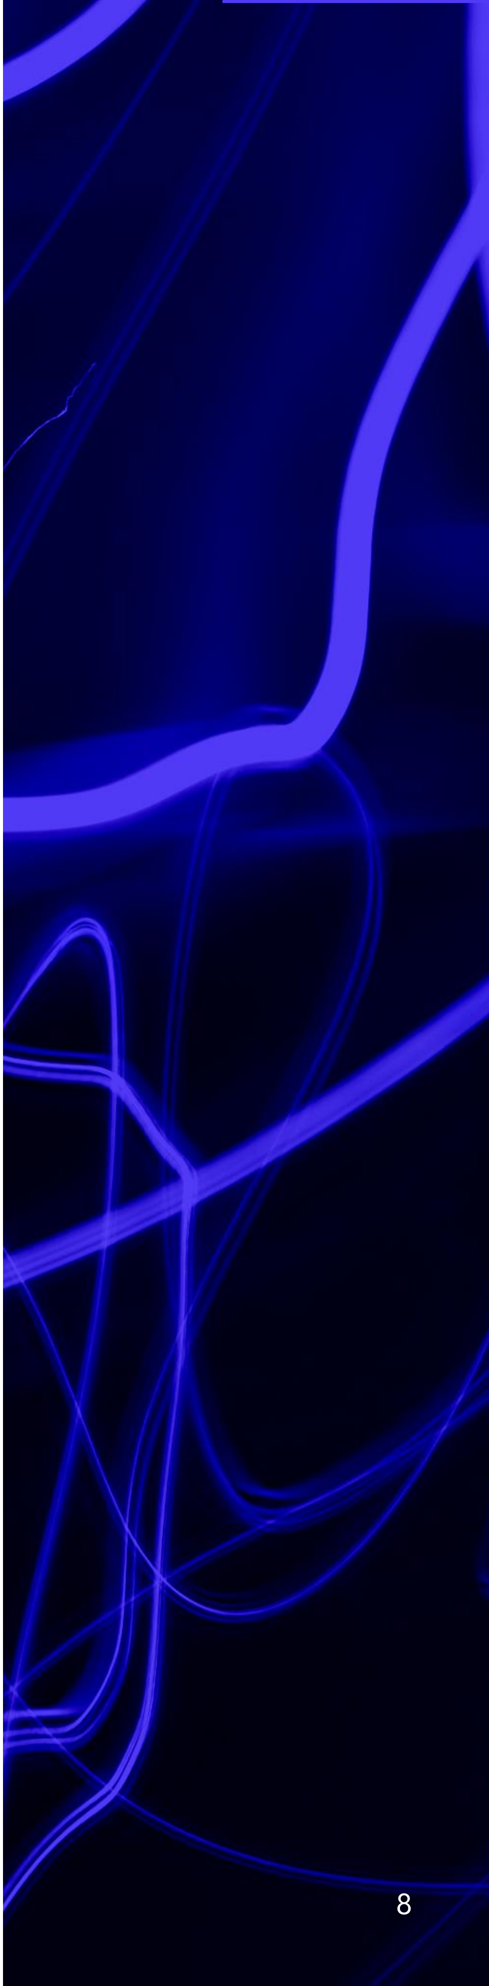

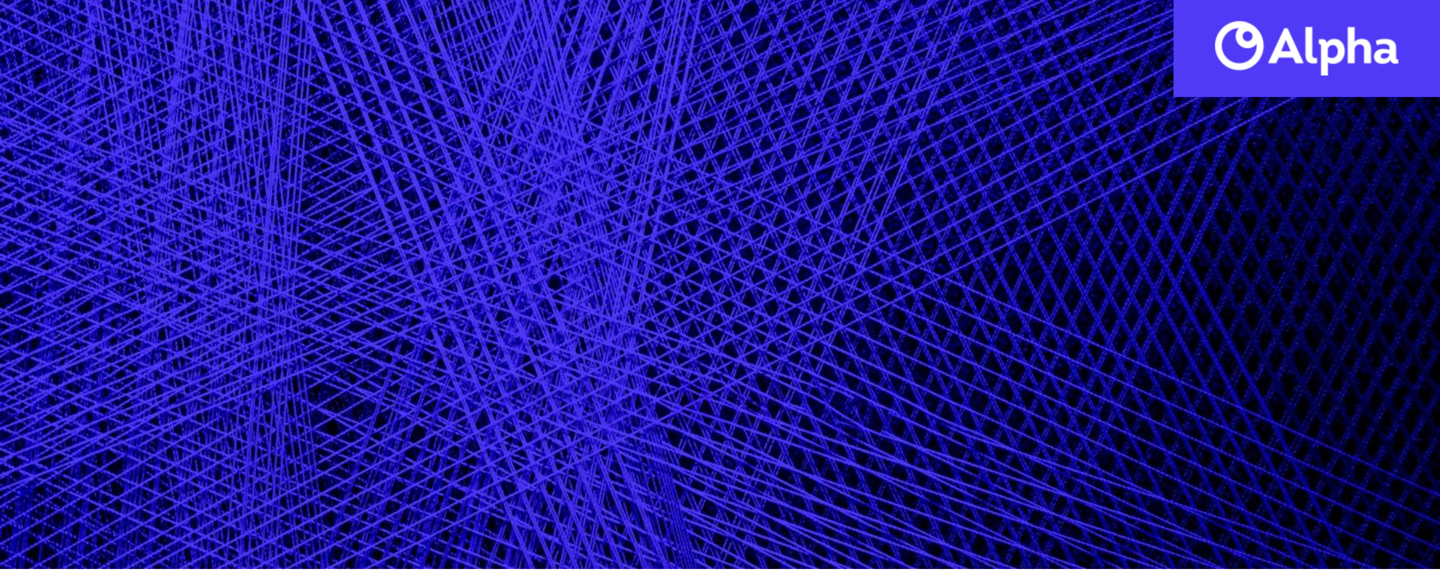

The Foundation Nobody Wants to Discuss

There is a conversation that should precede every AI initiative in real estate investment management. It is rarely had, because it is not commercially convenient for vendors, not appealing to boards under pressure to show AI progress, and not inherently exciting to those tasked with delivering it. But it is the conversation that determines whether AI deployments succeed or fail. That conversation is about process.

Agentic AI cannot operate effectively on fragmented, inconsistent, or unstructured processes and data. This is not a temporary limitation. It is a structural dependency. An AI agent requires defined inputs, predictable intermediate states, and auditable outputs, exactly what a well-governed process provides. Without that foundation, the risk of incorrect, non-auditable, or ungovernable AI behaviour increases regardless of the quality of the AI system itself.

Most real estate investment management firms have not resolved their process and data fragmentation. Fragmented data flows, dependency on individual knowledge, manual reconciliation, and legacy workarounds are not edge cases. They are the norm. The predictable resultof deploying AI on top of this environment is expensive pilots that cannot scale, outputs that cannot be trusted, and governance failures that arrive at the worst possible moment.

The path forward is sequential and disciplined. Process discovery and mapping, understanding the as-is state across the asset, investment, and fund lifecycles. Assessment and diagnostics, identifying the rekeying, manual workarounds, redundant approvals, and data quality issues that exist. Future-state design, redesigning around leading practice, with AI-readiness criteria embedded from the outset. Enablement and implementation, embedding redesigned processes into workflow tools and existing platforms. And finally, continuous improvement and AI readiness, establishing the governance structures and data quality standards that make safe agent orchestration possible. Firms that follow this sequence typically see significant reductions in cycle time before AI is deployed. The foundation delivers value independently, and delivers the confidence that AI, when deployed, will perform as intended and can be defended when it matters.

People Are the Point

Amid the debate about what AI can and cannot do, a simpler truth is in danger of being lost: real estate is fundamentally a people business. Relationships, judgment, trust, and presence in the moments that define outcomes are not features that can be automated, and they are not secondary to the technology strategy. They are the reason the technology strategy matters.

The risk of over-automation is not primarily job displacement. It is brand dilution. If every firm deploys the same AI agents producing the same generic outputs, the differentiation that comes from quality of service, quality of relationship, and quality of judgment disappears. AI should return time to people so they can do what only people can do: exercise judgment, manage exceptions, negotiate, and show up fully in the interactions that define relationships.

The single biggest barrier to realising value from AI in real estate is not technical. It is trust. People do not delegate judgment to systems they cannot understand, supervise, or intervene in. Building that trust requires transparency about what AI does, clear human review points for higher-risk actions, and a design philosophy that keeps people genuinely in control, not nominally so.

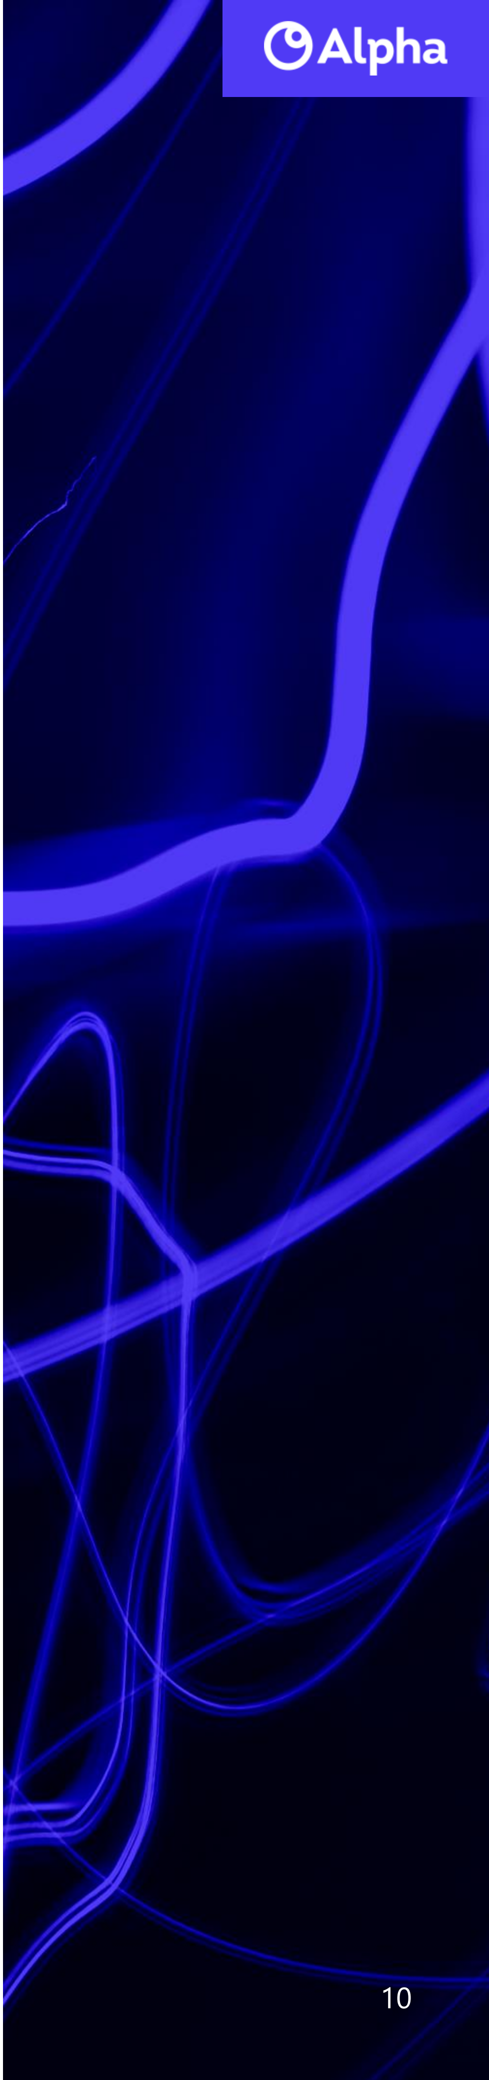

A Framework for Every AI Decision

Leaders evaluating AI initiatives should not be dependent on vendor briefings or technology teams for the fundamental questions. Those questions are not technical. They are strategic, commercial, and governance-oriented, and they have clear answers that should be established before any commitment is made.

Before proceeding with any AI initiative, five questions should be answerable with confidence:

1 Can We Audit the Output?

If an AI system cannot produce a traceable, verifiable explanation of how it reached a conclusion, linked to source data and a defined methodology, it is not deployable in a regulated or high-stakes investment context. This is not a standard to be negotiated. It is a minimum requirement.

2 Do We Understand the Full Cost Model?

Including token usage, agent architecture overhead, development and maintenance costs, and the cost of governance. If the cost model is not fully understood before deployment, it will become clear afterward.

3 Is Our Data Ready?

Fragmented, inconsistent, or unstructured data produces unreliable AI outputs regardless of the quality of the model. Data readiness is not a technical detail, it is a strategic precondition. If the answer is 'not yet,' the investment required to get there is the first investment.

4 Are Our Processes Standardised Enough to Automate?

A process that is unclear, inconsistent, or unmeasured cannot be automated effectively. AI will not fix a broken process, it will scale it. Process standardisation is the prerequisite, not the follow-on.

5 Is This Genuinely an AI Problem?

Or is it an automation problem, a data problem, or a process problem that does not require AI at all? The discipline of asking this question before reaching for AI is not conservatism. It is sound judgment, and it will save significant time and cost.

If you cannot answer all five questions confidently, you are not ready to deploy. You are ready to prepare. That is not a setback. It is the correct starting point, and the firms that treat it as such will deploy AI that works, rather than AI that looked good in a presentation.

The Firms That Will Get This Right

The firms that extract durable value from AI in real estate will not be the ones that moved fastest. They will be the ones that moved in the right sequence, building the foundations that make AI possible before deploying AI itself, applying it selectively where it genuinely adds value, and maintaining the discipline to ask hard questions when the commercial pressure is to simply proceed.

They will have resolved their data fragmentation and standardised their processes. They will have built AI on top of auditable, transparent foundations rather than instead of them. They will understand what their AI costs, what it produces, and why. And they will have kept people, their judgment, their relationships, their accountability, genuinely at the centre of how value is created and protected.

That level of discipline does not happen by accident, and it does not arrive via a vendor pitch. It requires an honest assessment of where you actually are: what your foundations look like, where the genuine gaps are, and what needs to happen before AI becomes an asset rather than a liability.

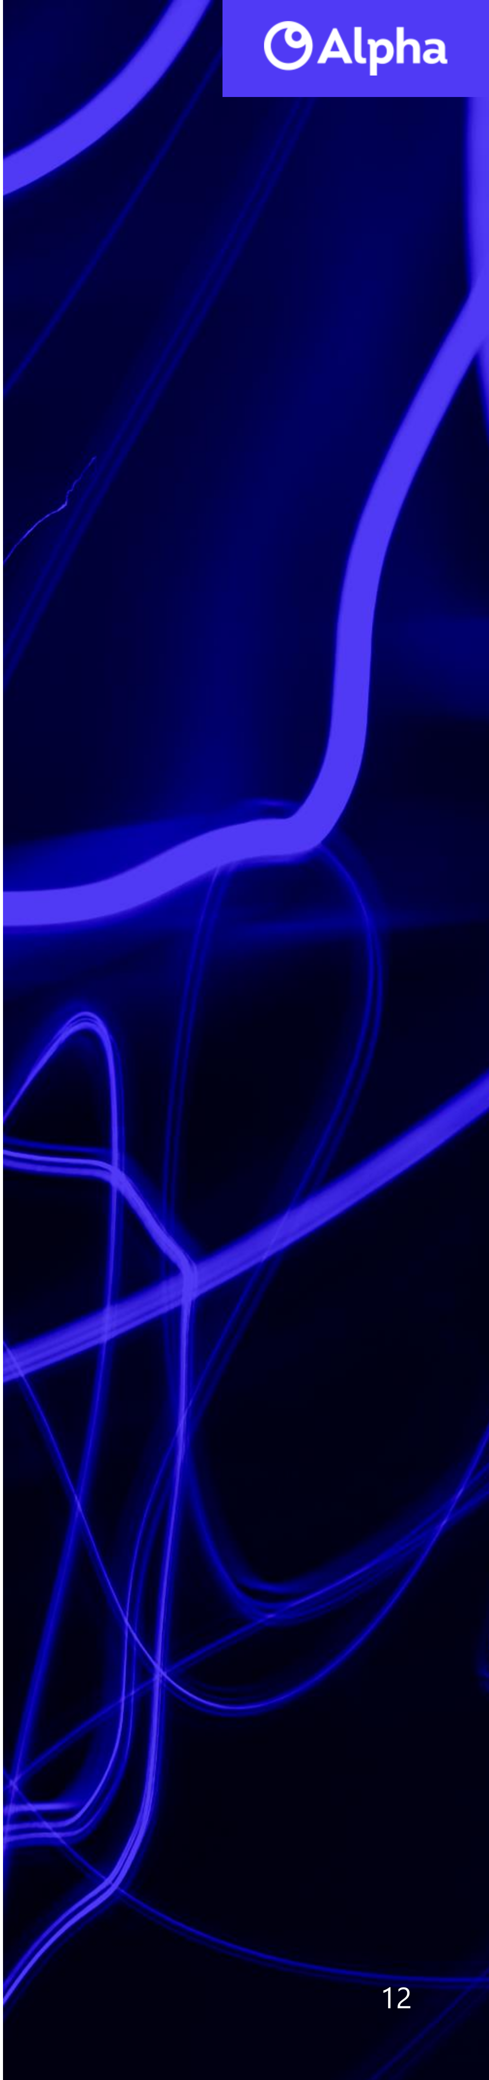

Alpha's AI Readiness Diagnostic

Alpha's AI Readiness Diagnostic is a structured, bounded engagement designed to give you that honest assessment. It is not a technology audit. It is not a vendor evaluation. It is a clear-eyed view of your AI readiness across five dimensions: process, data, technology, governance, and operating model, with a prioritised set of recommendations for what to address first and why.

It is the starting point the firms that get this right begin with. And it is available now.

To discuss the AI Readiness Diagnostic or this paper, contact:

Neil Elliott | Senior Partner

Neil.Elliott@alphafmc.com

Based in London, Neil has more than 25 years of financial services experience, including over two decades advising organisations across real estate and the wider private markets industry. Neil works with senior leadership teams at leading real estate investment managers to define and deliver their strategic transformation agendas. His experience spans operating model design, data and technology strategy, organisational change, system selection and the implementation of enterprise platforms.

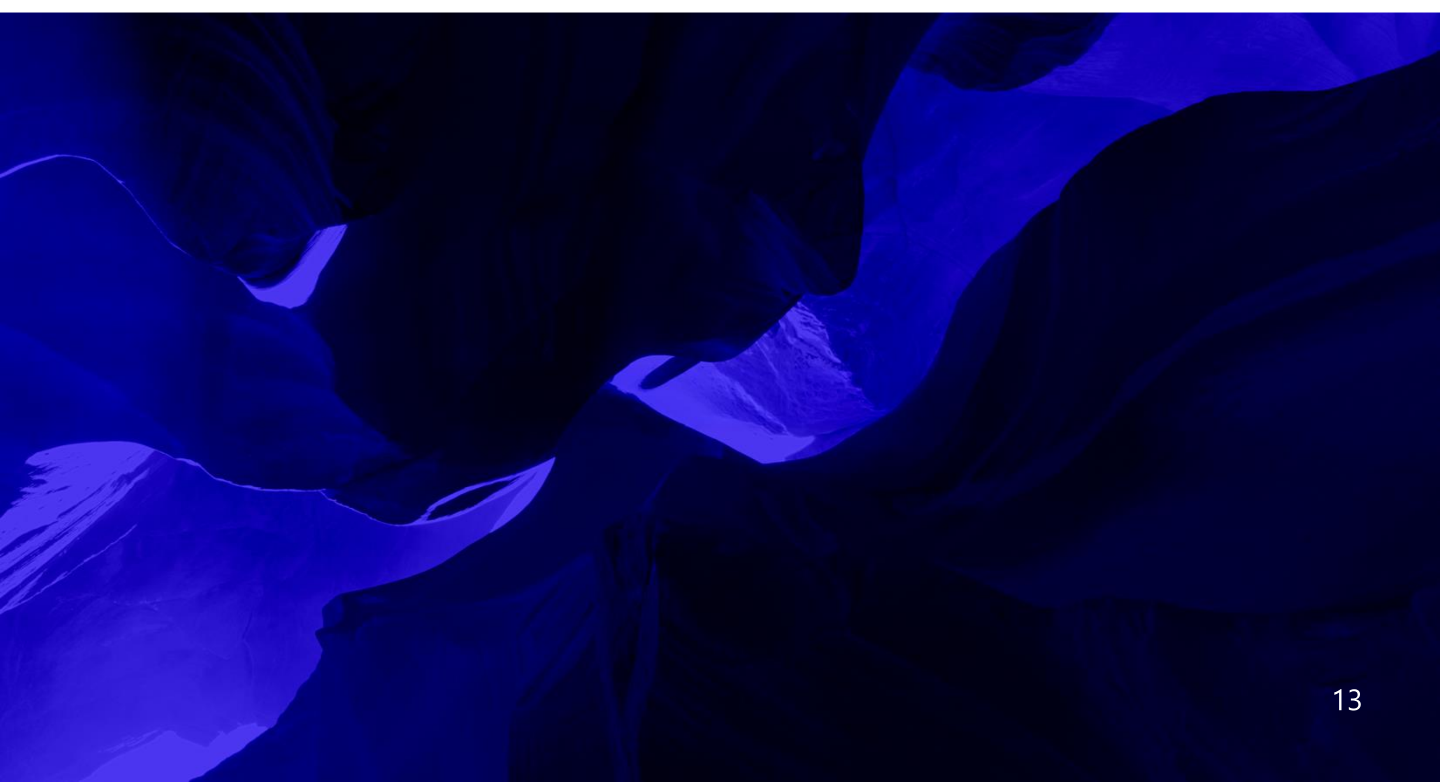

Confidence in the complex

For information or permission to reprint, please contact Alpha FMC at info@alphafmc.com. To find our latest content and insights, please visit alphafmc.com. Follow Alpha FMC on LinkedIn.