

AI in the Investment Office

Maintaining advantage as
the frontier shifts

Authors



Will Coulby
Head of Asset Management
Data Science



Dillon Edwards
Global AI Strategist



Ben Fisher
Global AI Strategist



Mack Kline
Head of Endowments,
Foundations and Healthcare



Michael Smith
Endowments, Foundations
and Healthcare Client Advisor


Foreword

Through our work with institutional investors and through our own experience at J.P. Morgan Asset Management, it is clear the playing field around Artificial Intelligence is leveling. Foundation models, agentic coding tools, and enterprise assistants are broadly accessible, and open-weight models are closing the gap on the frontier. What used to require a large engineering team is now available to any allocator willing to procure. While adopting AI has become easier, building or identifying sources of competitive advantage remains challenging.

This paper is for institutional investors trying to determine where advantage in AI actually lives, and what to do about it. We draw on three sources of information: the patterns we see across asset allocators, our own experience as a large investment organization, and the wider industry conversation. Three patterns stand out:

- 1 What counts as intellectual property (IP) has evolved. Alongside data and models, IP now includes the prompts, skills, context patterns, workflows, evaluations, and accumulated memory that make AI work for your organization. Each of these lives somewhere, either in your own fabric or inside a vendor's platform. Organizations retaining advantage make a deliberate choice about each: own it, control it, or simply access it.
- 2 The center of gravity is shifting from applications to fabric. Applications, the specific AI tools and use cases, remain important, and organizations have built many of them. But the fabric underneath (data readiness, institutional memory, evaluation, governance) is where advantage compounds. Fabric doesn't require building, but control matters. Organizations need to be deliberate about where their fabric lives, and conscious of the risks each choice carries.
- 3 The reward is agility. When your IP is portable, held in formats and locations you control, new AI systems slot in against what you've built rather than replacing it. Organizations with portable IP move with the frontier as it shifts every few months, compound the value of what they've built, and remain resilient when vendors change direction. This kind of agility is a matter of procurement discipline as much as engineering: contracts, data formats, and open standards do most of the work for organizations that make the choice deliberately.

Together, these patterns point to the same argument. Competitive advantage in AI is a question of ownership. What an organization owns and controls above the model layer decides whether advantage compounds over time or decays as the frontier moves. This makes competitive advantage a portfolio question, revisited on the cadence of the frontier rather than the budget cycle. The rest of this paper details how to turn these lessons into repeatable practices in your organization.



Six levers of competitive advantage

Six levers separate organizations retaining durable advantage from those running pilots. Each names a discipline that remains while the stack around it changes.

1 Onboarding and evaluation

To maintain pace with development, institutions need a fast way to trial new AI systems and test them against what success looks like for their own use cases. Only what passes gets integrated with the fabric that connects to their data, their context, and their evaluation setup. The key is having the structural flexibility to swap new capabilities in without giving up what you've built.

2 Data readiness

First- and third-party data needs to be easily exposed to any AI capability, with your calculations and internal ontologies applied. Data, calculations, and ontologies are IP the organization owns. The point is being able to bring them to any AI system on demand, rather than embedding them in one.

3

Institutional memory

The accumulated record of what your organization has decided, why, and what came of it is how we define institutional memory. Every investment office has this, few have made it fully usable to AI systems. Turning it into a live capability, rather than a static archive, requires linking through ontologies, with context, provenance, and permissioning. This is the single largest data asset most investment offices already own and rarely fully use.

4

Safety and evaluation as edge

Trust in deployed AI outputs is maintained through ongoing evaluation across the whole stack: the model, the assistant, and the harnesses, both individually and in combination. The evaluation criteria and test sets an organization develops are themselves IP. Applying them consistently lets you catch drift, measure how proprietary data affects performance, and defend choices to regulators or a board.

5


Governance at scale

While most of the technical building blocks are now available as services, governance cannot be outsourced. Getting governance right is a competitive advantage—it lets an organization disrupt and innovate while maintaining trust. The discipline is balancing speed with consistency at scale, so all teams can progress and take the best from each area, without fragmenting into disconnected tribes. Risk must also be aggregated and viewed across the organization, not only use-case by use-case.

6

Organizational readiness

AI is no longer the responsibility of a specialist central team. All teams need to be comfortable with it, which requires distributing knowledge and investing in training across the organization. Where automation removes the routine work that used to build judgment, training has to replace it deliberately, the way medicine and aviation build judgment when apprenticeship alone cannot. AI fluency is itself a competitive asset. The organizations creating lasting advantage are federating deliberately, treating training as infrastructure rather than an event.



Five practical principles for AI implementation

The six levers describe what compounds. Five principles sit behind them, applying to every institution, every use case.

1 Data quality is foundational

If the underlying data, calculations, and ontologies are incomplete, inconsistent, or poorly governed, the resulting AI output is unlikely to be reliable regardless of how sophisticated the model. Proprietary data is where differentiation lives, but only when it is well-organized, well-permissioned, and well-understood.

2

Evaluation is a discipline, not an event

LLMs can be fluent, confident, and wrong at the same time. That risk does not end when a system is deployed. Every institution needs an ongoing way to check that outputs are correct, that they stay correct as models and data change, and that someone is catching what the system misses. This responsibility sits with the institution, whether the AI was built, bought, partnered, or embedded.

3

Understand where human judgement is non-negotiable

Every AI-assisted decision needs someone accountable, with the authority to override and own the outcome. In investment contexts, fiduciary responsibility does not delegate to a model. Some work is critical because the human is the point: making an investment decision where the reasoning itself is the value, holding a client conversation where trust depends on presence, vetoing at moments when a human check is the safety mechanism. AI can help with this work without replacing the person doing it.

4

Leadership engagement is decisive

When senior leaders actively use AI, sponsor responsible experimentation, and reinforce learning, adoption moves from isolated activity to sustained behavior change. Culture is a pattern of observed behavior more than a communications exercise. When leaders use AI, test new approaches, share lessons, and iterate publicly, the organization develops new habits, and over time those habits become part of the operating model.

5

Progress requires experimentation

Institutions do not move from early testing to enterprise transformation in a single step. Progress comes through repeated cycles of use, feedback, adaptation, and scale. One portfolio manager framed this well—institutions must develop a willingness to allow experimentation, within guardrails. Investment teams determine where AI can enhance judgment without weakening accountability. Operating teams identify where automation improves quality and consistency rather than adding another layer of process. AI maturity is therefore an ongoing capability, developed through disciplined practice, use case by use case, function by function, and lever by lever.

How and where AI creates value across the investment office

To be useful, any AI capability must produce value. For example, extracting key terms from a DDQ is just an output. The desired outcome is a due diligence process that covers more managers, catches red flags earlier, or remembers what a prior fund cycle taught it. Often companies measure the features launched, or the vendor tools integrated, but the outcomes are what change the investment process.

That distinction matters more as AI accelerates. Without it, activity in the investment office rises without the quality of the work rising with it: more summaries, more copilots, more dashboards, and no change in how decisions get made or how confidently they can be defended. Every AI capability an organization

invests in should be justified, before it is procured or built, by the change it is meant to produce in the investment process itself: a decision that is faster, a coverage gap that closes, a risk that becomes visible, a repeated mistake that stops repeating. Levers compound because they carry that change forward across cycles. Tools without a defined change fade with the next release.

With the outcome discipline in mind, we highlight eight functions across the investment office as examples of where AI consistently creates value. The list below names the tools we and others have seen work. It is not exhaustive, but does provide examples of decisions each organization makes against its own investment process.

Across every one of these, the IP accumulated is what turns access to AI into a competitive investment process. The tools are increasingly the same across the industry. The investment process each organization builds around them, outcomes over outputs, is where “AI winners” emerge.

Over the last three years I've been working with our AI team to integrate AI into our investment process. Real progress only came when we switched from building single products to investing in the fabric and the IP we control. I now have tools that give me unparalleled understanding of risk and opportunity. I can curate and adjust them myself. I can go back through decisions I've made, understand where I got them wrong, and make sure I don't make the same mistake twice. Our IP now compounds, irrespective of how the frontier moves.

Leon Eidelman

*J.P. Morgan Asset Management Portfolio Manager
Emerging Markets and Asia Pacific Equities*

Idea generation and macro research	Synthesis of reports, earnings calls and news into daily briefs; thematic screening across filings for emerging risks; sentiment analysis on manager commentary and market data.
Manager due diligence and selection	Extraction of key terms from DDQs and PPMs; red-flag detection in fund documents, side letters and fee structures; performance and attribution benchmarking.
Portfolio construction and risk	Scenario and stress-test copilots; factor exposure and liquidity modeling; rebalancing optimization.
Ongoing manager monitoring	Parsing quarterly letters and capital call notices into structured data; anomaly detection in NAV and performance reporting; LP portal data aggregation across managers.
Legal, compliance and contracts	LPA and side-letter review; MFN clause comparison across fund families; regulatory filing review and compliance checks.
Deal sourcing and pipeline	Identification of co-investment and direct deal opportunities; CRM enrichment from approved internal and external data sources.
Reporting and investor relations	Drafted quarterly performance commentary; meeting summarization; stakeholder Q&A preparation.
Internal knowledge management	Retrieval over memos, correspondence, and prior decisions; linking accumulated experience to current work.

Sourcing

The cost and quality curve for AI has moved sharply. Capabilities that used to require an in-house engineering team (foundation models, enterprise assistants, agent frameworks, memory services, orchestration platforms) are now available as services, and the quality is materially better than what off-the-shelf offered even months ago. You can now get most of the benefit for very little cost, and the floor keeps rising.

For organizations without engineering resources, this levels the playing field. For those with engineering resources, replication remains a valid choice; it just no longer guarantees differentiation on its own.

For all organizations, differentiation now sits in the fabric around the model: the data, the memory, the evaluation, the workflows, the governance, your people and the accumulated IP that make AI work for your organization specifically. This fabric is where investment compounds.


Against this context the framework for choosing your digital strategy is not new:

Buy: Procure a capability from the market, either as-is or shaped to a defined purpose.

Partner: Co-develop with a specialist against your context.

Build: Construct the capability yourself, on your own terms.

And the questions behind them are the same as they have always been: what resources and capabilities do we have, what risks are we prepared to carry, and how quickly do we need to move.




Portfolio discipline and time compression

An organization's digital strategy has always been a blend of approaches across the portfolio: organizations seen as builders still partner or buy services, depending on how they assess the use case. However, the speed at which the market has shifted, and the pace at which capabilities have become cheaper and higher quality to procure, has triggered two key changes in how portfolio discipline is managed.

Firstly, the portfolio must be revisited more frequently as the market evolves. Capabilities that did not exist months ago now exist as services, and vendors bundle and unbundle constantly. What was worth building last year may be worth buying next year. What was safe to buy may need to be brought back inside as the fabric grows.

Secondly, irrespective of the mix, the priority now is maintaining focus on the fabric that holds an organization's IP and competitive advantage. Deciding consciously how the fabric is managed, maintained, and by whom across the portfolio, and thus the risks an organization carries around its own IP, is the core portfolio decision.

And, as stated earlier, the reward for managing the fabric deliberately, and revisiting the mix as the market moves, is agility.



The AI Maturity Diagnostic


The diagnostic that follows helps institutional leaders assess their organization across the six levers and identify where to invest next. It is not a curve to climb. Progress is uneven by design, and organizations typically sit at different points on different levers.

The output is a distribution

An organization's position on each lever, and the gaps between them—that distribution is where the next investment goes. For most organizations, the widest gaps are in the levers that compound most slowly: institutional memory, evaluation, and governance at scale. Building at the frontier is not the same as retaining advantage. What compounds is what the organization owns and controls across the six levers.

Each lever moves through five states: Exploring, Piloting, Workflow-integrated, Enterprise scale, and AI as fabric. The end state is not a single moment. It is the point at which AI stops being a project and becomes how work is done.

Lever	Exploring	Piloting	Workflow-integrated	Enterprise scale	AI as fabric
Onboarding, evaluation, and swap-out	Ad hoc use, driven by individual preference. No consistent evaluation.	Deliberate tests against defined use cases. Some comparison, mostly informal.	AI in daily use. New capabilities evaluated against internal work before adoption. Swap-out possible with some friction.	Deliberate portfolio managed centrally. Consistent evaluation methodology. Swap-out routine.	New capabilities integrated in days, not months. The fabric absorbs frontier improvements without disruption.
Data readiness	Data siloed, permissions ad hoc, calculations inconsistent across teams. Little reaches any AI system.	Use cases connect direct to datasets. Calculations and ontology defined per project, not shared.	Enterprise-wide calculations and ontology defined and permissioned. Both first- and third-party data reachable by approved AI systems.	Data, calculations, and ontology treated as a live capability. Any AI system can be pointed at any approved data source, on demand.	The distinction between data and AI-accessible data has dissolved. All data is AI-reachable by default, with governance built in.
Institutional memory	Historical decisions, memos, and correspondence sit in email, drives, and people's heads. Not usable by AI.	Selected corpora indexed for retrieval on specific workflows. Provenance and permissioning inconsistent.	Institutional memory linked through ontologies, with context and provenance. Retrievable across defined workflows.	First-class capability. AI systems draw on prior decisions and correspondence routinely, with permissioning and audit.	Accumulated experience continuously available to any AI-supported workflow. New decisions feed back into memory automatically.
Safety and evaluation as edge	No formal evaluation. Trust based on vendor claims or individual judgment.	Some testing against known cases. Evaluation happens at deployment, not after.	Ongoing evaluation across workflows. Drift monitoring for critical use cases. Same testing regime across sourcing choices.	Consistent evaluation criteria applied across every AI system in the portfolio. Drift, regression, and behavior shifts tracked continuously. Choices defensible to regulator or board.	Evaluation is continuous, automated where possible, and integral to how AI is deployed. Any AI-supported decision is defensible on its own evidence.
Governance at scale	Individual use, no shared framework. Shadow AI likely widespread and invisible.	Approved-use guidelines, basic data-handling rules, vendor selection standards for pilot workflows.	Governance framework in place, workflow-level accountability defined, human review sampling in production.	Federated model: teams work in their own ways, but shared calculations, definitions, and reference data are preserved. Focus shifts to how new capabilities alter the risk profile. Regulatory alignment and incident response operational.	Governance is embedded, not overlaid. The organization can move fast because guardrails are built in, not added on.
Organizational readiness	Individual pockets of AI use. No shared training. Fluency uneven and personality-driven.	Targeted training for specific teams. Some pockets of AI-native working. Adoption uneven.	AI use is normal across daily work in defined areas. Training is structured. Career pathways beginning to reward fluency.	Federated capability across teams. AI fluency is developed systematically through deliberate training programmes, not incidental learning. Leadership models AI use. Failure tolerance for experimentation is established.	The organization operates as an AI-fluent organization by default. Human judgment sits on top of a widely capable workforce. New hires expect to work AI-native.



How we are applying this ourselves

We have argued for a discipline. It is worth showing how we have run it in our own organization. None of it is finished, the frontier is still moving, and our portfolio changes with it, but the shape of the last three years explains where we now put our weight.

We have always built our AI capabilities inside Spectrum, our proprietary pan-asset-class platform, with a dedicated AI layer we call SpectrumIQ. Historically that meant a common backbone under target products and use cases. The pivot from products on a backbone to the fabric itself as the source of advantage crystallized for us in the first half of 2026, as the landscape matured and both our experience and the wider market sentiment aligned around the same conclusion. What follows describes how the six levers now look in practice.

Onboarding and evaluation

This was the first lever we prioritized, managed by a central team so any new model could be evaluated and onboarded quickly for general use. Model evaluation has now matured to the point where we can onboard new foundation models within days of release. Capability evaluation (assistants, agents, harnesses) moves more slowly, deliberately—the risk profile is more variable and each capability sits differently against the fabric.

Data readiness

We could not build anything useful without LLM-ready data, and the fastest wins were in unstructured data, which gave us proofs-of-concept that became products and made the work pertinent to investors immediately. We created a dedicated data team, and the corpus now stands at more than 40 million documents prepared for AI consumption and research.

Safety and evaluation

This edge followed naturally from data readiness. Evaluation is the core USP of any data science team, and it sits hand-in-glove with building and governance. Automated evaluation scripts run ahead of every deployment, which has been the largest single contributor to our speed to market without loosening our standards.

Governance at scale was prioritized early

We recognized that governance would become the ceiling on how fast we could move, and hired a Head of AI Governance to institutionalize the process before scale exposed the gaps. That decision has made regulatory oversight a design input rather than a retrofit.

Organizational readiness

Notably, this element remains less planned and more emergent. Where opportunity aligned, leadership pushed hard to make it stick. AI days, “shark tank”

events, a shared prompt library, and prompt training have combined to distribute fluency across J.P. Morgan Asset Management rather than concentrating it in a specialist team.

Institutional memory

An area that has been the greatest challenge. We started early with ingesting our internal research, but memory as a pervasive service, woven through context engineering at the organizational level, has needed focus, deep thought and experimentation to perfect. We have learned how it behaves at scale, and are ready to move to enterprise deployment. This is the lever we expect will absorb the largest share of our attention over the coming months.

Looking across the six, the pattern is consistent with the argument of this paper. Where we have moved fastest (data readiness, onboarding, evaluation, governance) we have built into the fabric. Where we are still working (institutional memory) is where the fabric is hardest to build, and where the compounding return will be largest. Portfolio discipline, applied to ourselves.



The winners in AI won't be the organizations that pick the 'right' model. They'll be the ones that build the right foundation. Our strategy has always been to invest in durable capabilities while staying technology agnostic, giving us the flexibility to integrate the best innovations as they emerge. That philosophy has enabled us to continuously evolve Spectrum into a platform built for change, not constrained by it.

Kristian West

*J.P. Morgan Asset Management
Head of Investment Platform*

Conclusion

The AI frontier will keep moving. Organizations retaining advantage are the ones deliberate about what they own above the model layer, and honest about where they sit on each of the six levers. Their portfolios shift as the market shifts, but the fabric underneath (the data, the memory, the evaluation, the workflows, the governance, and the IP their people build around them) is where their deliberate choices compound investment into agility. That is what separates pilots from durable advantage

The diagnostic gives institutional leaders a way to see this in their own organization: where the widest gaps sit, where the next investment goes, and a shared language for the choices that follow. It is meant to open conversation, not close it.

Like the technology, these conversations are not slowing down. We will keep sharing what we are seeing: what is working, and what is not. We hope this is useful in yours, and we would welcome the chance to discuss this further with your teams.

J.P. Morgan Asset Management Use of Large Language Models.

J.P. Morgan Asset Management utilizes Large Language Models (LLMs) internally in an effort to produce a greater level of operational scalability and efficiency across multiple lines of business. LLMs are not relied on to make business/investment decisions. The final business/investment decisions are the responsibility of designated person within JPMAM.

While the intent of LLM is to provide accurate and comprehensive content to internal users, LLM technology may occasionally generate inaccurate, incorrect, incomplete, misleading, or irrelevant information. As a result, LLM output is treated with the highest level of caution and scrutiny by J.P. Morgan Asset Management. LLM generated content is for informational purposes only and is not relied on for final business/investment decisions.

The views contained herein are not to be taken as advice or a recommendation to buy or sell any investment in any jurisdiction, nor is it a commitment from J.P. Morgan Asset Management or any of its subsidiaries to participate in any of the transactions mentioned herein. Any forecasts, figures, opinions or investment techniques and strategies set out are for information purposes only, based on certain assumptions and current market conditions and are subject to change without prior notice. All information presented herein is considered to be accurate at the time of production. This material does not contain sufficient information to support an investment decision, and it should not be relied upon by you in evaluating the merits of investing in any securities or products. In addition, users should

make an independent assessment of the legal, regulatory, tax, credit and accounting implications and determine, together with their own financial professional, if any investment mentioned herein is believed to be appropriate to their personal goals. Investors should ensure that they obtain all available relevant information before making any investment. It should be noted that investments involve risks, the value of investments and the income from them may fluctuate in accordance with market conditions and taxation agreements and investors may not get back the full amount invested. Both past performance and yields are not reliable indicators of current and future results.

To the extent permitted by applicable law, we may record telephone calls and monitor electronic communications to comply with our legal and regulatory obligations and internal policies. Personal data will be collected, stored and processed by J.P. Morgan Asset Management in accordance with our privacy policies at <https://am.jpmorgan.com/global/privacy>.

If you are a person with a disability and need additional support in viewing the material, please call us at 1-800-343-1113 for assistance.

J.P. Morgan Asset Management is the brand name for the asset management business of JPMorgan Chase & Co. and its affiliates worldwide.

© 2026 JPMorgan Chase & Co. All rights reserved

ad456add-9d78-11f1-ae7e-9334b148d692

J.P.Morgan
ASSET MANAGEMENT

FOR INSTITUTIONAL / WHOLESALE / PROFESSIONAL CLIENTS AND QUALIFIED INVESTORS ONLY –
NOT FOR RETAIL USE OR DISTRIBUTION