A-Team Insight
From A-TEAMGROUP



Generative and Agentic AI in Financial Markets

From Pilot to Platform: How Financial Institutions Are Deploying AI at Scale

Commissioned by:

LSEG

A-Team Insight
From A-TEAMGROUP

Table of Contents

Section 1
Introduction — A Market at an Inflection Point 3

Section 2
The State of AI Adoption — Buy Side vs Sell Side 4

Section 3
Buy Side in Focus — Hedge Funds vs Asset Managers 7

Section 4 (Theme 1)
Build vs Buy — Driven by Legacy, Not Strategy 10

Section 5 (Theme 2)
Real-World ROI and Where AI Is Delivering 12

Section 6 (Theme 3)
What Agents Actually Look Like in Practice 15

Section 7 (Theme 4)
The Data Challenge — Data Trust, Access and Alignment 17

Section 8 (Theme 5)
Unlocking Budget and Senior Buy-In 20

Section 9:
Geography — UK vs North America, and the Canada Opportunity 22

Section 10:
Size Matters — How Adoption Shifts by Firm Scale 24

Section 11:
Key Findings and Recommendations 26

About this Research 27

About LSEG 28

Section 1:
Introduction — A Market at an Inflection Point

Artificial intelligence has moved from the margins of financial markets to the centre of strategic planning. In 2026, the question is no longer whether firms should invest in AI — it is how quickly they can move from experimentation to operational deployment, and what infrastructure they need to get there.

This white paper presents findings from an A-Team Group research programme exploring Generative and Agentic AI in Financial Markets, commissioned by LSEG, and based on a survey of 200 senior AI decision-makers at buy-side and sell-side institutions across the USA, UK and Canada. Almost four-fifths (79%) of respondents are the primary AI decision-maker with budget authority at their organisation. The survey was supplemented by practitioner interviews with senior AI and data leaders, and by intelligence gathered from A-Team Group's ongoing programme of industry conferences, webinars and panel discussions.

The findings reveal a market in transition. Most firms have moved past the question of whether AI is worth pursuing (74% back first-mover advantage) but the majority are still spending the bulk of their AI budget on desktop productivity tools rather than structural, agent-driven transformation. The average firm allocates 57% of its AI budget to copilots and AI assistants, leaving 43% for agentic infrastructure. A small but growing cohort of institutions at the leading edge have already inverted that ratio, and their results — in cost reduction, capacity creation and alpha generation — are reshaping what the rest of the market believes is possible.

The picture is not uniform. Sell-side firms are moving faster than the buy side on agent deployment. Hedge funds are materially more advanced than asset managers. Canada is outpacing the UK. And across every segment, data readiness — not AI model capability — is the foremost concern. The firms that are scaling AI successfully are those that have made their data infrastructure AI-ready early, rather than those who have focused on adopting the most sophisticated models.

This white paper examines what the data shows, what practitioners are experiencing on the ground, and what the findings mean for institutions seeking to accelerate their AI programmes.

Section 2:
The State of AI Adoption — Buy Side vs Sell Side

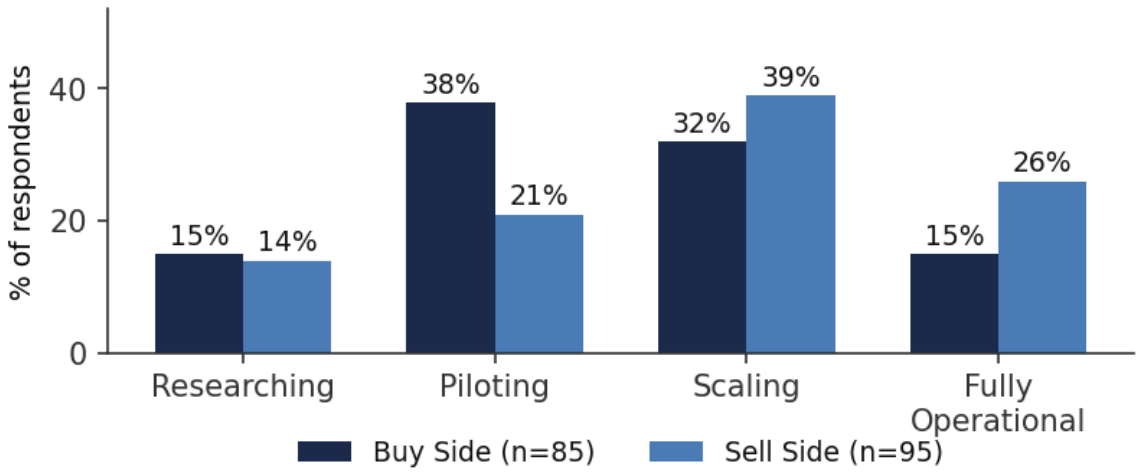
The headline finding from this research is one of divergence. At the broadest level, buy-side and sell-side institutions are pursuing AI with similar conviction. Most institutions back the first-mover advantage with 71% of buy-side and 76% of sell-side firms say the risk is worth taking, but the speed and nature of that pursuit differ considerably between the two segments.

Adoption Stage

With 65% of sell-side firms at scaling or fully operational stages on agentic AI, compared with 47% of buy-side firms. The sell side is moving faster across every adoption metric: 52% report their AI roadmap has changed significantly because of AI, versus 41% of buy-side respondents. 46% of sell-side firms are targeting their AI goals within six months, compared with 36% of the buy side.

This gap reflects structural differences between the two segments. Sell-side institutions, investment banks and brokerages are under intense margin pressure, and AI is being deployed primarily as an efficiency tool: cost reduction, workflow automation and operational capacity creation. The commercial mandate is clear and measurable. Buy-side firms, by contrast, are pursuing AI for revenue generation and alpha, which is a longer-horizon, harder-to-quantify objective that tends to produce more cautious deployment timelines.

Figure 1: *Agent adoption stage — % of firms (Buy Side n=85, Sell Side n=95)*

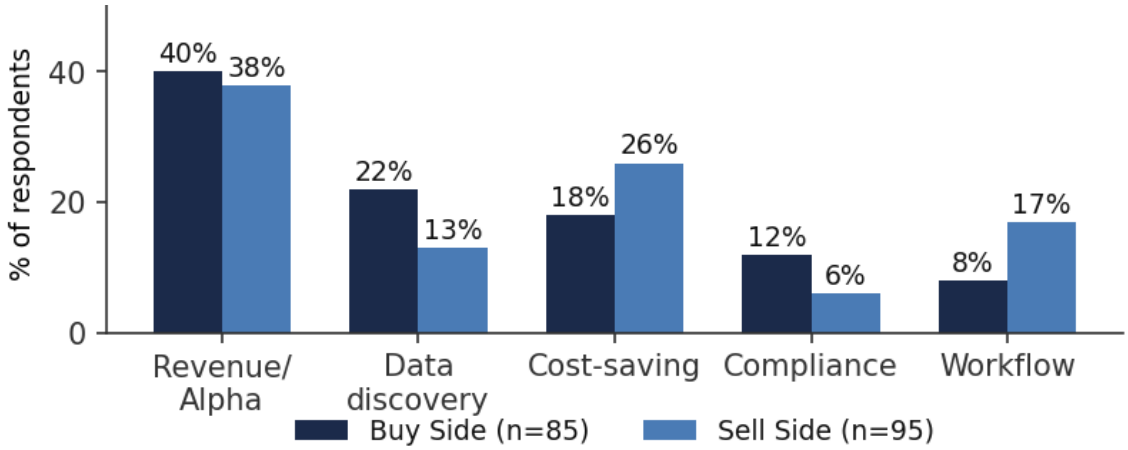


AI Philosophy

Revenue and alpha generation are the primary AI philosophy for both segments with 40% of buy-side and 38% of sell-side firms leading with this rationale, but the divergence begins immediately afterward. Sell-side firms strongly prioritise cost-saving (26%) as a secondary rationale, alongside workflow automation (17%). Buy-side firms diversify into data discovery (22%) and compliance (12%), reflecting the fiduciary obligations and the primacy of research quality in their business model.

This philosophical divide has a practical implication for how each segment is spending its AI budget. The 57% average desktop allocation (the proportion of AI budgets going to copilots and productivity tools) is virtually identical across both segments. But the underlying mix is different. Buy-side desktop spend is concentrated on research and alpha generation tools. Sell-side desktop spend skews toward workflow automation and cost-out applications. The same headline figure reflects fundamentally different strategic priorities.

Figure 2: *Primary AI philosophy — % of firms (Buy Side n=85, Sell Side n=95)*



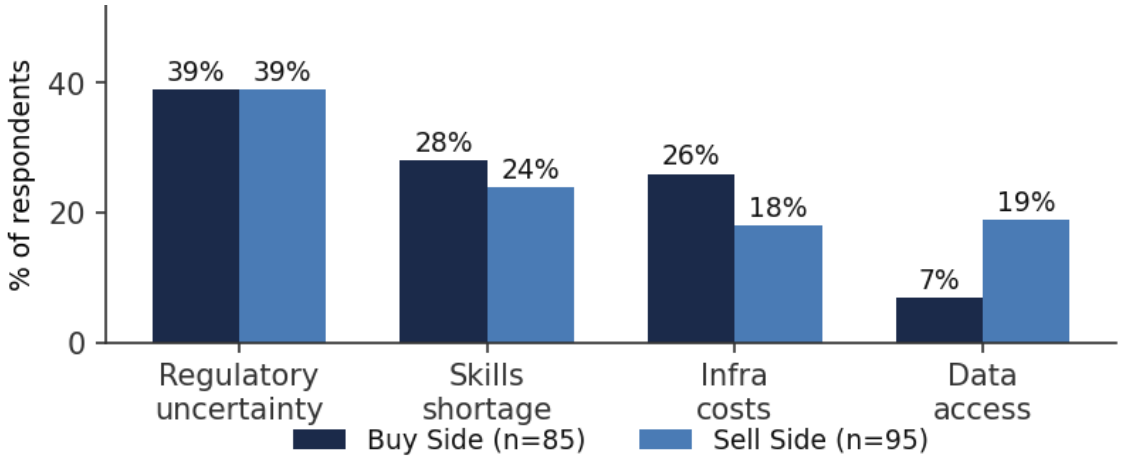
Build vs Buy

The build approach also diverges sharply, with 70% of buy-side firms preferring hybrid or proprietary builds, prioritising control over their data and algorithms. Sell-side firms lean more heavily on vendor platforms (41% select this as their primary approach), driven by commercial urgency and the need for speed-to-deploy. The distinction is not primarily about cost, it is about sovereignty. Buy-side firms, particularly quantitative strategies, are acutely aware of the risks of vendor model drift, deprecation and proprietary data exposure through third-party training infrastructure.

Barriers

Both segments cite regulatory uncertainty as their primary barrier to scaling AI (39% buy side, 39% sell side) but the secondary barriers differ. Buy-side firms are more constrained by skills shortage (28%), while sell-side firms uniquely stress data access (19% vs 7% of buy side). This is a structural difference. Sell-side institutions operating at scale have extensive external data pipelines and face genuine access and integration constraints while buy-side firms have the data but lack the skills to deploy it effectively.

Figure 3: *Primary barriers to scaling AI — % of firms (Buy Side n=85, Sell Side n=95)*



Data Quality

Data trust is the dominant quality concern for both segments, cited by 47% of buy-side and 43% of sell-side firms as their primary data quality challenge. AI reliability and data provenance are the shared critical dependency for scaled deployment in regulated markets. Firms cannot deploy AI reliably without first resolving the question of whether they can trust the data feeding their models.

Model Context Protocol (MCP) and Connectivity

The majority of firms across both segments plan to use the MCP, for both internal and external data connectivity, with 74% of buy-side and 67% of sell-side firms confirming this intent. This convergence on open agentic infrastructure standards is one of the strongest structural signals in the dataset. It suggests the industry is moving toward a common connectivity architecture for AI, and that firms planning for the medium term are factoring this into their data strategy.

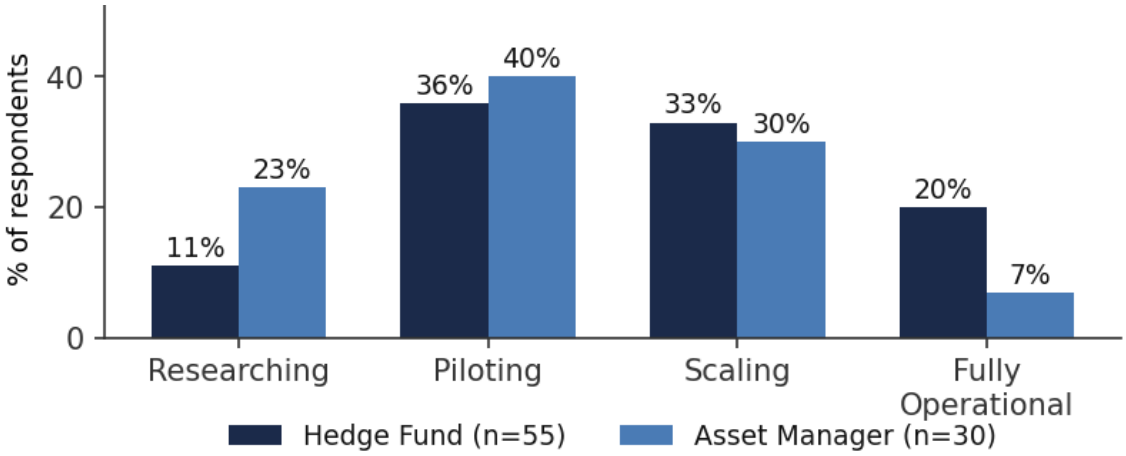
Section 3:
Buy Side in Focus — Hedge Funds vs Asset Managers

The buy side is not a single market. When respondents are segmented by investment strategy — grouping Quantitative, Macro and Multistrategy funds as Hedge Funds (n=55) and Discretionary and Passive managers as Asset Managers (n=30) — two fundamentally different AI trajectories emerge within the same broad category.

Adoption Stage

Hedge funds are materially ahead of asset managers on agent deployment, with 53% at Scaling or Fully Operational compared with 37% of asset managers. Hedge funds often move faster, encounter fewer capacity constraints, and have lighter legacy estates to work around. The absence of multi-decade technology debt is itself a competitive asset. It enables AI-native builds that asset managers, carrying the weight of larger operational estates, may struggle to replicate.

Figure 4: *Agent adoption stage — Hedge Fund vs Asset Manager*



AI Philosophy

Hedge funds lead strongly with Revenue/Alpha (47%) and Data discovery (24%). Asset managers are more fragmented: Revenue/Alpha (27%), Cost-saving (23%), Compliance (20%). This distribution reflects the asset manager's more complex mandate; managing risk alongside returns, with fiduciary obligations that make compliance a primary rather than secondary consideration in AI deployment decisions.

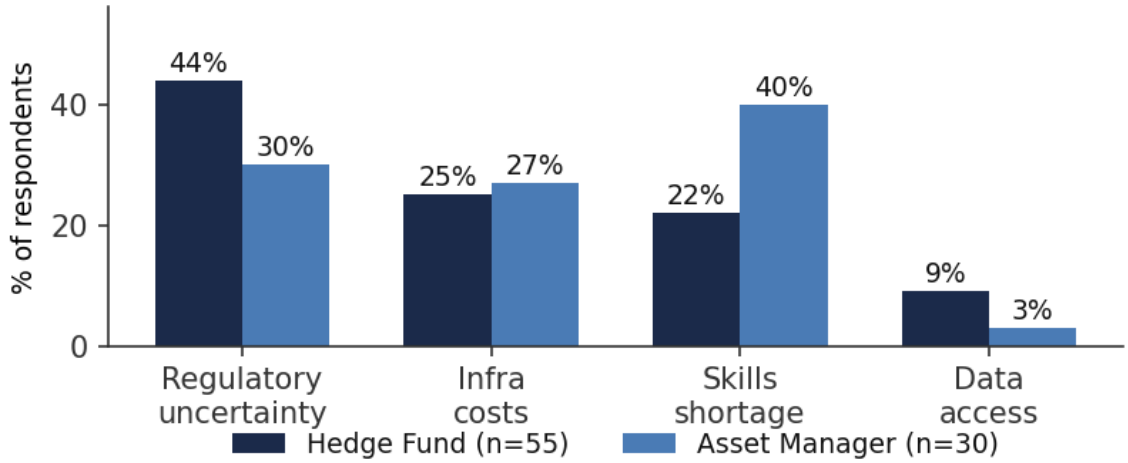
The revenue impact data reinforces this picture. Asset managers score Research/Alpha at 4.00 out of 5 — the highest single buy-side score in the dataset. This is the defining AI use case for asset managers: using AI to improve the quality and speed of research, and to generate insights that their investment teams cannot produce manually. Hedge funds score Research/Alpha at 3.78 — still high, but reflecting a more diversified deployment across compliance, data discovery and execution.

Challenges

The starkest difference between the two sub-segments is the nature of their primary challenges. Skills shortage is the primary barrier for asset managers (40%), at nearly double the hedge fund rate (22%). Regulatory uncertainty is the hedge fund's primary constraint (44%). These are fundamentally different problems requiring fundamentally different vendor conversations. A technology solution that leads with talent augmentation and ease of deployment will resonate with an asset manager audience. The same message is less compelling for a quantitative fund whose constraint is regulatory approval rather than human capacity.

A-Team Insight practitioner interviews with senior figures at leading buy-side institutions reinforce this picture. The head of software development at a mid-size discretionary asset manager described capacity (not budget, not technical capability) as the single most important constraint on their AI programme. The CEO of the AI division at a quantitative buy-side fund, by contrast, described a 90-day board approval window and an aggressive 20% cost reduction and 25% revenue uplift target, with minimal discussion of talent constraints. The same category, different worlds.

Figure 5: *Primary AI barriers — Hedge Fund vs Asset Manager*



Desktop Budget

Desktop budget allocation is virtually identical between the two sub-segments: hedge funds at 56.7% and asset managers at 57.3%. The headline masks a qualitative difference, however. Hedge fund desktop tools are predominantly research and signal-generation focused. Asset manager desktop tools include a higher proportion of compliance and reporting applications. The similar budget figure reflects different underlying use cases.



Section 4 (Theme 1): Build vs. Buy — Driven by Legacy, Not Strategy

The decision to build proprietary AI tools, use vendor platforms or adopt a hybrid approach is frequently framed as a strategic choice. In practice, it is more often an infrastructure choice, driven primarily by the state of a firm's legacy technology estate.

Legacy as the Deciding Factor

Practitioners at lean buy-side firms are explicit: the absence of legacy technology is the most important driver of their build approach, not strategic preference. Firms with minimal technical debt can build AI-native from the ground up: faster, cleaner, without integration overhead. Firms with extensive legacy estates are effectively forced into hybrid approaches, not because they prefer vendor platforms but because their existing infrastructure constrains their options.

A-Team Insight interviews with practitioners at buy-side firms of different scales illustrate the spectrum. A mid-size discretionary asset manager with minimal legacy described an AI-native build as the only practical option: the legacy estate was an asset, not a burden, precisely because there was so little of it. A Tier 1 global investment bank's equities division, carrying decades of legacy infrastructure, described a hybrid approach as the only realistic path, with significant governance and compliance overhead accompanying every deployment decision.

AI Sovereignty

For buy-side firms, particularly quantitative strategies, the build preference is primarily about control, rather than cost. Practitioners report strong wariness of vendor model drift, deprecation cycles and the risk that proprietary algorithms or trading signals could be exposed through third-party training infrastructure. Private AI on fixed, internally controlled frameworks is an explicit sovereignty decision. The concern is not whether vendors are trustworthy today, but whether the relationship can be maintained reliably over the multi-year timescales that matter for quantitative strategies.

Selective LLM Restraint

One of the more striking themes from A-Team Insight's practitioner interviews is the emergence of active efforts to limit the use of large language models. Senior data leaders at multiple Tier 1 institutions described deliberate programmes to reduce or redirect LLM deployment, citing energy consumption, hallucination risk

in client-facing applications, and what one described as an ego-driven and immature phase of adoption in which teams rushed to apply LLMs to problems they were poorly suited to solve. Several practitioners described reverting to simpler, more reliable Python-based solutions after LLM failures in production, including at least one case involving a client-facing chatbot that was replaced after providing incorrect answers.

The emerging consensus among the most advanced practitioners is that LLMs are powerful but not universally applicable, and that the most effective AI deployments combine LLMs with purpose-built tools in a selective, workflow-appropriate manner. This is a meaningful shift from the 2024 instinct to apply LLMs broadly and quickly.

Democratised Access at Scale

At the other end of the scale, one major global institution interviewed by A-Team reports that 94% of employees have access to an AI agent workbench, tens of thousands have been trained to create and test agents, and anyone can build in a sandbox environment, with a governance board controlling what reaches production. This is a fundamentally different model to the centralised AI teams that remain the norm at most institutions. It treats AI capability as an organisation-wide competency rather than a specialist function, and the speed and scale of deployment that results are measurably ahead of peer institutions that have not made the same democratisation investment.

Survey Data: Build vs Buy

Surveying the market as a whole, 70% of buy-side firms prefer hybrid or proprietary builds. On the sell side, 41% rely on vendor platforms as their primary approach, with the remainder split between hybrid (25%) and proprietary build (32%). Faster insight into the value of AI tools is the most frequently cited partner evaluation criterion for both segments (cited by 67% of buy-side and 71% of sell-side firms) ahead of cost reduction (54% buy side, 63% sell side).

Section 5 (Theme 2): Real-World ROI and Where AI Is Delivering

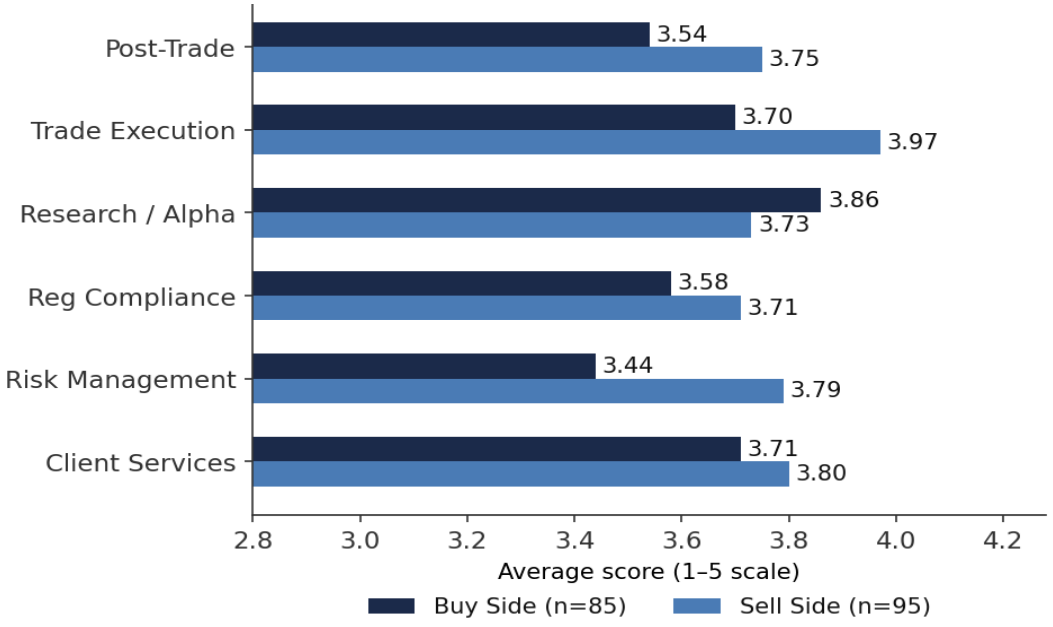
The question of where AI is generating measurable return on investment is the most practically important in this research. The survey data provides impact scores across six business areas; the practitioner interviews ground those scores in real operational examples; and the evidence from both sources points toward a consistent set of conclusions about where AI is working, where it is not, and what determines the difference.

Revenue Impact Scores

Respondents were asked to rate the expected revenue impact of AI across six areas on a one-to-five scale. For the sell side, Trade Execution scores highest at 3.97, reflecting the direct, measurable profit and loss impact of AI optimisation in high-volume execution flows. For the buy side, Research/Alpha scores highest at 3.86 overall and 4.00 among asset managers specifically, the highest single buy-side score in the dataset. The sell side also rates Risk Management highly at 3.79, consistent with the systemic risk monitoring priorities of larger institutions.

The largest gap between the two segments is Trade Execution: sell side 3.97 versus buy side 3.70, a difference of 0.27 points. This reflects the structural difference in how the two segments use trading infrastructure. For sell-side institutions processing large volumes of orders across multiple venues, AI optimisation of execution has immediate and quantifiable impact. For buy-side firms, whose primary value creation is in investment decision-making rather than execution, the same use case is less central.

Figure 6: *Expected revenue impact by area — avg 1–5 score (Buy Side n=85, Sell Side n=95)*



Where AI Is Delivering — Practitioner Evidence

The practitioner interviews provide several concrete examples of AI delivering measurable operational value at scale. In the consumer banking division of a Tier 1 global investment bank, three collaborative AI agents (one focused on product, one on client profiles, one on communications) reduced the time to generate targeted email campaigns from twelve weeks to seven minutes. This is classified internally as an efficiency play rather than a revenue driver, but the capacity created is significant.

At a major global institution, AI agents are deployed as what practitioners describe as digital employees, assigned to the corporate directory with manager relationships, performance reviews and the ability to be let go if they underperform. These agents self-assign to engineering tickets, fix code autonomously within defined parameters, and report to their assigned manager for approval before changes reach production. This represents a genuinely novel deployment model that sits between traditional automation and full autonomy.

In the operations division of a Tier 1 global investment bank, a smart workbench built on LLMs reads unstructured documents (PDFs, handwritten notes, multilingual documents including Chinese and French) for a team of several thousand staff. Previously done manually, this processing is now scaled across jurisdictions, with the LLM enabling multi-language capability that was previously impractical at the required scale.

Quantified ROI Benchmarks

The practitioner interviews also reveal quantified ROI targets that illuminate how different types of institutions approach the business case for AI investment. The CEO of the AI division at a quantitative buy-side fund described a 20% cost reduction and 25% revenue uplift target, with a 90-day board approval window, a commercially aggressive posture reflecting the fund's conviction that AI is a direct alpha generation tool. The Head of Equity Capability at a Tier 1 global investment bank described an 18% cost target and 12% revenue target, with a nine-month board approval window: more conservative, reflecting a larger institution's governance requirements. The AI Change Lead at the operations division of a Tier 1 global bank described a cost-out mandate with a 10% ROI target and no explicit revenue uplift objective, consistent with a workflow efficiency rationale.

Insights from A-Team's events and webinars suggest a broader market shift in ROI timelines: AI investments are now being questioned within twelve months, far faster than equivalent cloud or data infrastructure investments, which were typically evaluated over two to three years. Zero-based budgeting for AI and data functions, requiring each line item to be justified from scratch rather than carried forward from the previous year, is already live at multiple Tier 1 institutions in 2026.

Where AI Is Still Falling Short

The most significant gap between impact expectation and deployment reality is autonomous trade execution. Trade Execution scores 3.97 for expected revenue impact on the sell side (the highest score in the dataset) but practitioners consistently describe deployment likelihood in this area as very low. Liability, market impact risk and regulatory requirements keep humans in the loop for execution decisions, even where the potential value of AI optimisation is well understood.

Data mesh failures and cross-functional misalignment are the most common operational reasons that AI pilots fail to scale. When data from different business functions uses inconsistent identifiers, timestamps or granularities, AI models cannot combine them effectively, and the problem is organisational, not technological. Several practitioners described pilots that delivered strong early results but stalled when they encountered the cross-departmental data alignment challenges that had not been apparent at the pilot stage.

Client-facing LLM applications represent a third area where the gap between expectation and reality is significant. At least one senior practitioner described a client-facing LLM chatbot that began providing incorrect answers and was subsequently replaced with a Python-based solution. The lesson drawn was not that AI cannot be used in client-facing contexts, but that LLMs are unsuitable for applications where the cost of a wrong answer is borne by a client rather than an internal team.

Section 6 (Theme 3): What Agents Actually Look Like in Practice

The term agentic AI covers a wide spectrum of deployment models, from fully supervised AI that makes recommendations for human review through to fully automated systems operating without human intervention. Understanding where firms are actually operating on this spectrum, and where they are deliberately choosing not to deploy is one of the most practically useful findings from this research.

The Autonomy Slider

A-Team Insight uses a framework of three deployment positions. In the Loop describes deployments where humans take all final decisions and AI assists. On the Loop describes deployments where AI acts autonomously within defined parameters, with humans reviewing by exception. Out of the Loop describes fully automated deployments for specific low-risk, high-volume processes.

The survey asked respondents to rank five business areas by their likelihood of autonomous agent deployment. The ranking data across all 200 respondents reveals a clear hierarchy of preference and avoidance.

Where Agents Are Prioritised

Regulatory Compliance is consistently the most favoured area for autonomous agent deployment, ranked first or second by 65% of buy-side and 49% of sell-side respondents. Practitioners confirm that compliance use cases are the most palatable for governance teams precisely because they are rules-based, auditable and bounded in their consequences. An AI agent that monitors transactions for compliance purposes can have its outputs reviewed by exception and its reasoning reconstructed from explicit rule sets. This is significantly easier to govern than an agent making investment or client-facing decisions.

Risk Management is the second priority for buy-side firms, ranked first most frequently by 29% of buy-side respondents. Sell-side firms show a more distributed set of priorities, with Research/ Alpha tied for first place ranking at 20%.

Where Agents Are Avoided

Client Services is consistently ranked last, most likely to receive a rank-5 (least likely) designation from both buy-side (28% rank it last) and sell-side (23%) respondents. Practitioner evidence is consistent with this data: client-facing AI carries reputational risk that firms are not prepared to accept, and the LLM rollback example described in the previous section has reinforced institutional caution about this area.

Trade Execution presents the most interesting pattern in the data: it scores both first (buy side 24%, sell side 23%) and last (buy side 20%, sell side 22%) most frequently. Firms are either fully committed to autonomous execution or actively avoiding it. There is almost no middle ground, and the determining factor appears to be the nature of the firm's trading activity. High-volume, low-value execution flows are natural candidates for automation; discretionary, high-value decisions are not.

The Digital Employee Model

The most advanced deployment model identified in A-Team Insight's practitioner research is what one major global institution describes as digital employees. AI agents are assigned to the corporate directory with a manager relationship, a performance review cycle and the ability to be let go if they underperform. They self-assign to engineering tickets, autonomously resolve defined classes of issues, and submit their proposed changes for human approval before production deployment. This is an On the Loop model: the agent acts, the manager reviews, but at a scale and with an organisational integration that goes well beyond typical enterprise AI deployments. Tens of thousands of employees at this institution are trained to create and test their own agents within a sandbox environment governed by a central review board.

The practitioner evidence is that this model, democratised access to agent creation, combined with centralised governance of production deployment, is the key structural differentiator between firms that are scaling AI and those that remain in pilot mode.

Section 7 (Theme 4): The Data Challenge — Data Trust, Access and Alignment

The single most consistent finding across both the survey data and the practitioner interviews is that data readiness is the defining enabler of AI deployment at scale. Firms that have invested in high-quality, AI-ready data infrastructure are scaling faster, deploying more confidently and demonstrating stronger returns than those that have not. The technology to build sophisticated AI models is widely available — the competitive advantage now lies in the quality, governance and connectivity of the data that feeds them.

A-Team Insight identifies three distinct data problems that firms are navigating, each requiring a different solution.

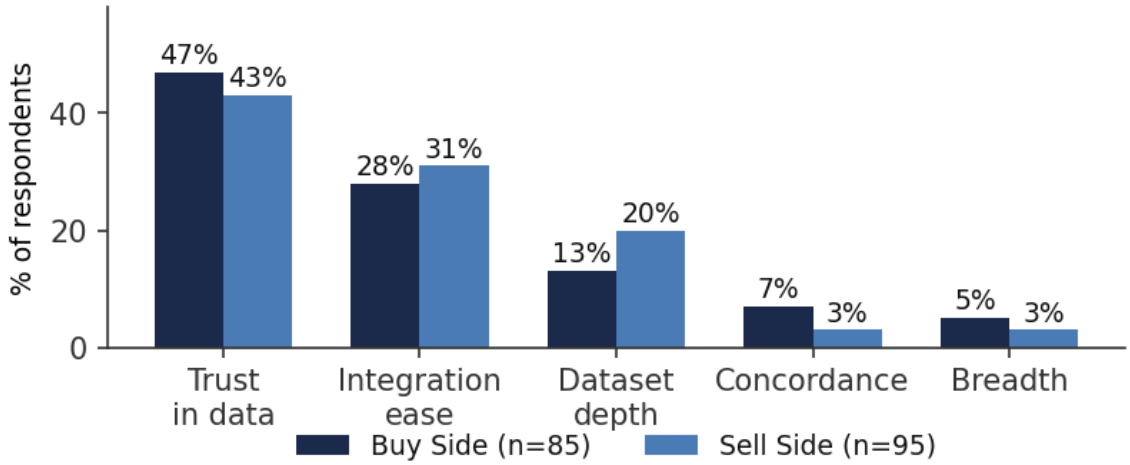
Data Trust

Data trust is the primary quality concern for both segments: 47% of buy-side and 43% of sell-side firms cite it as their principal data quality challenge. This is about provenance, lineage and reliability: the ability to know where a piece of data came from, how it has been transformed, and whether it can be relied upon for a decision that may subsequently need to be reconstructed for regulatory purposes.

Under regulatory frameworks in the UK, firms are expected to maintain appropriate data governance, accountability and oversight when deploying AI systems. AI models that cannot be interrogated at the level of their input data are not deployable in regulated contexts. Practitioners at senior levels describe building semantic layer wrappers, Retrieval-Augmented Generation (RAG) systems and knowledge graphs specifically to assess and document metadata quality before allowing that data to be used in AI-driven processes.

The most effective approach identified in A-Team Insight's practitioner research is to extend existing trusted vendor relationships into the AI layer, rather than establishing new data relationships at the AI level. One Tier 1 global institution described choosing its AI data integration partners from its existing, established data vendor relationships precisely because trust takes years to build and cannot be shortcut through procurement processes.

Figure 7: *Primary data quality concern — % of firms (Buy Side n=85, Sell Side n=95)*



Access and Integration

Data access and integration is the second top concern – the challenge of whether data exists at the required granularity and whether it can be connected to AI infrastructure without prohibitive engineering effort. This is a technology and infrastructure problem, more directly solvable than data trust through tooling and open standards.

The clearest signal of market direction is the MCP adoption data: 71% of buy-side and 67% of sell-side firms plan to use the Model Context Protocol for both internal and external data connectivity. MCP is converging rapidly on a market standard for AI data connectivity, and firms planning their medium-term data infrastructure are treating it as a given.

Practitioner interviews at Tier 1 global banks reveal another dimension of the access challenge: fragmented data mesh conditions in which transaction data held by different business units is not in a consistent enough condition for agent queries to work reliably across the organisation. One practitioner described agents that functioned correctly when working with data from a single business unit failing entirely when asked to combine data from multiple units with different identifiers and timestamps.

Cross-Functional Internal Alignment

The third and most structurally difficult data challenge is internal misalignment: the situation where data from different business functions uses different identifiers, timestamps, granularities and conventions, making it impossible for AI models to combine those data sources without significant manual intervention. This is an organisational problem. Technology alone cannot solve it. It requires business

function negotiation, agreed data standards and, in some cases, structural organisational change.

A-Team Insight webinar participants from Tier 1 banking institutions have been explicit on this point: AI deployment forces a reckoning with data governance questions that have been deferred for years. Every attribute is now potentially reachable by AI, which means the informal inconsistencies that were manageable when data was accessed manually by human analysts become systemic blockers when AI models attempt to process data at scale.

UK vs North America — Inverted Problems

The geography of the data challenge adds another layer of nuance. UK firms rank integration ease as their primary data quality concern (estimated 31% of respondents), suggesting cleaner underlying data but fragmented legacy systems that make connection difficult. North American firms rank data trust as their primary concern (estimated 47%), suggesting more connected data pipelines but lower confidence in the provenance and reliability of what is flowing through them. These are genuinely different problems requiring different vendor solutions and different investment priorities.

Section 8 (Theme 5): Unlocking Budget and Senior Buy-In

The barriers to scaling AI at most institutions are less often technical than organisational. The question of how to secure and maintain senior leadership support for AI investment, and what proof points move the CFO conversation from scepticism to approval, is one of the most practically useful findings from this research.

Match the ROI Window

The single most important lesson from the practitioner interviews is that the ROI window (the timescale within which returns must be demonstrated to maintain budget) varies enormously between institution types, and that failing to match your pitch to the specific window of the institution you are addressing, is the most common reason that otherwise compelling AI business cases fail.

A-Team Insight interviews identify three distinct ROI windows from three practitioner profiles. The CEO of the AI division at a quantitative buy-side fund described a 90-day board approval window: a commercially aggressive posture reflecting the fund's conviction that AI has immediate, measurable impact. The Head of Equity Capability at a Tier 1 global investment bank described a nine-month window: longer, reflecting larger governance requirements but still within a twelve-month commercial horizon. The Digital Transformation Lead at a major sell-side institution described a three-to-five-year window, treating AI investment as long-term strategic infrastructure rather than a near-term commercial lever.

The broader market context for these windows is shifting rapidly. A-Team Insight research confirms that AI investments are now being questioned within twelve months. Far faster than the two-to-three-year evaluation horizon that applied to equivalent cloud investments. Zero-based budgeting for AI and data offices, requiring every line item to be justified from scratch each year is already live at multiple Tier 1 institutions in 2026.

Frame as Legacy Displacement

One of the most effective CFO framings identified in practitioner interviews is legacy displacement: rather than presenting AI

as a new cost, presenting it as a replacement for existing legacy technology maintenance expenditure. Practitioners at global institutions describe redirecting 70 to 80% of their AI spending from legacy technology maintenance budgets. The CFO conversation shifts from requesting new investment to reallocating existing spend, a structurally different and typically easier conversation. A-Team Insight research participants confirmed that AI investment must increasingly come from existing envelopes, as BAU budgets remain flat and AI headcount growth must be self-funded from productivity gains.

Use a Value-Effort Framework

A practitioner at a Tier 1 global institution described using a two-by-two effort-versus-value matrix to prioritise AI investments and build the internal business case. High-effort, high-value use cases, complex modelling and balance sheet optimisation, earn long-term budget. High-effort, low-value use cases are cut. Low-effort, meaningful-value use cases like meeting summaries and code suggestions, are prioritised for early deployment because they build cultural buy-in across the organisation at low cost. The framework is simple, but it provides a repeatable mechanism for justifying each AI investment in language that finance and operations leadership can engage with.

Train the Organisation

A-Team Insight practitioner evidence is unambiguous on one point: firms that democratise AI access across the whole organisation move faster, sustain buy-in more effectively, and build the cultural foundation for scaling that centralised AI teams cannot replicate. One major global institution trained tens of thousands of employees to create and test agents. The cultural change that resulted: from AI as something that happens to employees, to AI as something that employees participate in creating, is described by practitioners as the primary differentiator between institutions that are scaling AI and those stuck in extended pilot phases.

Peer Proof Points

The survey data on partner evaluation criteria confirms what practitioners report directly: peer proof points (specific, named examples from comparable institutions) carry more weight in the internal buy-in conversation than any benchmark study or generic ROI claim. When a named competitor announces a live AI capability that has produced measurable results, the internal approval cycle shortens. A-Team Insight events consistently confirm that visibility of what peers are actually doing, not what analysts say is theoretically possible, is what moves the internal conversation from scepticism to urgency.

Section 9:
Geography — UK vs. North America, and the Canada Opportunity

UK vs. North America

The most striking geographic divergence in the dataset is between the UK and North America on first-mover conviction. UK firms back first-mover advantage at a rate of 64%, compared with 79% of US and Canadian firms. Among UK firms, 23% say the risk is not worth it, compared with 13% in North America. This 15-point gap in first-mover conviction reflects a materially different risk appetite that is translating into measurably different adoption pace.

On agentic AI, 58% of US and Canadian respondents are Scaling or Fully Operational, compared with 54% of UK respondents. Among US and Canadian firms, 12% are targeting their AI goals within three months, compared with 6% in the UK. Combined with the first-mover conviction gap, these figures suggest the adoption lead could widen as UK firms continue to work through how existing regulatory requirements apply to AI.

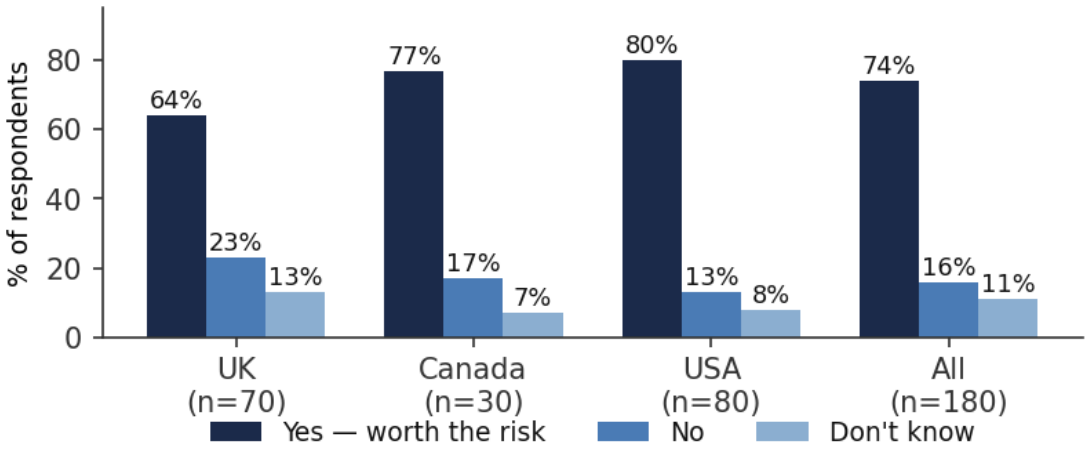
Regulatory uncertainty is an important factor shaping UK firms' approach to scaling AI. It is cited as the primary barrier to scaling AI by 41% of UK firms, compared with 37% in North America. The UK's principles-based approach provides a strong foundation, and regulating AI through existing financial services frameworks strikes the right balance between safety and innovation. However, there is still some uncertainty in how existing regulatory requirements apply to AI in practice.

The UK's established regulatory framework places significant emphasis around governance, auditability and oversight of AI decisions, creating additional governance considerations for firms deploying AI at scale. For example, while SM&CR framework does not impose AI-specific requirements, it reinforces the importance of clear human accountability for activities involving AI. This has no direct equivalent in the US regulatory environment, although U.S. firms increasingly face AI governance obligations through a combination of federal supervisory expectations and state-level requirements.

The data quality picture is also inverted between the two markets. UK firms rank integration ease as their primary data quality concern, suggesting cleaner underlying data but fragmented legacy systems. North American firms rank data trust as their primary concern, suggesting more connected data pipelines but greater anxiety about provenance and reliability. These represent genuinely different problems, requiring different vendor responses and different investment priorities.

One area where the UK leads North America is Trade Execution impact scores. UK firms rate it at 3.97 (out of 5) for revenue impact, compared with 3.76 for US and Canada. This reflects the concentration of global execution infrastructure in London, and the tangible profit and loss gains that AI optimisation is already delivering in that context.

Figure 8: *First-mover advantage — worth the risk? By geography*



Canada — A Market Worth Separate Attention

Canada (n=30) is a sufficiently large sub-sample to identify distinct regional patterns. The findings reveal a market that sits between the UK and USA on most metrics but shows some characteristics that are distinctly its own.

Canadian respondents back first-mover advantage at a rate of 77%, higher than the UK rate (64%) and close to the USA (80%). Canadian institutions demonstrate North American risk appetite, consistent with a broader pattern of lower regulatory anxiety seen across US and Canadian respondents compared with the UK.

On agentic AI, 60% of Canadian respondents are Scaling or Fully Operational, the highest of the three markets surveyed. Canada appears ahead of the UK on agent adoption maturity and broadly matches USA pace, which is notable given the smaller scale of Canada's domestic financial market.

Data access is a prominent challenge in Canada, cited by 20% of respondents as their primary data quality concern, matching the US sell-side rate and significantly higher than the UK figure. This likely reflects the smaller Canadian ecosystem for specialist financial data providers, creating genuine access constraints that are less acute in larger markets.

Canada is also the most commercially oriented of the three markets in its AI philosophy. Revenue/Alpha is the leading priority for 43% of Canadian respondents, with 23% leading with Data discovery: more commercially aggressive than the UK and less cost-focused than the US sell side. Compliance is the lead priority for only 10%, the lowest compliance-first proportion of any market in the dataset.

Section 10:
Size Matters — How Adoption Shifts by Firm Scale

Buy Side by AUM

The relationship between AUM size and AI adoption on the buy side is not linear. The most active tier is not the largest, and the most cautious tier is not the smallest.

The under-$10 billion tier (n=25) punches significantly above its weight, with 52% at Scaling or Fully Operational, higher than any other buy-side AUM tier. Firms with minimal legacy estates and lean operating models are building AI-native from the ground up, using AI primarily for efficiency and data discovery. Research/Alpha scores 3.92 (out of 5) for this group. First-mover conviction is strong at 68%.

The $10 billion to $200 billion tier (n=39) is the most commercially active, with 51% at Scaling or Fully Operational. Revenue/Alpha leads as the primary AI philosophy (44%) and Research/Alpha scores 3.95 (out of 5), the highest score for any buy-side AUM tier. These mid-size firms have enough scale to invest meaningfully in AI and enough agility to deploy it without the governance overhead of the largest institutions.

The $200 billion to $1 trillion tier (n=17) is the most cautious in the dataset, with only 29% at Scaling or Fully Operational, the lowest adoption rate across all AUM tiers. Revenue impact scores are the lowest across all six areas. Skills shortage and regulatory uncertainty are co-equal primary barriers, each cited by 47% of respondents. This tier has scale large enough to carry significant systemic exposure, creating governance requirements that materially slow AI adoption, without the resource advantage of the very largest institutions.

AUM tiers on the buy side broadly mirror the Hedge Fund versus Asset Manager dynamic described in Section 3. The aggressive $10 billion to $200 billion tier reflects the hedge fund profile: lean, fast-moving, skills-constrained. The caution visible at $200 billion to $1 trillion is more characteristic of large discretionary asset managers, where systemic governance obligations slow adoption. AUM and fund type do not map perfectly: large multi-strategy funds and small passive trackers can occupy the same AUM tier with very different behaviours — but the directional correlation holds.

Sell Side by Headcount

The sell-side size picture is clearer and more consistent: larger institutions are further ahead on AI adoption, and the relationship between scale and adoption maturity is generally monotonic.

Firms with more than 50,000 employees (n=5) are the furthest ahead in the dataset, with 80% fully operational on agents and impact scores of 4.2 to 4.6 (out of 5) across all six business areas, the highest in the entire dataset. Data access, rather than regulation or skills, is their only remaining constraint. Scale enables investment in AI infrastructure at a level that smaller firms cannot match.

The 500 to 2,500 employee tier (n=35) is the fastest-scaling mid-tier, with 71% at Scaling or Fully Operational, the highest proportion of any sell-side tier. Revenue/Alpha leads as primary philosophy (37%) and impact scores are consistently strong: Trade Execution at 4.17, Client Services at 3.97 (out of 5). These firms are moving aggressively, with 49% targeting their AI goals within six months.

The 2,500 to 10,000 employee tier (n=23) is the most mature mid-tier, with 70% at Scaling or Fully Operational. Trade Execution scores highest for revenue impact at 3.83 (out of 5). Regulatory uncertainty is the primary barrier.

The 10,000 to 50,000 employee tier (n=8) shows a distinct shift in priorities: Risk Management becomes the top revenue impact priority at 4.25, overtaking Trade Execution at 3.14. Large institutions focus AI on systemic risk and compliance first, execution second.

Firms with fewer than 500 employees (n=24) show the widest variance in the dataset. Some are fully operational, others have not started. Cost-saving dominates as primary AI philosophy (38%), and data access is the primary barrier (46%), reflecting the external data dependency of smaller sell-side operations.

Section 11:
Key Findings and Recommendations

The following summarises the eight most significant findings from this research and the practical implications for AI decision-makers at financial markets institutions.

1. Sell side leads on agentic deployment, but convergence is coming
On agents, 65% of sell-side firms are Scaling or Fully Operational, compared with 47% of buy side. The sell side's urgency is higher across every metric. But buy-side institutions, particularly hedge funds, are closing the gap, and the data suggests convergence within the next twelve to eighteen months for the most aggressive buy-side operators.

2. Hedge funds are materially ahead of asset managers
The constraints are fundamentally different: regulation for hedge funds, skills for asset managers. Vendor conversations and product roadmaps need to reflect these different barriers rather than treating the buy side as a single market.

3. Research/Alpha defines the asset manager AI use case at 4.00
The highest single buy-side score in the dataset. Asset managers are using AI primarily to improve research quality and speed, representing a cultural shift from defensive data governance to offensive alpha-seeking.

4. Data readiness, not model capability, is the foremost concern
Across every segment and geography, practitioners are emphatic: you cannot deploy AI reliably without first making your data AI-ready. Fixing metadata, building semantic layers, aligning internal standards and establishing provenance frameworks is prerequisite work, not parallel work.

5. The 57% desktop allocation masks a strategic divide
The average firm spends 57% of its AI budget on desktop tools in both buy and sell side. But the underlying mix is different: buy-side desktop spend is concentrated on research and alpha tools, sell-side on workflow automation. The headline similarity conceals fundamentally different strategic rationales.

6. Canada is the most commercially oriented market and the most advanced of the three geographies surveyed
Canada's less complex regulatory burden relative to the UK, combined with North American commercial risk appetite, is producing faster adoption. Data access is cited as a key challenge for Canadian firms.

7. Regulatory Compliance dominates autonomous agent deployment intent, Trade Execution polarises
Firms are either all-in or actively avoiding autonomous execution. Strategy, not technology, determines which side of that divide they sit on.

8. Training the whole organisation is the most powerful adoption accelerant
Firms that democratise AI access move faster, sustain buy-in more effectively, and build cultural foundations for scaling that centralised AI teams cannot replicate. Culture, not technology, is the primary scaling differentiator.

About this Research

This white paper is based on findings from A-Team Insight's 2026 research programme: Generative and Agentic AI in Financial Markets.

Survey

Structured questionnaire covering AI philosophy, adoption stage, impact expectations, governance and data strategy. Field dates: June 2026. Total respondents: 200 (100% completion rate). 79% are the primary AI decision-maker with budget authority at their organisation. Segments: Buy Side n=85, Sell Side n=95. Geography: USA n=80, UK n=70, Canada n=30. Buy-side investment strategy breakdown: Hedge Fund (Quantitative, Macro, Multistrategy) n=55; Asset Manager (Discretionary, Passive) n=30. All figures rounded to nearest whole percentage point.

Qualitative Research

In-depth interviews with senior AI and data leaders at buy-side and sell-side institutions, conducted by A-Team Insight researchers throughout 2025 and 2026. All interview subjects are anonymised in this report and referenced by firm profile and role rather than by name or institution. Findings are also contextualised against intelligence gathered from A-Team Insight's ongoing programme of industry events, webinars and panel discussions with capital markets practitioners.

About LSEG

LSEG is one of the world's leading providers of financial markets infrastructure and delivers financial data, analytics, news and index products to more than 44,000 customers in over 170 countries.

Through the three pillars of our AI Strategy, LSEG is setting the standard for trusted AI in finance.

- Trusted data – our open, partnership approach fuels AI in financial services with world-leading, high-quality, accurate data
- Transformative products – delivering new AI products that enable our customers to work faster, smarter and with confidence
- Intelligent enterprise – using AI to drive efficiency, innovation and growth across our operations

We help organisations fund innovation, manage risk and create jobs by partnering with customers at every point in the trade lifecycle: from informing their pre-trade decisions and executing trades to raising capital, clearing and optimisation.

Backed by more than three centuries of experience, innovative technologies and a team of over 26,000 people in 65 countries, we are driving financial stability, empowering economies and enabling you to grow sustainably.

Find out more: lseg.com.

A-TEAMGROUP

For Marketing: **www.a-teamgroup.com**

For Content: **www.a-teaminsight.com**

A-Team Group is not a licensed investment advisor and therefore does not offer investment advice. The opinions and information provided in this material reflect our judgment as of the publication date and may change at any time without prior notice. This material is intended solely for informational purposes and should not be interpreted as investment advice. Although the information herein is derived from sources believed to be reliable, A-Team Group makes no warranty as to its accuracy or completeness.