



SPONSORED BY:  Arcesium®

JUNE 2026



AGE OF AI II

UNDERSTANDING HEDGE FUNDS' AI CAPABILITY VERSUS USAGE

EXECUTIVE SUMMARY

Oh good, another AI report. Before you close the tab, a promise: this one is about the part nobody put in the brochure. In our Age of AI report last year we asked the question whether hedge funds are investing in artificial intelligence. And tried to understand their rationale for the hefty sums they were putting into the technology. They did invest, and the efficiency gains were real.

This year the question is harder and could be seen as less flattering. The tools have been bought, the platforms switched on, the licences paid for. So why are so few who work at hedge funds using them?

That is what this report is really about. Two-thirds of managers tell us most or all of their staff use AI tools regularly. But is it making the running of hedge funds more efficient just yet? When someone needs a quick answer from an operational system, only one in ten reaches for an AI assistant. The technology has arrived. The habits have not. And the two are drifting further apart, because adoption has outrun the trust, the training and the measurement that are supposed to hold it together.

This is not a story of failure. It is a story of a maturing market discovering that buying capability and building usage are different problems, and that the second is mostly about people. The firms closing the gap are not the ones with the most tools. They are the ones treating adoption as something to be managed rather than something that happens on its own.

MANAS PRATAP SINGH
HEAD OF HEDGE FUND RESEARCH


CONTENTS

METHODOLOGY	3
KEY FINDINGS	4
PART I THE ADOPTION GAP	5
EXCLUSIVE INTERVIEW WITH ARCESIUM FIELD CTO	8
PART II THE USER EXPERIENCE PROBLEM	9
PART III THE NEW INTERFACE LAYER	12
PART IV TRUST, MEASUREMENT AND THE ROAD AHEAD	14
CONCLUSION	16

METHODOLOGY

This report draws on Hedgeweek's Q2 2026 Hedge Fund Manager Survey, completed by hedge fund firms across the major global domiciles, AUM bands and flagship strategies, alongside interviews with named and unnamed industry sources conducted in Q2 2026. AUM cross-tabulations compare firms managing \$100m to \$1bn against those above \$1bn. Some questions allowed multiple responses, so figures do not always sum to 100%.

Chart 1.1 Manager Survey Demographics



KEY FINDINGS

1

Broad but shallow

Two-thirds of managers (67%) say most or all staff use AI tools regularly, but only 10% reach for an AI assistant when they actually need a quick operational answer. The tools are deployed; the habits are not. Adoption has spread across firms without yet embedding in daily workflows, leaving a gap between reported usage and real reliance.

2

Smaller firms are going deeper

Depth of adoption seems to fall as firms get bigger. Among managers running \$100m to \$1bn, 39% say all staff use AI regularly. Above \$1bn, that drops to 13%. Larger firms have more tools but more committees, more data and more friction, so usage thins as it scales. Nimbler firms continue to lead, echoing last year's finding.

3

Trust, not usability, is the brake

Asked the single biggest barrier to wider adoption, just 1% blame complex or unintuitive tools. The real obstacle is trust: 44% cite concerns about data accuracy and hallucination, more than three times any other factor and consistent across firm sizes. The friction is in the output, not the interface.

4

The strategy vacuum

38% of firms have no formal strategy for driving AI adoption, the most common answer, and 28% have no formal training programme. Among firms above \$1bn the gap widens sharply: 47% have no strategy and 53% have no training. Tools have been bought faster than the management scaffolding to support them.

5

Measuring nothing

More than half of managers (52%) don't measure whether their AI investments deliver value at all, rising to 80% at firms above \$1bn. Trust and measurement reinforce each other: it is hard to build confidence in AI without evidence of what it delivers, and hard to justify measuring it without the trust to act on the results

PART I: THE ADOPTION GAP

From investment to impact





Chart 1.1 Depth of regular AI use across all staff, by firm size
size Depth of regular AI use across all staff, by firm size



A year ago, the hedge fund industry was in the acquisition phase of its relationship with AI. The reasons were practical and well rehearsed: scale operations without adding headcount, reduce errors, do more with the team you already have. Those drivers have not gone away. What has changed is that the spending has now happened, and the bill has come due in a currency money cannot settle, which is behaviour change.

On paper, adoption looks healthy. Some 67% of managers say most or all of their staff use AI tools regularly, with a further 22% reporting that usage sits with a small group of power users. Read quickly, that is a remarkable penetration rate for a technology still measured in single-digit years of serious enterprise deployment. Read slowly, it describes something shallower. Regular use and embedded use are not the same thing, and the survey data pulls them apart the moment you look at firm size.

Among managers running between \$100m and \$1bn, 39% say all staff use AI tools regularly. Among those above \$1bn, that figure collapses to 13% (Chart 1.1). The larger firms are far more likely to describe their usage as

occasional, the kind of thing people reach for when they remember it exists rather than the default way work gets done. This is counter-intuitive only if you assume bigger budgets buy deeper adoption. They do not. They buy more tools, more approvers and more places for momentum to dissipate.

The pattern echoes one of last year's more surprising findings, that smaller and nimbler funds led on comprehensive implementation. Matthew Katz, Senior Vice President and Field CTO at Arcesium, framed the mechanism then in a way that holds up now. Large firms, he noted, are organisationally more complex, with more committees and more approvers, but the deeper constraint is data: the more of it you have, the harder clean governance becomes, and the harder it is to get information into a state an AI model can actually use. Smaller teams can establish that discipline incrementally, experiment, learn and pivot. Bigger ones have to coordinate it across the enterprise first. A year on, the same dynamic now expresses itself not in whether firms have deployed AI but in whether their people use what has been deployed.

“Over the past twelve months, the conversation has shifted from aspiration to application, and that shift has been genuinely useful,” founder, CIO and CEO of Demeter Tactical Investments, Jeffrey Sexton told *Hedgeweek*®.

What has changed, he explained, is the quality of the tooling available to systematic managers for the work that sits around the investment process. This includes, “compliance monitoring, data processing, operational workflow, and the speed at which large volumes of market information can be structured and made usable. Those improvements are real and we have incorporated them where they add genuine value”.

WHERE USAGE CONCENTRATES

The concentration of usage in a handful of power users is roughly constant across the size bands, at around a quarter of firms regardless of AUM (Chart 1.2). This is worth dwelling on, because it is the quiet structural feature of hedge fund AI adoption. Most firms are not waiting for a tool. They have one. What they have not done is move it from the enthusiasts to everyone else, and that distance, from the early adopter's desk to the median employee's

daily workflow, is where the gap lives. “One in four hedge funds keep AI use concentrated in power users regardless of AUM”

The good news buried in the data is that outright resistance is rare. Only a small fraction of firms have tools sitting idle or have not deployed at all. The problem is not refusal. It is the far more ordinary friction of people who are busy, slightly sceptical and perfectly able to do their jobs the way they always have. Closing that gap is a management problem dressed up as a technology problem, and the firms that recognise the difference are the ones making progress.

But that barrier is slowly being overcome. According to Sexton, the investment management community is becoming more precise about what AI is genuinely suited to do within a systematic process. That precision, he explained, benefits managers who have built their process on research rather than on tooling, and it benefits allocators who are developing more sophisticated frameworks for evaluating how systematic strategies like theirs actually work.




Chart 1.2 Concentration of usage in power users



One in four hedge funds keep AI use concentrated
in power users regardless of AUM



MATTHEW KATZ

Matthew Katz, Senior Vice President, Field CTO, Arcesium



When it comes to AI and hedge funds, what has changed most in the last 12 months?

Two big themes. The first is ROI. People have started the journey, the spend is now serious, and leadership is asking what they are getting for it. A year ago, the conversation was about getting people involved and building usage. Now it is about discernment: not all usage is equal, and not all of it produces great results. Firms have seen the effects of both good and bad usage, and they are starting to work out where the value actually sits.

The second is sophistication. Everyone began with the simplest uses of AI. Now people are leveraging up into genuinely agentic workflows, coordinated groups of actions that follow data through a process. The real design question becomes: how do I build something that involves my judgment where it matters, and leaves me out where it does not?

Our data shows two-thirds of firms say staff use AI regularly, but only one in ten reaches for an AI assistant for a quick operational answer. Is genuine adoption missing?

I would want to see that split by seniority, because I suspect a lot of it is exactly that. Operations people, especially senior ones, have been under pressure for decades to move fast with what they have. A senior person knows instantly where the answer is and can find it in seconds, so a new tool is not necessarily faster than the muscle memory they built over twenty years. The tool lifts people with less context highest, because it spares them the pain and the tears the rest of us went through to learn. For the most senior staff, AI is less a daily assistant and more something you reach for when you need to do deep thinking or build a big plan. If it is a repetitive task, that is where you will see them want to automate.

Hallucination is once again cited as the single biggest barrier to wider adoption by hedge funds. How is Arcesium supporting its clients with that issue?

Our philosophy is that the numbers need to be the numbers. If you are asking a tool that can give you a different answer each time, that is not how you want to deal with numbers. So wherever possible, you use generative tools to create deterministic ones. AI is excellent at

turning unstructured data into structured data, but once you have something structured, you want it to run as deterministic code, because it is faster, cheaper and it does not hallucinate. With **Arcesium Intelligence***, our agent studios, embed our context and our learnings, and we only do the financial industry, so an agent does not need to be a free-roaming model that invents new things. An agent is something that accomplishes your goal. When we produce investor fact sheets from client data, we are not making them up on the fly. We use generative tools to get over the hump of the work, then hit the mark every time through a deterministic pathway.

What does trustworthy AI actually look like to you?

We should not anthropomorphise these systems. They are statistical models with code wrapped around them. But the ways you trust an AI are the ways you trust a person. Is it open about its sources? Does it cite where the information came from? Has it been wrong before? You build trust incrementally, the way you would with a new hire: small tasks first, then more important ones as it earns them, and guardrails written in when it

makes a mistake. If you insist it always shows its sources, you get a chance to look at something and think, that feels off, and go check, rather than a binary choice to trust it or not.

Training and onboarding is the second most common barrier to widespread adoption by hedge funds. How should firms approach it?

The same way you onboard people. Write things down. The more you document how the work should be done, the more both humans and AI can ingest it. When something goes wrong, do not just say “you did badly”; go back and correct the documentation, because then it works for the next person and the next agent too. You have to be more explicit with AI than with people, because it has no common sense and no memory, it learns everything anew each time. But if you put in that work, co-working and mentoring rather than one-shotting whole systems, you get back great results. I am actually writing about what business users can learn from how developers have worked with these tools.

* Launched in May 2026, Arcesium Intelligence is an AI platform embedded in Arcesium's Opterra® and Aquata® systems that lets institutional investment firms deploy agentic AI across operations and data workflows. It launches with 100+ pre-configured agents plus a self-service studio for building custom ones, all governed by a domain-aware orchestration layer providing institutional-grade compliance, auditability, and governance.

PART II: THE USER EXPERIENCE PROBLEM

Why capable tools go unused

If adoption is broad but shallow, the obvious next question is what is stopping it from deepening. The answers managers give are revealing, mostly for what they do not say.

Almost nobody blames the tools. Asked to name the single biggest barrier to wider adoption, just 1% of managers cite tools being too complex or unintuitive (Chart 2.1). This is the great misdirection of the user experience debate. Vendors and buyers alike tend to assume the friction is in the interface, that if the buttons were better arranged and the dashboards more elegant, usage would follow. The data says otherwise. The friction is real, but it hides inside other answers.

The single largest barrier, named by 44% of managers, is concern about data accuracy and hallucination, more than three times any other factor (Chart 2.1). This is a trust problem, not a usability one, and it is strikingly consistent across firm sizes, cited by 36% of mid-market firms and 33% of institutional ones. People are not avoiding AI tools because they are hard to operate. They are avoiding them because they are not yet sure they can rely on the output, and in an industry where being precisely wrong is a fireable offence, that caution is rational.

The second barrier divides more sharply by size. Training and onboarding is cited by 17% of firms overall, but rises to 40% among those above \$1bn (Chart 2.2). The institutional end has a structural gap here that the mid-market does not: 53% of larger firms have no formal AI training programme, against 18% at the smaller end (Chart 2.3). The bigger the firm, the more likely it is to have handed staff powerful tools with no systematic way of teaching them to use them well, then wondered why uptake stalled.

Chart 2.1 Single biggest barrier to wider AI adoption




Chart 2.2 Training & onboarding cited as a barrier, by firm size




Chart 2.3 Firms with NO formal AI training programme, by firm size




Chart 2.4 Approach to driving AI adoption





Chart 2.5 Firms with NO formal adoption strategy, by firm size



STRATEGY AS PRIORITY OR AFTERTHOUGHT

Beneath both barriers sits a more fundamental absence. Some 38% of firms have no formal strategy for driving AI adoption at all, the single most common answer. A further 27% take a decentralised approach in which each team owns its own adoption, and only 9% rely on a top-down mandate (Chart 2.4). Among firms above \$1bn, the strategy vacuum widens to 47% (Chart 2.5). Tools have been procured faster than the organisational scaffolding to support them, and the result is the gap this report keeps returning to.

This is where the experience of those who have gone furthest becomes instructive. Thomas Rice of Minotaur Capital, who built his firm's proprietary system Taurient two days after founding the firm, is blunt that capability alone does not produce trust. The bar for relying on AI-generated investment research, he says, is extraordinarily high, as it should be, and clearing it takes an extensive test and evaluation framework where methodologies are validated against the outputs you actually want. His firm's discipline, he argues, is what has let it avoid the common trap of over-relying on flashy but flawed implementation. Notably, Minotaur did not reduce headcount as AI supercharged its

research. It used the technology to move faster, citing its ability to react to a Munich Security Conference speech ahead of peers, while keeping human judgment over everything the system produced.

Vincent Berard of THEAM Quant locates the same problem one layer down, in explainability. Machine learning has been part of his firm's products for years, he notes, but the question of how to use it in actual investment decisions remains genuinely unsettled, because clients want to know why a model is long the S&P and short gold today, and being told it is the output of a two-layer neural network does not

satisfy anyone, least of all during a difficult month. Robustness and explainability, in his telling, are not compliance niceties. They are the precondition for anyone trusting the tool enough to use it. His firm wraps its strategies in strict governance, with independent robustness testing and no model changed overnight without committee approval, the institutional answer to the trust problem that the survey shows holding adoption back.



PART III: THE NEW INTERFACE LAYER

The login is still winning

This is the new interface layer, the one built on enterprise assistants, copilots and the connective tissue increasingly described as MCP, meeting the old reality of tabs and logins, and mostly losing. The technology to ask a system a question in plain language and get an answer exists. The default behaviour, even at firms that have deployed such tools, is still to do it the old way. Among firms above \$1bn in our sample, not one reached for an AI assistant first, defaulting instead to colleagues or manual reports.

That last finding deserves a caveat, given a smaller institutional sample, but the direction is unmistakable and it matters. The promise of the new interface layer is that it collapses the distance between a question and an answer, removing the navigation, the export, the second login. The data shows that promise is real in capability and almost entirely unrealised in practice. The firms that build the layer are not yet the firms whose people reach for it by reflex.

THE VANGUARD AND THE HOLDOUTS

Enterprise AI assistants remain early in their deployment curve. Only 10% of firms have fully deployed one connected to operational systems. A further 22% are piloting and 27% are evaluating, while the largest single group, 42%, is not yet considering one at all (Chart 3.2).

The institutional picture is the more interesting one, because it is bifurcated. Firms above \$1bn are simultaneously more likely than mid-market peers to be fully deployed, at 13% against 7%, and more likely to be sitting the whole thing out, at 47% against 36% (Chart 3.3). A small

institutional vanguard is well ahead. A larger institutional cohort has not begun. There is very little in the middle. The shape of that future interface is already visible at the most advanced firms, and it is not a chat box. Marcus Storr of Feri describes seeing in-house tools whose agents write their own agents, where a fund arrives in the morning to find that overnight the system has read every piece of investment bank research received in the previous twelve hours, summarised it, cross-referenced it against the fund's long and short positions and cut the overnight news flow into digestible pieces in a consistent format. The interface, in other words, is increasingly the absence of one: the work arrives done. Storr is careful to note that this remains a support tool rather than a decision-maker, AI that filters, analyses, summarises and formats so that humans can decide, though he suspects that boundary will move, with portfolio managers soon asking the machine for its opinion as a kind of crash-test dummy for their own.

Chawkat Nammour of Bainbridge offers the ground-level version of the same shift. He has automated a quarterly tone tracker that once meant manually copying management commentary from earnings transcripts into a spreadsheet, using an AI assistant to categorise commentary by product line and surface sentiment trends over time. His web scraper, built by an analyst in three or four days, produces a daily social media feed. None of this is a futuristic agent swarm. It is one analyst rebuilding his own workflow around tools that now sit close enough to hand to be worth reaching for, which is precisely what the median firm has not yet achieved.

Chart 3.1 How staff get a quick operational answer




Chart 3.2 Enterprise AI assistant deployment status





Chart 3.3 Assistant deployment: institutional bifurcation





PART IV: TRUST, MEASUREMENT AND THE ROAD AHEAD

The measurement vacuum

The trust problem has a mirror image, and it is measurement. More than half of managers, 52%, do not currently measure whether their AI investments are delivering value at all. Among firms above \$1bn, that figure rises to a striking 80% (Chart 4.1). Where firms do measure, the most common yardstick is efficiency gains in specific workflows, cited by 41% (Chart 4.2), the same instinct that drove last year's investment in the first place.

The trust gap and the measurement gap are not two problems. They are one problem feeding itself. It is hard to build confidence in AI output when you have no internal evidence of what that output is worth, and hard to justify investing in measurement when you do not yet trust the tool enough to act on what the numbers say. The firms stuck in this loop have the tools and the usage statistics but no way of telling whether any of it is working, which is its own kind of answer.

Chart 4.2 How firms measure AI value (where they measure)

41% of hedge funds see efficiency gains in specific workflows when measured

Breaking the loop is, again, a matter of foundations rather than features. The thread running through every advanced practitioner in this report, and through last year's, is that the work which makes AI trustworthy is

mostly unglamorous data work done before the model ever runs. As Sexton puts it, the measurement vacuum we discuss here "likely reflects a challenge of definition as much as one of measurement. The question of what an 'AI investment' is supposed to deliver needs to be answered before ROI can be assessed meaningfully."

WHAT CLOSING THE GAP LOOKS LIKE

A pattern emerges across the data and the interviews, and it is consistent enough to state plainly. The firms closing the gap between AI capability and daily usage are not distinguished by the sophistication of their tools. They are distinguished by three things the survey shows most firms still lack: a deliberate adoption strategy rather than organic drift, a formal way of teaching people to use what has been bought, and a means of measuring whether it works. None of these is a technology. All of them are management.

The platform sprawl that frames the whole problem makes the stakes concrete. Some 61% of firms have operational staff working across three to five distinct systems every day, and a further 13% navigate six or more (Chart 4.3). This is the friction the new interface layer exists to dissolve, the daily tax of moving between logins and screens to assemble an answer that a connected assistant could return in a sentence. The opportunity is sized and obvious. The take-up is not yet there. The gap between those two facts is the entire subject of this report, and on current evidence it is a gap that will be closed by the firms that treat adoption as a discipline, not a download.

Chart 4.1 Firms NOT measuring whether AI delivers value, by firm size





Chart 4.3 Number of daily operational systems per firm



CONTRIBUTORS:

Manas Pratap Singh

Head of Hedge Fund Research

manas.singh@globalfundmedia.com

FOR SPONSORSHIP & COMMERCIAL ENQUIRIES:

Please contact

sales@globalfundmedia.com

Published by: Global Fund Media Ltd, 1 Fetter Lane, London, EC4A 1BR, United Kingdom

© Copyright 2026 Global Fund Media Ltd. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher.

Investment Warning: The information provided in this publication should not form the sole basis of any investment decision. No investment decision should be made in relation to any of the information provided other than on the advice of a professional financial advisor. Past performance is no guarantee of future results. The value and income derived from investments can go down as well as up.