



The Role & Future of AI on the Trading Floor

Spectre or Saviour ?

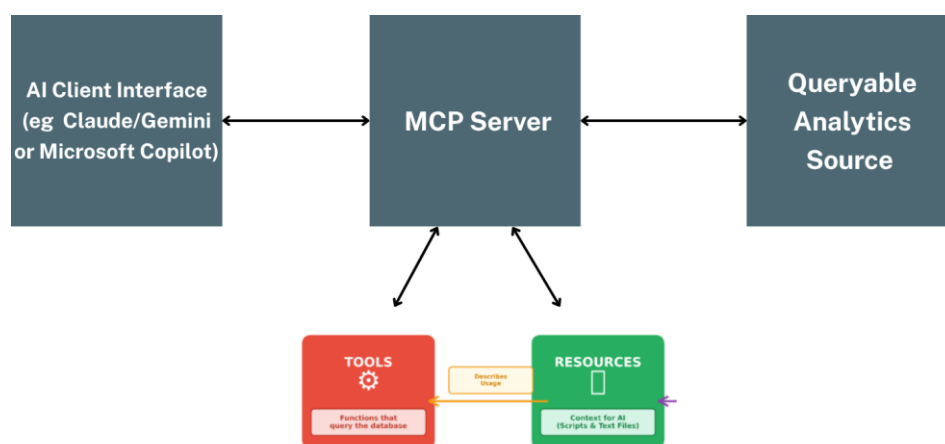
There has been no bigger topic than AI at the forefront of capital markets minds in the recent past. Geopolitical turmoil aside, AI has consumed a huge chunk of intellectual bandwidth across markets participants and regulators, with no market left out of the global conversation.

The European Union, the Bank of International Settlements, IOSCO and most developed market regulators have all formally consulted on AI in the past three years alone, but there is as yet no real consensus on how AI might sensibly be addressed to protect the integrity of our interconnected and interdependent markets.

Agents of Change

Nevertheless, AI developments like large language models (LLMs) and natural language programming (NLP) do offer real promise. Far from replacing traders – an outcome that is very much in doubt today – judicious use of AI can bridge the gaps and bottlenecks that still plague the front office in 2026.

Perhaps of most promise is Model Context Protocol (MCP). Introduced by Anthropic in late 2024, MCP was quickly adopted by all the major AI companies as the value of a standard integration became clear. For financial institutions, MCP has been one of the key AI developments, offering a real opportunity to unlock the true value of data for the trading floor.



An open-source protocol for agentic AI – that is, an AI framework that allows AI agents to take semi-autonomous action, make decisions and execute workflows – MCP connects AI agents to external tools, systems and data sources. It aims to replace ad-hoc, fragmented integrations with a single protocol. For financial firms, this offers the potential to build an AI architecture across multiple market and internal data sources, providing an integrated, top-to-bottom view into the full data universe.

Democratising Data Access

The trading function is a delicately balanced triangle with traders at one corner, quants at another and IT at the third. They have a common language – data – but there are very few firms who would confidently say they are using all the data they have to its full capacity.

The reason for this is both simple and complex. In most firms, there are only a handful of people – typically quants – who know the schemas, APIs and dataset quirks well enough to transform data into insights and answers. This however creates a natural bottleneck. Quant teams are under pressure to meet a constant stream of ad-hoc requests, which in turn takes their focus away from the more valuable work they could be doing. Meanwhile, trading teams can be left waiting for answers when the quant team is busy, which brings opportunity cost as the chances to optimise intraday trading slip away.

This is where AI agents can step in.

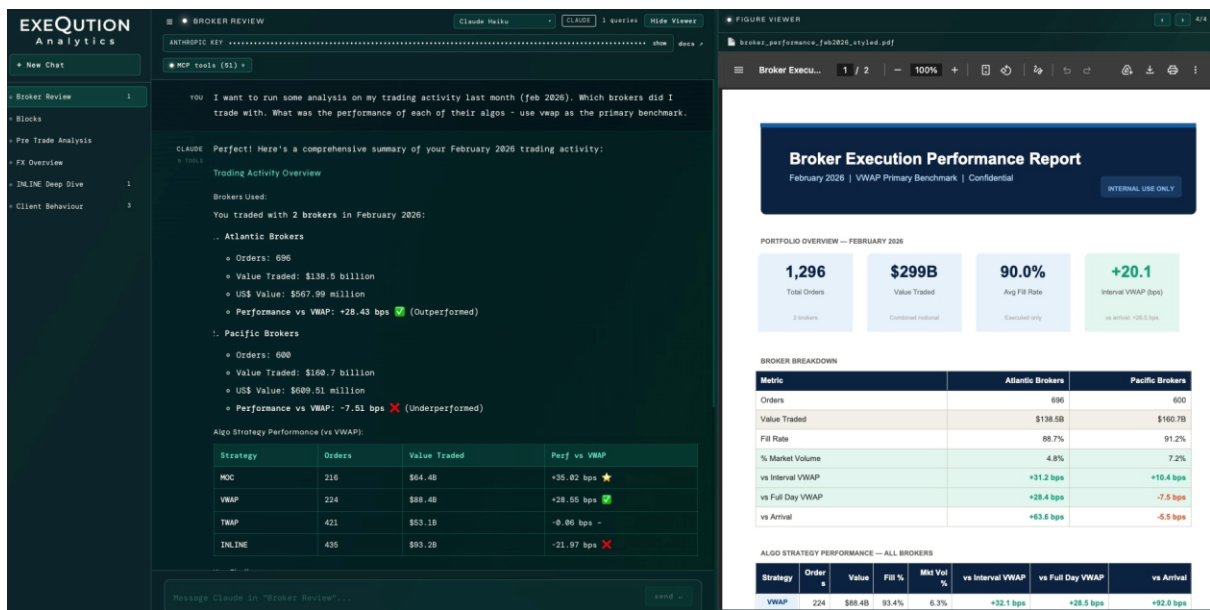
Traditionally, data and analytics dashboards have been delivered to the trading desk with pre-defined metrics – which are slow and time-consuming to develop as markets evolve and requirements change. What agents deliver is the capability for various roles across the trading lifecycle to access intelligence and insights dynamically.

Agentic Analgesic

Agents can alleviate pain in several areas of the trading floor. For Heads of Trading, a good agent can give both the helicopter view necessary to manage the floor, and the ability to dive down and rapidly diagnose and fix problems as they arise, without the need to call IT or quant teams to pull or analyse data.

For both Heads of Trading and Sales Traders, this real-time ability to interrogate data provides an enormous advantage when dealing with clients. Have a client meeting? Have the system pull the report on their trading results year-on-year, including discussion points and recommendations. This ability for anyone on the trading desk to interrogate data saves enormous amounts of time - not only can reports be prepared in minutes rather than hours or days, but they can capture the nuances of current market conditions and the particular areas of interest of any given client, even as they change over time.

For traders on the desk, agents enable dynamic execution quality monitoring in real time. The buy-side execution trader is under constant pressure to demonstrate best execution across multiple asset classes while managing real-time routing decisions, but their data is locked behind systems that require tech or quant resource to access. An agent of this kind delivers that intelligence directly onto their screens.



On the sell-side, the value of the trading desk is increasingly measured by their ability to provide smart, data-driven analysis of execution performance and market behaviour - but accessing that data at the speed the role demands is a persistent challenge. An agent closes this gap, delivering realtime views of trading performance at any time of day.

The benefits flow outside the trading function too. For compliance and risk teams, the ability to query trading data in natural language provides the context to investigate potential issues, rather than waiting for specialised resources, accelerating investigation timelines and allowing the people closest to regulatory requirements direct access to the data that proves adherence or surfaces concerns.

All this delivers extra bandwidth to both quant and IT resources. Relieved from the burden of urgent ad-hoc data and analysis requests, these valuable teams can focus on higher value tasks, while also benefitting from faster data analysis to support their work.

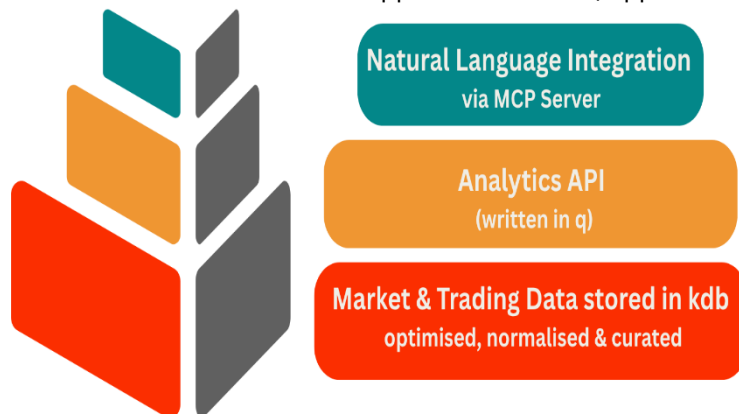
But what about hallucin(AI)tions?

One significant issue remains – one that has yet to be definitively addressed by Silicon Valley’s AI evangelists – and that is the propensity of AI to make things up. Hallucinations are a very serious concern for the front office – one fictitious result, confidently presented, could unleash a cascade of trading decisions that could theoretically spell, if not outright disaster, then at least a very bad day on the desk. This is probably why a 2025 IOSCO report observed that only about 40% of firms were using AI to support trading decisions, and 60% as part of algorithmic trading. Considering machine learning – captured under the definition of AI in this case – has been in use in these contexts for two decades or more, it is likely that the newer forms of AI, so loudly proclaimed as the future of almost everything, are being treated with extreme caution – and rightly so.

So then what is to stop an AI agent – particularly one accessed by NLP, and with full access to precious and proprietary data sets – from simply making up the answers?

The answer lies in the API. By converting natural language queries into API calls, the agent can be constrained, limited to producing responses that are directly drawn from the data. In this context the agent is a translator, rather than performing computations itself, allowing humans to ‘speak’ directly to data – and for the data to speak back.

The API layer puts strong guardrails around the agent. It enables a controlled contract - approved functions, approved data, and auditable execution paths - ensuring results are trustworthy. This type of architecture ensures that every answer is grounded in actual data returned from known, reliable sources, not the agent's probabilistic generation. Users get to enjoy the ease of asking questions in natural language, while the firm can have complete confidence that the agent's responses are factual, traceable, and consistent with the same data sources that power critical trading decisions.



Within this context, a potential hallucination looks like the agent inventing a parameter that doesn't exist, or passing something to a parameter that it doesn't understand. This can be solved for by ensuring that the parameter will respond with an error. An 'intelligent error' process can then be applied to help the agent autocorrect in most circumstances. This concept of a 'constrained agent' delivers the convenience and speed of AI analysis, drawn from the central source of truth, without the risk of hallucinations.

Outperforming on the edge

The enormous benefit of having a trustworthy agent accessing the full data universe is the opportunities it offers for speed and creativity. Markets favour survival of the quickest – the ability to shorten the path from ideation to iteration allows firms to evolve to meet market conditions much faster. Relieving quant teams of the drudge work in making this happen adds more speed again, as they instead focus on the value-adding end of the process.

In today's world data is no longer the constraint – there is an almost boundless universe of data available. Instead, it is the ability to access and interpret the data on offer that constrains performance. By eliminating the technical barrier between questions and answers, AI can transform latent organisational or market knowledge into active decision-making. It can turn curiosity into insight, insight into action, and action into competitive advantage. In markets where understanding arrives before opportunity disappears, that transformation isn't incremental – it's existential.

This new paradigm is not about finding an edge over the competition, but the next evolution of this – finding brand new boundaries to performance, those edges that others cannot find. This is a very new and emerging area, and there are many use cases here yet to be found. And all this unrealised value is hiding in the data that already exists – it just needs to be unearthed.

In the quest to find the edge of performance, the largest data set and the most powerful analytics infrastructure is of limited use if only a handful of people can access it. Integrating AI into the trading function means everyone, regardless of technical skill, can interrogate data, validate hypotheses, and act on market signals in real time. The result is a firm where innovation can emerge from every desk, and the speed of insight can match the speed of markets for the very first time.

ExeQution Analytics' agentic AI tool for trading, Eolas, was launched in March, 2026. For more information visit www.exequtionanalytics.com