TabbFORUM Report: The State of AI in the Capital Markets

Sean Creamer

Financial markets have long embraced automation and machine learning, with high-frequency trading algorithms transforming the speed and efficiency of trading in the 2010s. In this report on the State of AI in the Capital Markets, financial reporter Sean Creamer explores how advancements in AI, including large language models and natural language processing, are revolutionizing every aspect of capital markets, from trading strategies to compliance and portfolio management.

Financial markets practitioners are no strangers to automation and machine learning. In the 2010s, high-frequency trading algorithms changed how the buy-side and sell-side approach the capital markets by reducing the latency between trades, increasing alpha for early adopters.

Machine learning has only grown stronger over the last decade. With the introduction of large language models (LLMs), natural language processors (NLPs), and generative pre-training transformers (GPTs), the new age of artificial intelligence (AI) stands to reshape the capital markets once again.

Across the front, middle, and back offices, AI's ability to process vast amounts of data in real-time is revolutionizing multi-asset trading strategies, automating compliance tasks, and helping to balance portfolios. From optimizing meetings with clients to speeding up the trade settlement process, the growing influence of AI across fixed income, equities, and commodities presents unique opportunities and challenges for firms in the capital markets.

Rich Allen, Managing Director of Business Development at Numer.ai, notes, "Machine learning and data science have grown exponentially in the past ten years, not just in terms of technological capabilities and frameworks but also in terms of accessibility. Ten years ago, far fewer people were familiar with machine learning or AI and could apply it to a tough data science problem like predicting stock market returns."

Streamlining Data in Fixed Income Markets

These tools improve operational efficiency and provide deeper context in asset classes like the municipal bond space.

> Traditionally, if you wanted to gather historical market insights for a bond, an analyst would have to sift through vast amounts of messages from their trading counterparties, often searching through different identifiers and naming conventions, according to **Eugene Grinberg**, the CEO and co-founder of SOLVE Fixed Income. It's a time-consuming and fragmented process, but with SOLVE's AI-driven quotes aggregation, you can input a CUSIP and instantly retrieve all available market color related to a bond.

Their SOLVE Quotes™ platform streamlines the data collection process, allowing for easy comparisons by finding bonds from the same issuer or similar securities. The tool enables users to create watchlists of bonds they're interested in and cross-reference those with bonds currently on the market. Then, their SOLVE Px™ AI delivers predictive pricing for bonds based on vast amounts of market, reference, and proprietary data.

"We're currently focusing on two main areas: transforming unstructured text into data—data that can

be searched, charted, and used to train predictive models. For example, in the municipal bond space, we're pricing 900,000 bonds, while only a fraction of that, around one-sixth or one-seventh, is actively quoted," says Grinberg. "If you're trying to find a fair price on a bond that trades frequently, it's straightforward since public information is available. However, pricing becomes more complicated if a bond hasn't been traded in six months. This challenge is eliminated with AI-predictive pricing, as it processes all 900,000 bonds in real-time."

According to Grinberg, AI's ability to analyze vast amounts of unstructured data has been a game changer for bond market operators. Market participants share thousands of quotes in the fixed-income space daily in highly unstructured, conversational ways. Traders send many emails daily indicating bids and offerings across the thousands of securities they're filing.

The ability to ask natural language questions and receive fast, accurate responses helps analysts save vast amounts of time, which can be used for generating trade ideas, according to **Jim Kwiatkowski**, the CEO of LTX, a Broadridge company. He noted that via the BondGPT tool the firm released, users are garnering improved insights by synthesizing information from disparate data sources, but analysts can build new trading models that they wouldn't typically use because of the sheer possibility of achieving these insights.

> "We have incorporated issuer filing data, the ability to seek bonds rated X, priced Y, in Z sector, and then the ability to narrow this search to 'firms whose earnings have grown in each of the last eight quarters.' It would take much time to do this manually, says Kwiatkowski. "Further, traders expect to be able to speak the minimum required to be understood, so we ensured that the application understands bond market jargon common to trading desks."

Natural Language Processing in Financial Markets

The problem is that there hasn't been an effective way of aggregating quote data between traders and digitizing it to create a more visualized, searchable way to increase transparency in the market. And so much of the fixed income market is rarely quoted and traded. That's why Grinberg and his team launched their SOLVE Px on top of the aggregated quotes data, to create transparency across both liquid and illiquid segments of the market, including credit, structured products, and municipal bonds.

"We're using natural language processing for the data capture piece. We're using machine learning for error detection and cleansing of incorrect or mislabeled data points. Often, quotes are mislabeled by traders," says Grinberg. "Somebody may pull something at a price when it's really at yield, or someone may have fat-fingered a bad price. The model sees all that and may eliminate the outliers or reclassify incorrectly labeled data points. And most importantly, the NLP does a phenomenal job of dealing with highly conversational and unstructured communications."

Leaning on the syndicated loan space, Grinberg gave an example of how two traders might refer to the same loan.

"Traders will use nicknames in the syndicated loan space when speaking via email. One trader will call it 'McDonald's Terminal B,' and the other trader will call it 'McJunkin TLB.' The AI needs to understand that they're referencing the same loan," says Grinberg. "In addition to that, the NLP needs to be able to kick out data points that may represent the underlying characteristics of the bond while accurately capturing anything related to pricing."

AI-Driven Multi-Asset Strategies

Regarding equities, AI has helped scale both market and unstructured data. **Sandeep Varma**, the CEO of Equity Data Science, noted that with the introduction of large language models, fundamental investors can apply the same style to quantitative and qualitative sources combined. Traders, analysts and portfolio managers can have AI tools scan ticker tags and then scan through reams of unstructured notes to identify any mentions of the original ticker or relevant ones from categories like automobiles, semiconductors, or artificial intelligence.

> "When you try to ask questions [of LLMs] in a more natural language way, you get answers that synthesize both the quantitative and qualitative data. You're not just bombarded with a bunch of data on your screen," says Equity Data Science's Varma. "But along with the answer, it will also give you the reference data. Fundamental investors need transparency to understand where the answer is being derived. That is why we think combining AI and human intelligence will be the norm going forward as we merge a mosaic of many different types of content and data."

So long as an AI tool is optimized to the suitable dataset and parameters around a trading vehicle, these tools can operate agnostic of an asset class. Even exchange-traded funds benefit from the computing power of AI, according to **Joshua Pantony**, the CEO & co-founder of Boosted.ai.

"We helped Evolve ETFs use generative AI to launch an AI-powered fund, or "AI squared" as they call it, where it's looking at stocks set to do well with continued AI adoption and built with AI across a huge variety of metrics," says Pantony. "We helped an asset manager's research team cut about 95% of the time they spent poring over earnings season documents by automating their workflow of reading and synthesizing data into digestible pieces for their teams."

The ability to scale unstructured data is one of the biggest boons of AI in the capital markets, according to **Ashok Reddy**, the CEO of KX Systems. Traders can use NLPs and LLMs to parse through massive amounts of data to find opportunities in multi-asset strategies.

"If I'm going to introduce $100 million for market making with fixed assets, like three-year or five-year bonds, typically you would rely on historical data. But now, by dynamically incorporating the latest order book data and changes in real-time, along with unstructured data, you can make more informed decisions."

Reddy explains that different asset classes require different approaches, whether you're involved in market making, buy-side, or sell-side. Increasingly, hedge funds are creating derivatives and employing multi-asset strategies.

> "This strategy combines fixed assets, equities, and options to offer growth and risk management," says KX Systems' Reddy. "The multi-asset strategy is gaining popularity because it allows for a balanced approach: achieving growth while covering risks. AI can be applied to each asset class individually, but it also helps unify them into a broader multi-asset strategy."

Personalizing Client Services

Front-office operations benefit from AI's capacity to generate trading insight and help personalize client services. The power here is that the AI or the LLMs can scan through that and find patterns or sentiments in data because they can simultaneously look at multiple reports. When a company does a quarterly earnings call, AI tools can review all that text, from speech to text. This can influence trading execution, algorithmic trading, or modifying strategies based on data.

Analyzing large datasets in real time allows firms to offer hyper-personalized services to clients, from tailored investment to multi-asset portfolio strategies.

For instance, AI @ Morgan Stanley Debrief can summarize meetings between wealth management clients and advisors at Morgan Stanley Wealth Management, according to **Koren Picariello**, the Head of Generative AI Strategy & Execution at the firm. With client consent, the model can analyze the transcript, identify key topics discussed and outline action items.

"For instance, if you're my financial advisor and we spend time discussing my 401k contributions, IRA contributions, and a 529 account for one of my children, the summary might say: 'Koren and Sean spent a significant portion of the meeting discussing Koren's retirement and 529 planning," says Picariello. "The next steps could include, for example, 'Koren to fund the 529 account with $15,000 this year and consider a catch-up contribution to her IRA.' They enhance the client interaction and provide clear next

steps for both the advisor and the client, ensuring accountability and progress after the meeting."

AI for Investment Research and Risk Management

Despite the advantages, issues with staffing, data privacy, and integration with legacy systems remain top-of-mind with fund managers, technology officers and information security officers. However, the benefit of early adoption means firms can leverage AI to improve operational efficiencies and manage more risk.

Middle- and back-office functions use AI tools to increase reporting and education about processes.

AI-driven tools automate these processes by helping firms stay on top of evolving regulatory requirements.

Traditionally, when new regulations are introduced, firms take a manual approach to processing the requirements. According to **Irene Galperin**, the Senior Advisor for Financial Services at InterSystems, they assess what analytical models are needed, where the necessary data resides, and how to ensure compliance. The manual process will happen a few times to establish data governance and ensure all compliance steps are in place. Once the process is figured out, firms take one of two approaches: either hire dedicated staff to handle the ongoing compliance or automate the process.

The benefit is that if you can ingest and process more data, you can account for more variables and develop more sophisticated risk models. Real-time risk management isn't necessary for every department or trading team, but Galperin noted that the more data you can feed into a model, the more sensitive and insightful it becomes.

> "We automate the data collection process, streamline workflows, and ensure compliance with governance frameworks. We also adhere to global and local regulations regarding data privacy, making sure data is anonymized before analysis," says InterSystems' Galperin. "Another critical aspect of risk management, especially for internal compliance monitoring, is the ability to perform historical analysis. Many financial services firms sit on vast amounts of data but often struggle to access or analyze historical data—whether from financial markets or their own internal data. InterSystems can assist with tracking historical data, enabling firms to perform richer analyses and maintain a historical data store."

According to reporting from Bianca Chen at Business Insider, AllianceBernstein (AB) is going all in with using AI for risk management. Headed up by **Andrew Chin**, AB's head of investment solutions and data science, the team aims to use AI to scour over 400 company reports and filings daily to discover potential risks in relevant companies. These risks could be in AB's client portfolios or an investment directly from the firm. Natural language processing underscores this whole model, as this AI model processes human language in ways that allow computers to understand and create investment signals. Those alerts can unearth subpar performance or predict positive movement.

"One signal compares regulatory filings for company strategy or management differences. The reason is that stocks underperform when there are a lot of differences, which means that the company has a lot of dramatic changes. Usually, that means they're facing challenges," Chin said in the BI report.

One aspect of melding AI into risk management is that providers work with firms to layer in their existing tech stack. For instance, Equity Data Science enables their AI to work with major model providers like Axioma, MSCI, and Wolfe Research, as well as proprietary models from their clients, according to Varma.

"In terms of performance attribution, it could be driven by market data, factors, or alpha, allowing them to see where their performance is coming from," says Varma. "For example, in risk management, a portfolio manager could use AI to ask questions like, "Identify the most relevant ideas from a risk perspective." They could also ask which names in their portfolio contribute the most to risk, or which names they should add or remove to either increase or reduce portfolio risk."

With the right data, you can train AI models to look for market-specific risks. So now, financial institutions can transition from reactive to proactive risk management, according to **Mark Katz**, the Client Strategy and

Technology Officer at Hitachi Vantara. He added that customers want to know how to use these tools to optimize baskets of securities. This comes down to scouring through the century of market data to inform results.

> "There's 100 years of market data, so there's a lot to train the models on," says Hitachi Vantara's Katz. "They understand what types of trade baskets, under what conditions, have yielded the maximum response over the last 100 years, and the specific trigger conditions that yielded the maximum return."

Companies are turning to AI to enhance operational efficiency, and there's a growing trend of low-code and no-code AI solutions to handle repetitive tasks. This leaves more complex business application designs to human developers.

AI benefits risk management, as it can analyze over a century of market data. Firms can use AI to gain a precise understanding of risk factors that lead to financial losses. According to Katz, the goal is to replace outdated blunt approaches to risk management, such as those widely used after the 2008 financial crisis, with refined, data-driven strategies.

Most of the excitement for AI is on the investment research side, according to **John Brennan**, the CEO and Founder of the Alternative Investment Technology Executives Club (AITEC), who noted that there's a massive potential for leveraging tools that read through the vast amounts of data and articles. Research analysts spend hours, if not days, creating summaries by reviewing research and white papers. AI tools expedite that process and create summaries. And then, they'll use that information to develop or revise a model they're working on for an investment thesis.

Balyasny Asset Management is one firm that's diving head-first into AI technology. According to a report by Business Insider reporter Bianca Chen in her story, "How $21 billion Balyasny is building an AI Equivalent of a senior analyst," the firm rolled out their BAM ChatGPT tool that can plug into every internal and third-party dataset provided by Balyasny. This includes transcripts, sales and sell-side commentaries, and broker research.

"We're very focused on proactive insights. How do we actually bring these capabilities in a way that we're actually doing the analysis, in some cases in advance of folks even understanding they should be asking the question," said **Charlie Flanagan**, head of applied AI at BAM. "We're pretty laser-focused right now to move [the AI] from junior analysts to senior analysts," Flanagan said of the AI agents' abilities.

While there's a push to bring AI tools to the level of an analyst, it is hard to make a case for fully removing a human from the portfolio management equation. According to Boosted.ai's Pantony, collaborative AI meshes AI and human insight and combines the best of both worlds: AI's ability to process enormous data sets and uncover trends and human judgment for interpreting and acting on those insights.

> "We believe portfolio management benefits from collaboration because human oversight is critical for contextual understanding and managing unforeseen events," says Boosted.ai's Pantony. "This approach enhances decision-making rather than replaces it, allowing managers to maintain control while leveraging the speed and precision of AI. In finance, human traits like judgment, experience, and market intuition are critical to success and we see AI as giving asset managers more time to focus on those nuances."

Addressing AI Hallucinations

A big issue firms worry about is hallucinations from AI. These anomalies happen when an AI provides a response that contains false or misleading information, fully believing the answer is fact.

From life sciences to financial services, end users worry about misleading or incorrect hallucinations caused by inadequate training data, faulty assumptions the model makes or simple biases in the data.

At LTX, one of the most common concerns Kwiatkowski addresses about GPT technology is

hallucinations. The worry is that GPT, by design, strives to be accommodating, which doesn't suit financial market participants who require accurate and verifiable information.

"We looked to the success of ChatGPT in terms of a simple, easy-to-use natural language interface but focused on curating the data very carefully," says Kwiatkowski. "To meet the needs of our users, we need to ensure that only the highest quality sources of data go into providing answers and that there is no creativity coming from the generative aspect of GPT and creating hallucinations."

According to AITEC's Brennan, this is the frontline of AI usage. Reducing hallucinations as much as possible is paramount for asset management firms that get cagey about sharing information with each other. Larger asset management firms are testing their AI tools internally, but smaller firms must depend on the vendors and their backtesting.

According to KX Systems' Reddy, hallucinations can occur when the data an AI model has been trained on is insufficient to make accurate predictions or generate results, especially when no data is available. For example, just because the New York Stock Exchange closes at a specific price doesn't mean the stock will open at the same price the next day. Gaps can occur, or data may be missing due to collection issues or timestamp problems. Sometimes, no data is available simply because the market is closed, such as on a holiday.

"Understanding the reasons behind missing data is critical, especially when predicting outcomes in wealth management or stock market trades. If the AI generates a prediction without sufficient data, this is a hallucination," says Reddy. "To avoid this, the data must meet certain standards. It should be relevant, reliable, representative, and adhere to privacy and regulatory requirements."

This is why most capital market customers don't use general AI in production environments like real-time trading, Reddy notes. Instead, firms use AI for less critical tasks, such as summarizing documents or helping analysts review reports, but not for decision-making in trades, as the risk of hallucinations is too high and trust in the results is lacking.

However, not all hallucinations are bad. According to Reddy, firms can generate synthetic data to detect risks, anomalies and fraud. Generative AI can backtest these scenarios by replaying numerous possibilities to train models, even if some are unrealistic.

"Instead of finding similar stocks or events, AI can simulate every alternative, which reduces the risk of false positives. AI can generate diverse scenarios that haven't been encountered before, minimizing false positives," says Reddy. "This approach is not necessarily a hallucination; it helps eliminate potential risks by imagining every possible scenario. While hallucination is negative when making decisions without a solid basis, it can also be a powerful tool to generate new data and reduce risk."

Balancing Build vs. Buy in AI Solutions

Whether you're trying to avoid hallucinations or test against them, having trusted data sources is critical as more financial services firms adopt AI technologies, according to **Steven E. Orr**, the CEO and Founder of Quasar Markets. To ensure the AI tools provided by Quasar have low hallucination rates, Orr works with clients to provide datasets that best match the firm's trading activity.

> "We thought the best way to do that was to attach it to actual databases maintained by industry bodies or government entities. So these range from FRED (Federal Reserve Economic Data) and the Securities and Exchange Commission's EDGAR (Electronic Data Gathering, Analysis, and Retrieval) databases."

The idea here is that offering access to tightly maintained datasets cuts down the likelihood of hallucinations. Thus, by choosing the right AI tool and databases to inform the tool, end users can customize and build out an AI system that fits their current tech stack and proprietary data.

An important part of utilizing AI tools in the capital markets is attribution. End users need to be able to have their offerings provide a trail of context to explain how they built an answer to a prompt. Just like how blockchain ledgers show who held a specific

cryptocurrency token, AI needs to provide attribution for any answers it provides.

However, depending on how many tasks an AI tool has to run, that can lead to latency issues, according to InterSystems' **Galerpin**. Many clients keep their data in cloud storage environments, and they want to know how they can eliminate latency while improving data quality and consistency.

"Every time you move data from source systems and your data provider's API into a warehouse and then out of the warehouse for analysis, there are quality, consistency, and latency issues," says Galperin. "Our solution is to go directly to the source system. Via our low-code applications, like InterSystems® TotalView for Asset Management™ or the InterSystems® Data Fabric Studio™, we provide a low-code, self-serve way to facilitate this data management approach."

When it comes to AI adoption, one of the issues firms on the buy-side and sell-side face is ensuring their data is AI-ready. Galperin notes that data in most financial services firms is spread across silos making it challenging to feed AI models with unified data. A smart data fabric offers an elegant solution that does not require a lot of time and resources to deploy.

"Not all firms want to dedicate developer resources to create and maintain the systems needed for a smart data fabric approach," says Galperin. "So we launched low code options for firms that want to take advantage of these new data management methodologies."

An issue facing asset management firms buying a solution is that some platforms come stuffed with features they don't need. Or data feeds that don't pertain to the trading they're engaged in. Plus, having extra models or datasets uses a lot of energy.

The Electricity 2024 report from the International Energy Agency noted that data center electricity usage will exponentially rise in the future due to AI and cryptocurrency usage. To keep costs low and latency high, tech providers like Quasar Markets are taking a simple approach to ensure clients get what tools they need.

> "We wanted to offer more flexibility so users can opt for additional features or keep it simple. We also let them customize the platform and trade wherever they want, without locking them into one ecosystem," says Orr. "TradingView has excellent charting software, so we integrated it into our platform. We thought it would be best to let them trade on Fidelity and Trading Central. We are agnostic to what platform the users want."

Asset managers are concerned about being able to connect to existing order and execution management systems. For instance, LTX trading capabilities are integrated with major order management systems (OMS) and execution management systems (EMS) like Aladdin, CRD, TradingScreen, and Flextrade, and its BondGPT tool is accessible via API integration ensure seamless integration with existing trading workflows.

"We offer the ability for institutional investors to integrate their own data in BondGPT+ – trade history on LTX and other trading platforms, recommended lists, restricted lists, and other information that they use in their decision-making in their trading workflows," says Kwiatkowski.

According to Brennan of AITEC, this buy-and-build model is very popular with asset management firms. He noted that there are firms that have built their own LLMs within their own environment, building the server farms and setting up parameters so they have full control of the data going in, ensuring there's no way data is getting exfiltrated.

While smaller firms have to buy and build, larger firms with the staffing and financial resources can hire the best-in-class developers to customize an AI tool for their specific needs.

> "Some firms are using products that are out there in the industry marketplace that offer compliance-focused environments so that you have full control over it, whether your data is cloud-based or physical," says Brennan. "Firms that don't have the resources or want to expand their options are using third-party tools, and there are several out there today

> like ELLA from Eze Castle Integration or the BlueFlame AI."

AI Interfaces

The end-user experience is another factor asset and wealth managers consider when picking tools. Many people associate user interfaces with GPT-style interfaces, where a user enters a query, and the AI generates a response. In some cases, that is what the user-facing UI ends up being.

For instance, Morgan Stanley partnered with OpenAI to leverage the base knowledge of GPT-4, which has been trained to understand human language. The tool, AI @ Morgan Stanley Assistant, allows advisors and their teams to query the tool. The model uses the vastness of human language to respond to the question, but only by sourcing answers from Morgan Stanley Wealth Management's internal content.

> "Traditional virtual assistants rely on pre-written question-answer pairs, limiting their ability to respond," says Picariello. "Our AI-enhanced assistant can answer any question by sourcing answers from internal data, allowing professionals across wealth management to quickly find and summarize information from over 150,000 internal pages, delivering relevant results in seconds."

Conversely, some firms want AI's power in the background, optimizing the systems their traders know and use each day. That's why SOLVE Fixed Income designed their SOLVE Px tool to offer multiple interfaces for users. One is an institutional app with charting and comparing tools. There's an Excel add-in interface to meet the needs of clients depending on Excel, so the software plugs directly into the tool. Lastly, there's an API plug to speed up data transmission between data sources and the Px model.

How AI Polices Itself

Whether firms are building their own AI models or are leveraging products from vendors, stress-testing AI models is an important aspect of adopting these tools. If an AI recommends a basket of investments or a trigger point for buying and selling based on faulty or hallucinated data, and the trade results in financial loss, there could be significant consequences.

According to Hitachi Vantara's Katz, an essential tool in the battle against hallucinations is retrieval-augmented generation data, which can check any answers an AI provides. RAG Data is a hybrid AI technique that combines information retrieval with text generation. When an AI model generates an output, it retrieves relevant information from a predefined, trusted dataset and incorporates it into the generation process. This method helps ensure that the generated content is based on factual, reliable data rather than relying purely on the model's internal knowledge, which could be limited or lead to hallucinations.

"RAG (Retrieval-Augmented Generation) data is a known authoritative dataset used to check AI-generated results before being shared with end users or other systems. At Hitachi, we ensure customers have robust, authenticated RAG data, which helps eliminate obvious errors," says Katz. "However, subtler inconsistencies may still occur, which we address through explainable AI using observability and telemetry. The models learn from both mistakes and successes by continuously feeding back corrections, improving accuracy over time."

However, RAG data isn't bulletproof. Moving to simpler language models from generalized large language models is a valuable risk mitigation strategy. More targeted pre-trained models are starting to appear in the common open-source libraries.

To offset this, many AI models for financial services are working towards layering in explainable AI (XAI). Most AIs operate as a black box, and while they produce results, the reasoning behind the results is unclear, Katz notes. The idea behind an XAI model is that it asks an AI about a result produced, enabling models to explain their reasoning. XAI is not fully developed yet, but there's an effort at Carnegie Mellon and other institutions to change the structure of the models to incorporate XAI into them.

> Having an AI explain its reasoning is crucial for financial practitioners. But another area that concerns asset managers is how AI models decide what data can inform an answer. Unstructured data, as the example

> above proved, can contain blindspots such as cultural or social biases. Generally, unstructured data can come from social media feeds and other human-generated content, which can add jargon and bias. The answer to this is responsible AI (RAI).

According to Microsoft, RAI has six pillars: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability.

"Responsible AI involves training on massive, unstructured data sets like social media feeds, which can introduce human biases or blind spots. The model can learn these issues, leading to biased results," says Katz. "To address this, test frameworks and QA checks are being built into pre-trained models to detect and correct such biases. Training against hallucinations is also crucial, ensuring the models are thoroughly tested and guided to produce accurate results."

Regulatory Oversight for AI

While technological advancements are happening in leaps and bounds, innovation isn't going unnoticed by regulators. The Securities and Exchange Commission introduced a new rule aimed at regulating the use of AI by investment advisers and broker-dealers. The primary concern is that these technologies might introduce risks or conflicts of interest for investors.

"The SEC has issued guidelines and wants to know if firms are using AI, which has made some firms cautious. They're also encouraging transparency with investors about AI usage," says AITEC's Brennan. "I credit the SEC for evolving—they now have experts who understand the technology and risks, and they respond quickly, especially with AI developments."

Firms using AI must demonstrate that they're using the tools responsibly and transparently. The SEC wants to understand how AI models could influence market behavior and pricing as automated, AI-powered trading begins to take hold. As such, firms must disclose their AI usage to stakeholders if they're using AI models to generate reports or make investment decisions. ▲

• • • •

***Sean Creamer** is a writer focusing on AI, financial, and marketing technology. He's worked as a journalist for HFMTechnology, Wall Street Letter, and Operational Risk and Regulation. As a marketing tech writer, he penned stories for EMARKETER. In addition to writing about real estate, Sean is an outdoor enthusiast who manages publicity for the Kayak and Canoe Club of New York, a whitewater enthusiast club.*

