Debevoise & Plimpton

Debevoise In Depth

Agentic AI in Retail Investing: Navigating Regulatory and Operational Risk

October 29, 2025

The Rise of AI-Driven Investing in Retail Finance

Generative artificial intelligence ("GenAI") innovations are rapidly transforming the formulation, analysis, and delivery of investment advice. Many broker-dealers and investment advisers are embracing GenAI to support one or more parts of the investment lifecycle—synthesizing investment research; undertaking trend analysis, anomaly detection, and pattern recognition for risk modeling and market surveillance; and performing large-scale data extraction and analysis.

One new focus is agentic AI: the use of AI to complete more than one task, either in series or in parallel, without any human involvement. The use of agentic AI in the investment selection process could be one of the most transformative yet challenging applications of GenAI in financial services. Gen AI investing experimenters have enthusiastically chronicled their novel investing experiences. For instance, a journalist documented his successful experience in prompting ChatGPT to design a "clever and highly aggressive" diversified high-growth portfolio that gained 10% in two weeks,[1] and a Reddit user crafted a profitable ChatGPT options trading prompt with profitability, loss, risk, and sector weighting parameters.[2] In an experiment run by a personal finance comparison website, a ChatGPT-designed basket of stocks generated nearly 42% in returns over a two-year period, beating comparable popular investment funds by as much as 23% during this timeframe.[3] Some financial services firms are likewise starting

[1] Thomas Smith, *I Gave ChatGPT $500 of Real Money to Invest in Stocks. Its Picks Surprised Me*, Fast Company (Sept. 22, 2025), https://www.fastcompany.com/91405657/chatgpt-invest-stocks.

[2] Civil Learning, *The Guy Who Let ChatGPT Trade for Him — and Somehow It* Worked, Medium (Oct. 9, 2025), https://medium.com/coding-nexus/the-guy-who-let-chatgpt-trade-for-him-and-somehow-it-worked-a5e81a911741.

[3] *Finder's ChatGPT Investment Fund is Still Outperforming the UK's 10 Most Popular Funds*, European Business Magazine (March 11, 2025), https://europeanbusinessmagazine.com/business/finders-chatgpt-investment-fund-is-still-outperforming-the-uks-10-most-popular-funds.

to experiment with using AI agents in the investment selection process—namely, testing whether GenAI tools can autonomously research, analyze, and then select potential investments.[4] Some investment pundits have suggested that financial advisors may eventually evolve into—much like the options trading Redditor—"*AI prompt engineers*" who design the perfect instructions for an AI to pick stocks, rather than serving as traditional stock-pickers.[5] That said, a different Reddit user recently went viral by directing OpenAI's ChatGPT to manage a live micro-cap stock portfolio—yet after initially outperforming the S&P 500 by nearly30%, at the 12-week mark the portfolio plunged below the S&P by approximately 24% and, as of the time of this publication, the portfolio is approximately 35% below the S&P.[6]

Handing over investment decision-making to agentic AI tools without meaningful human oversight and review creates significant regulatory and compliance risk. As we have discussed, the Trump SEC has declared that it plans to pursue AI-focused retail fraud and has already brought fraud charges involving AI misuse. FINRA has also issued cautionary guidance on GenAI usage by broker-dealers. Stitching together several individual AI tasks within a multi-part investment workflow can additionally create compounded risk that is greater than the sum of the risks posed by each individual step in the process—which in turn amplifies overall regulatory risk.

Moreover, as GenAI tools become more powerful and cheaper, they will be able to analyze vast sums of data on an ongoing basis and make time-sensitive recommendations for trading strategies. For this reason, it will be difficult for humans to review and sign off on the soundness of GenAI-generated recommendations in the time window that maximizes the value of the trade and before other investment professionals using AI spot and capitalize upon the same opportunity.

Accordingly, the prospect of agentic AI for retail investment selection and recommendation—with limited human intervention in certain key stages of the process—raises a critical question for compliance and legal teams at financial services firms: how to harness GenAI's analytical power to improve investment performance (and boost profits) in an understandable and explainable manner without compromising investor protection or violating regulatory obligations. In this post, we discuss the regulatory framework—and associated risks—applicable to agentic AI

[4] Jose Antonio Lanz, *AI Trading Bots Are Booming—But Can You Trust Them with Your Money?*, Yahoo! Finance (Aug. 3, 2025), https://finance.yahoo.com/news/ai-trading-bots-booming-trust-150102860.html.

[5] Greg Isenberg (@gregisenberg), Twitter (July 30, 2025, 9:09 AM), https://x.com/gregisenberg/status/1950544309515637126.

[6] Nathan Smith, *ChatGPT's Micro-Cap Portfolio: Week 17*, Substack (Oct. 26, 2025), https://nathanbsmith729.substack.com/p/chatgpts-micro-cap-portfolio-week-d4e.

investment and trading applications, as well as risk mitigation strategies for financial services firms to consider when evaluating such tools.

New Tools but Same Rules: Regulators Emphasize "Technology-Neutral" Obligations

The SEC and FINRA have each stated unequivocally that a registrant's regulatory obligations remain unchanged when leveraging GenAI to generate retail investment recommendations and advice. The existing framework squarely applies to agentic AI.

Broker-Dealer and Investment Adviser Registration. As a threshold matter, agentic AI investment and trading applications will need to consider broker-dealer registration under the Securities Exchange Act of 1934 (the "Exchange Act") and/or investment adviser registration under the Investment Advisers Act of 1940 (the "Advisers Act") depending on the applications' features and design.

Section 15(a) of the Exchange Act generally prohibits any "broker" or "dealer" from using the mails or any means or instrumentality of interstate commerce to induce, attempt to induce, or effect any purchase or sale of a security unless such person is registered as a broker-dealer with the SEC. Section 3(a)(4) of the Exchange Act defines a broker as a person "engaged in the business of effecting securities transactions for the account of others." Put simply, a "broker" is therefore an agent—not a principal—in securities transactions for profit. Further, any developer or operator of an agentic AI tool that is designed, operated, or sold to facilitate trade placement for others could be deemed a broker depending on whether the tool provides functionality that is deemed to involve "indicia" of effecting securities transactions for users.

Section 202(a)(11) of the Advisers Act defines an "investment adviser" as any person or firm that: (1) for compensation; (2) is engaged in the business of; (3) providing advice, making recommendations, issuing reports, or furnishing analyses on securities. Section 203(a) of the Advisers Act requires a firm that meets the definition of "investment adviser" to register with the SEC unless it meets certain exemptions or exclusions or falls within certain prohibitions. An agentic AI tool that generates and provides investment advice could potentially trigger the Advisers Act registration provisions.

Duties Under FINRA Rules and the Advisers Act. A broker-dealer's recommendations are subject to multiple disclosure-, care-, and conflict of interest-related obligations under Reg BI and FINRA rules, and a registered investment adviser's recommendations are subject to fiduciary duties of care and loyalty under the Advisers Act.

FINRA has stated that existing legal obligations involving the retail investment process—such as compliance with sales and supervision standards—apply with full force to a broker-dealer's use of GenAI. In June 2024, FINRA informed member firms that all existing securities laws and FINRA rules "continue to apply when member firms use [generative AI] or similar technologies in the course of their businesses, just as they apply when . . . firms use any other technology or tool."[7] For example, if registered representatives of a broker-dealer start relying on ChatGPT to provide stock recommendations, those recommendations must still be in the best interest of the customer under Regulation Best Interest under the Exchange Act ("Reg BI").[8] Likewise, FINRA Rule 3110 (Supervision) requires firms to maintain a reasonably designed supervisory system; if GenAI tools are part of the workflow, a broker-dealer's policies and procedures should address technology governance controls such as model risk management, data privacy, and ensuring that the tool's outputs are reliable and accurate.[9]

For registered investment advisers, GenAI-driven investment recommendations are still "investment advice" under the Advisers Act, and investment recommendations originating from agentic AI workflows are still subject to all fiduciary duty and otherwise applicable substantive Advisers Act obligations. As with a broker-dealer's care obligation under Reg BI, a registered investment adviser risks breaching its duty of care by following an AI's suggestions blindly.

These regulatory obligations mean that broker-dealers and registered investment advisers using AI tools in the investment selection process will need to prioritize explainability to be able to perform the assessments needed to determine whether an AI-generated recommendation satisfies regulatory standards. Such AI tools need to be capable of identifying which specific pieces of information contribute to an investment selection and to what extent; presenting that data in a format that can be timely reviewed and verified by a human investment professional; and demonstrating consistent and accurate predictive investment performance over time. Specifically, a broker-dealer using AI in the investment selection process must still be able to demonstrate under Reg BI that a recommendation was in the client's best interest upon

[7] FINRA, *Regulatory Notice 24-09, FINRA Reminds Members of Regulatory Obligations When Using Generative Artificial Intelligence and Large Language Models* (June 27, 2024), https://www.finra.org/rules-guidance/notices/24-09.

[8] Patrick Donachie, *AI-Generated Recommendations Can Still Fall Under Reg BI, FINRA Exec Warns*, WealthManagement.com (May 18, 2023), https://www.wealthmanagement.com/regulation-compliance/ai-generated-recommendations-can-still-fall-under-reg-bi-finra-exec-warns.

[9] FINRA, *Regulatory Notice 24-09, FINRA Reminds Members of Regulatory Obligations When Using Generative Artificial Intelligence and Large Language Models* (June 27, 2024), https://www.finra.org/rules-guidance/notices/24-09.

a consideration of costs, the investor's investment profile, and comparison to the risk and performance of reasonably available alternatives.

Risks of Overreliance on AI in Retail Investment Recommendations and Advice

Several major risks arise when AI is deployed as part of the investment lifecycle without meaningful human review or transparency. Simply put, there is no "GenAI made me do it" defense under the federal securities law or FINRA rules; from a compliance perspective, the firm and its human investment professionals own the recommendation and its consequences.

For instance, a generic, one-size-fits-all AI prompt that is not tailored to the profile of a retail customer or a client would likely fail to incorporate required consideration of an individual client's risk tolerance, liquidity needs, or investment objectives. A broker-dealer that simply relies on this output without further diligence would likely violate its obligation under Reg BI to ensure that it has a "reasonable basis" that the recommendation is in the customer's best interest, and a registered investment adviser would likely violate its fiduciary duty of care, as these obligations require consideration of the specific client's needs and investment objectives. Similarly, a broker-dealer or registered investment adviser that simply accepts the output of a GenAI stock selection tool without understanding the rationale or conducting further due diligence could also trigger violations of these obligations.

These risks are compounded in the GenAI context because many GenAI models can review thousands of pages of complex financial data in seconds, but operate as partial "black boxes," unable to identify exactly which pieces of information they relied upon to reach their decision. This lack of explainability can be problematic for registrants if they cannot explain and justify investment decisions to customers, clients, and regulators. Critically, this includes a duty to recommend an investment in the best interest of the client based on reasonably available alternatives. For instance, if a registrant operating an agentic AI tool cannot articulate why the tool recommended a particular security—for instance, what fundamentals or factors the tool relied on and what alternatives were available, including relative costs—it will be difficult to demonstrate that the recommendation was made consistent with the registrant's legal obligations before recommending it. Finally, agentic AI hallucinations have the potential to expose a firm to claims under the antifraud provisions of the federal securities law, Reg BI, and FINRA rules.

Moreover, because the federal securities laws require accuracy of disclosures made to investors in connection with the provision of securities recommendations and investment advice, broker-dealers and registered investment advisers must ensure that they do not inadvertently mislead clients about the use of agentic AI in connection with the investment process. A regulator could challenge a failure to disclose to a customer or a client that GenAI was involved in the generation of an investment recommendation, or disclosures minimizing or misstating the role of GenAI in the investment process and associated risks, if the facts about GenAI usage would be material to the customer or client's decision-making. GenAI-driven investment processes that are unprofitable or result in widespread investment losses could face especially significant regulatory risk.

The SEC's BlueCrest Capital matter illustrates potential disclosure risks. As we have discussed, the SEC charged BlueCrest Capital with failing to disclose that an algorithmic trading program was managing a significant portion of client assets—and that the algorithm was underperforming the firm's human traders. The SEC faulted BlueCrest for only providing generic references to "quantitative strategies" and for failing to disclose the risks and limitations of its algorithm, such as increased volatility and delayed execution, which materially affected client returns.

Finally, firms must ensure that they effectively supervise the use of GenAI in connection with the provision of investment advice. For example, if GenAI is providing support to research analysis or serves as a model portfolio provider, a registrant should document any AI outputs that influence investment recommendations to substantiate those recommendations. Additionally, firms must calibrate trade surveillance and compliance monitoring systems to detect patterns driven by GenAI-suggested trades. As we have also discussed, effective supervision of GenAI in connection with investment advice also requires building controls to prohibit off-systems GenAI usage by financial professionals.

Strategies to Mitigate AI-Related Risks in Investment Advice

As they continue to incorporate AI into investment selection, registrants should consider implementing several safeguards to reduce regulatory and legal exposure when using AI.

- **Clear Policies and Pre-Approval for AI Use**: Firms should update their policies and procedures to address whether and how investment professionals may use GenAI tools in providing investment recommendations or advice. For example, a policy might allow using GenAI for research or idea generation only, but prohibit trading on GenAI outputs without meaningful human review. Firms may also consider requiring compliance approval for using a particular GenAI application or use case so that the firm can assess the tool's reliability, data handling, and alignment with regulatory requirements. In practice, that could mean the firm's risk or IT team evaluates a GenAI model for accuracy on historical data, tests for biased outputs, and ensures that it will not learn or leak confidential information, before any financial professional relies on it. A related step is maintaining an updated inventory of approved (and prohibited) GenAI tools, similar to how firms track other IT resources, in order for the firm to stay in control of what technology is influencing client advice.

- **Training Financial Professionals on GenAI's Capabilities and Limits**: Firms should provide training to financial professionals that GenAI is a tool to augment, not replace, their expertise and judgment. This process includes educating them on the known risks—for example, that large language models can hallucinate, exhibit biases based on their training data, or become less effective outside the scenarios on which they were trained.[10]

- **Human Review and Approval ("Human in the Loop")**: To satisfy regulatory expectations of oversight and care, firms should institute checkpoints to ensure that a qualified investment professional independently evaluates any GenAI-generated investment advice before acting on it:

  - Financial professionals should be trained to double-check factual assertions from a GenAI tool against primary sources, and validate that any investment thesis makes sense.

  - Firms can also coach financial professionals on crafting better prompts to get more useful GenAI outputs, while making clear that the financial professional is ultimately responsible for the recommendation that goes to the customer or client.

  - Firms should consider having financial professionals internally document their own rationale for each retail investor recommendation that was influenced by GenAI and have a supervisor or second financial professional sign off in sensitive cases. A "human in the loop" can evaluate fiduciary and regulatory considerations that a GenAI tool may not be able to evaluate—for instance, by overruling the GenAI tool if its picks would over-concentrate a portfolio or conflict with other client-specific considerations.

[10] FINRA, *Podcast: An Evolving Landscape: Generative AI and Large Language Models in the Financial Industry* (Mar. 5, 2024), https://www.finra.org/media-center/generative-ai-llm.

  - If at any point the human professionals do not understand the GenAI's decision-making process, that is a red flag to pause and evaluate the deployment of GenAI. The human investment professionals need sufficient insight into the GenAI's tool's rationale behind any investment recommendation to independently assess its soundness.

- **Robust Disclosure and Client Communication**: As we have discussed here and here, the SEC has charged multiple firms with making fraudulent misstatements and omissions about AI use. Firms must ensure that disclosures about the role of GenAI in the investment process are accurate and complete, so as not to overpromise or understate their use of AI.

- **Monitoring and Surveillance of GenAI-Influenced Activity**: Just as firms monitor email, trade logs, and other financial professional activities, they should consider extending surveillance to the use of GenAI in the investment process—effectively treating GenAI like a third-party research provider whose input is auditable.

  - For example, compliance could periodically review samples of recommendations that had some GenAI basis to test the accuracy, quality, and appropriateness of the recommendation.

  - Surveillance systems could also be tuned to search for anomalies that could be tied to GenAI usage—for instance, an uptick in trading small-cap technology equities across many client accounts might warrant investigation to determine whether a GenAI tool is driving the trading and whether the security selection is suitable for all clients.

  - Much like firms that have created mobile messaging compliance programs in response to the SEC and CFTC off-channel enforcement sweeps, firms should adopt programs to ensure that financial professionals can access vetted GenAI tools to lower off-channel GenAI usage.

Conclusion: Embracing Innovation with Accountability

GenAI tools are rapidly proliferating across the investment landscape but pose particularly acute regulatory and compliance risks for retail-focused firms. Compliance officers and in-house counsel at retail-focused financial services firms should stay ahead of this trend by engaging with financial professionals who are exploring AI; updating policies and procedures to address the use of GenAI in investment recommendations and advice; and implementing strong oversight, documentation, and risk controls.

* * *

Please do not hesitate to contact us with any questions.

Andrew J. Ceresney
Partner, New York
+1 212 909 6947
aceresney@debevoise.com

Charu A. Chandrasekhar
Partner, New York
+1 212 909 6774
cchandrasekhar@debevoise.com

Avi Gesser
Partner, New York
+1 212 909 6577
agesser@debevoise.com

Robert B. Kaplan
Partner, Washington, D.C.
+1 202 383 8060
rbkaplan@debevoise.com

Julie M. Riewe
Partner, Washington, D.C. and
San Francisco
+1 202 383 8070
jriewe@debevoise.com

Jeff Robins
Partner, New York
+1 212 909 6526
jlrobins@debevoise.com

Kristin A. Snyder
Partner, San Francisco
+1 415 738 5718
kasnyder@debevoise.com

Achutha N. Raman
Law Clerk, New York
+1 212 909 6106
anraman@debevoise.com

This publication is for general information purposes only. It is not intended to provide, nor is it to be used as, a substitute for legal advice. In some jurisdictions it may be considered attorney advertising.