Hedge Fund Law Report

hflawreport.com

Sep. 24, 2026

Insider Trading

Can AI Trading Models Commit Insider Trading?

By Daniel Michael and Andrea Griswold, *Skadden*

Fund managers are increasingly deploying large language models (LLMs) and other artificial intelligence (AI) tools capable of accessing and analyzing nonpublic information at unprecedented scale. This trend raises pressing questions about how existing insider trading and material nonpublic information (MNPI) handling rules apply when the entity processing restricted data is not a human analyst but an AI model operating across vast datasets.

The stakes are significant. Broker-dealers and investment advisers must maintain policies and procedures reasonably designed to prevent misuse of MNPI, and managers risk regulatory scrutiny if AI tools foreseeably could use restricted data improperly – even absent actual trades. The question is not whether a machine can commit insider trading but whether people and managers have designed AI systems that foreseeably allow MNPI to influence trading or recommendations.

This article examines how the existing securities law framework applies to AI tools that access nonpublic information, identifies regulatory risks beyond traditional insider trading claims and offers practical guidance on controls fund managers can implement now – before AI-specific regulatory developments arrive.

For more on other new insider trading issues, see "Managing Prediction Market Insider Trading Risks" (Aug. 13, 2026); and "Insider Trading Enforcement Moves Beyond Equities" (Jul. 16, 2026).

How Existing Insider Trading Principles Apply to AI Tools

Insider trading under federal securities law requires these five elements:

1. a purchase or sale of a security;
2. trading "on the basis of" inside information;
3. information that is material and nonpublic;
4. a breach of a fiduciary or similar duty of trust and confidence; and

5. the requisite scienter.

Each element raises distinct considerations in the AI context.

Purchase or Sale

Continuing to hold securities, even if informed by MNPI, does not satisfy this element – there must be an actual purchase or sale. For AI tools, the relevant inquiry is whether the model's output triggers or informs a transaction in securities.

The "On the Basis of" Standard

Courts have diverged on what "on the basis of" requires. Some circuits treat knowing possession of MNPI at the time of a trade as sufficient, while others require proof that the trader actually used the information in making the trading decision. The SEC's Rule 10b5-1 reflects a preference for the knowing-possession standard, providing that a purchase or sale is "on the basis of" MNPI if "the person making the purchase or sale was aware of the [MNPI] when the person made the purchase or sale."

For AI models with access to MNPI, this standard is potentially expansive: if a model is "aware" of restricted information at the time it generates a trading recommendation or executes a trade, the knowing-possession standard could be met regardless of whether the model demonstrably "used" the information in its reasoning.

Materiality and Nonpublic Status

Information must be material – that is, important to the total mix of information available to a reasonable investor – and not broadly disseminated to the public. These standards apply to AI-processed information just as they do to information assessed by human analysts.

Breach of a Duty and Scienter

The duty element distinguishes illegal insider trading from legitimate informational-advantage trading. Corporate insiders owe duties to shareholders not to trade on material company information, and the misappropriation theory prohibits trading on information obtained under a duty or obligation of confidentiality, extending liability to tippees. At least in some federal circuits, a tipper need not know the precise legal character of the breach but must understand that disclosing the information would violate a confidence.

Applying these principles to AI, a fund manager may face regulatory risk when employees design, permission, deploy or supervise a model knowing – or consciously avoiding knowing – that it can use MNPI in breach of a duty. This does not require ascribing scienter to the model itself. The knowledge and conduct of employees acting within the scope of their responsibilities (data ingestion, model development, compliance approval, portfolio management, trading and surveillance) may, under certain circumstances, be attributed to the fund managers.

For example, an employee who knowingly grants a trading model access to MNPI subject to a confidentiality obligation possesses knowledge relevant to the fund manager's risk analysis. The SEC could establish scienter by focusing on the humans and governance processes surrounding the model, including:

- how data sources were reviewed and approved;
- whether anyone knew MNPI was included in the model's accessible data;
- how use restrictions and permissions were established and monitored; and
- whether a human-in-the-loop reviewed outputs before trading.

See "Benchmarking Fund Managers' Adoption and Governance of Generative AI" (Nov. 6, 2025).

How the Mosaic Theory May Leave Room for Nonpublic Data

The mosaic theory is important for AI-enabled research and trading. Analysts have long combined publicly available information, immaterial nonpublic information and industry data – such as credit-card spending data, app-usage metrics, satellite imagery and alternative data – to form investment conclusions. LLMs can amplify this analytical approach by identifying patterns across large, disparate datasets that exceed practical human review limits.

Nonpublic information may appropriately be part of such analysis if it is not material and not subject to a use restriction that would make trading on it improper. However, the risk is real: reasonable minds can disagree on materiality, and the mosaic becomes problematic if a single piece of information within it is both nonpublic and material. Thus, fund managers deploying AI tools that ingest diverse data streams must carefully evaluate whether any component of the model's accessible data crosses this threshold.

See "How to Apply Alt Data Best Practices to AI Systems" (Oct. 10, 2024).

Regulatory Risk Beyond Insider Trading

The SEC's enforcement toolkit extends beyond traditional insider trading claims. Section 15(g) of the Securities Exchange Act of 1934 requires broker-dealers, and Section 204A of the Investment Advisers Act of 1940 requires investment advisers, to establish, maintain and enforce written policies and procedures reasonably designed to prevent the misuse of MNPI. The SEC has brought enforcement actions when information barriers were not properly established, documented or respected – without alleging that actual insider trading occurred.

See "SEC Charges Hedge Fund Manager With MNPI Failures Related to Consultant" (Jan. 30, 2025); and "SEC Charges Broker-Dealer Over Information Barrier Failures" (May 9, 2024).

If a trading desk, investment team or AI tool has access to MNPI when misuse was foreseeable, the fund manager may face scrutiny over the reasonableness of its controls. This reasonableness standard is informed by prevailing industry standards adopted by market participants. Because industry standards for AI governance of

MNPI are still coalescing, there is no agreed-upon roadmap or checklist, creating both risk and opportunity for firms developing their frameworks now.

See "ACA Study Finds Widespread, but Limited, Implementation of AI" (Jul. 16, 2026).

Instructing a model not to trade in an issuer's securities after receiving MNPI about that issuer is an important control, but additional measures may be warranted. Non-disclosure agreements (NDAs) may prohibit using deal information for any financial benefit, and similar restrictions may appear in vendor terms, alternative-data licenses, portfolio-company agreements or joint-venture documents. The risk is not limited to trading in the securities of the issuer that supplied the information.

Some observers believe the SEC may be less interested in pursuing novel theories such as "shadow trading" (using MNPI about Company A to trade in Company B's correlated securities), but the Commission's interest in enforcing MNPI policies-and-procedures provisions – independent of actual misuse – appears undiminished.

See "Mitigating 'Shadow Trading' Risk in the Post-*Panuwat* Environment" (Aug. 28, 2025).

Practical Controls for AI Access to MNPI

Controls Beyond Model Instructions

Written instructions and system prompts directing a model not to use certain information for specific purposes are useful starting points, but managers should consider additional technical and procedural controls. In a 2023 U.K. study, researchers found that an AI bot disregarded instructions not to trade on certain information because it perceived that the risk of inaction outweighed the insider trading risk – and denied doing so when questioned. Even if newer models demonstrate more reliable instruction-following, managers should consider cross-deal contamination risk: AI agents ingesting data from multiple active matters into a shared knowledge base could enable inference of relationships or retrieval of confidential insights across workspaces.

Data Segregation and Permissions

Fund managers can extend traditional information-barrier concepts to AI architecture by:

- segregating data by deal team, strategy, issuer, account or contractual restriction;
- limiting which models or users can access restricted data;
- encrypting restricted data separately; and
- maintaining approval workflows before a model can use sensitive information for research, recommendations or trading.

Cloud and infrastructure tools may support these controls, but the legal question is whether the manager can demonstrate that MNPI access was limited; monitored; and aligned with confidentiality and trading obligations. Managers should also consider which employees have access to internal models and whether those employees

could access data ingested by those models – ensuring, for example, that AI models designed to respect information barriers between investment banking and equity research are matched by steps ensuring that employees on one side of a wall cannot access the model serving the other side.

See this four-part series "How Can Hedge Fund Managers Structure, Implement and Enforce Information Barriers to Mitigate Insider Trading Risk Without Impairing Securities Trading?": Part One (Jan. 16, 2014); Part Two (Jan. 23, 2014); Part Three (Jan. 30, 2014); and Part Four (Feb. 6, 2014).

Explainability, Audit Trails and Human Review

"Explainable AI" refers to techniques that give professionals a way to understand, trust and manage AI-generated decisions. Maintaining records of model inputs, reasoning steps and outputs can support compliance review, supervisory escalation and the ability to demonstrate that a particular trade was not informed by restricted information.

Explainability is not a cure-all, particularly when the knowing-possession standard applies, because mere awareness of MNPI may suffice regardless of whether the model demonstrably relied on it. Nevertheless, robust audit trails can help demonstrate that AI governance was designed to prevent and detect improper use of restricted information.

See "Risk Alert Cites Compliance Issues Regarding Advisers' Handling of MNPI" (May 19, 2022).

Contractual Restrictions, Data Provenance and Staleness

A disciplined inventory of the duties and restrictions attached to nonpublic information is a key safeguard. Vendor agreements, alternative-data terms, NDAs, portfolio-company agreements and joint-venture agreements may impose confidentiality or use restrictions affecting whether data can be used in AI-assisted research or trading. Fund managers should track the source of each restricted dataset; the scope of any use restriction; the securities or strategies affected; the personnel and systems permitted to access it; and the date the information was received. Data provenance also matters because stale information is less likely to remain material, while recently received information may warrant heightened approvals, access limitations and surveillance.

Key Governance Questions

Fund managers deploying AI tools with access to nonpublic information should consider the following questions now:

- What nonpublic data sources can the tool access?
- Which contractual or fiduciary restrictions apply to those sources?
- Can the tool use restricted information to inform or generate trading ideas, recommendations or portfolio decisions?

- Are deal-specific, issuer-specific or st gy-specific barriers reflected in the model's permissions?
- Are outputs logged in a way that compliance personnel can review?
- Can surveillance systems detect when AI-informed trading occurs in securities connected to restricted information?

See this two-part series "The Algorithmic CCO: AI's Role in Shapi he Future of Hedge Fund Governance": Part One (Feb. 13, 2025); and Part Two (Feb. 27, 2025).

ng t

Conclusion

te

Policies-and-procedures rules requiring broker-dealers and investment advisers to properly handle MNPI may pose more direct regulatory risk than traditional insider trading claims because they focus on the reasonableness of controls and the foreseeability of misuse – not just on whether a completed trade occurred. The reasonableness standard can work in fund managers' favor, however. It does not impose perfection or strict liability for every control failure but asks whether the manager designed, implemented, tested and enforced controls appropriate to its business, technology, data sources and trading activity.

Thoughtfully designed controls that encompass training, permissions, approvals, testing, documentation and surveillance may help demonstrate reasonableness even if an incident later raises questions. Although the timing of AI-specific regulatory or enforcement developments remains uncertain, managers can anticipate the applicability of existing MNPI obligations now.

ra

Managers that develop, test, deploy or surveil AI models should treat MNPI governance as part of the model lifecycle. The best-positioned managers will be those that can explain what data their models can access, why access is permitted, how use restrictions are enforced and how the firm detects and escalates potential misuse. Such managers will be better positioned to avoid costly technological redesigns when future regulatory guidance arrives.

See "The Core Benefits and Burdens of AI Use in Financial Services" (Apr. 23, 2026); and "Alternative Data and AI Becoming Integral to Investment Processes, Survey Finds" (Mar. 26, 2026).

Daniel Michael and Andrea Griswold[1] are partners in the white collar defense and investigations group at Skadden, Arps, Slate, Meagher & Flom LLP in New York. They focus on securities enforcement matters, internal investigations and regulatory compliance, with particular experience advising hedge funds and other investment advisers. Prior to joining Skadden, Michael was chief of the SEC Enforcement Division's Complex Financial Instruments Unit, and Griswold was the deputy U.S. attorney and chief of the Securities and Commodities Fraud Unit at the U.S. Attorney's Office for the Southern District of New York.

[1] The opinions expressed in this article are those of the authors and do not necessarily reflect the views of Skadden or its clients.

© 2026 Mergermarket Limited. All rights reserved.