



Research Article

THE FACELESS INSIDER: CAN AI TRADE?

Author: Mahak Laddha^{1*} 

¹ Shri Vile Parle Kelavani Mandal's Narsee Monjee Institute of Management Studies, Indore (M.P.) (India)

KEYWORDS:

AI, Insider Trading, Laws, UPSI, MNPI, Technology

ARTICLE HISTORY

Received Date: 11/07/2026

Revised Date: 15/07/2026

Accepted Date: 12/08/2026

Published Date: 06/09/2026

DOI: 10.69557/2czdkk06

CORRESPONDING AUTHOR

Mahak Laddha

mahakladdha9@gmail.com

ABSTRACT

"A nation that legislates only after loss has been suffered, purchases wisdom at the price of trauma. True governance foresees tomorrow before tomorrow exacts its cost."

‘The law is like the old grey streets of London, is a thing of worn stone and hidden corners.’ It does not move for the man who shouts, nor for the woman who weeps. Instead, it moves only for the careful step, the quiet word and the patient hand that knows its weight. But in the financial markets of tomorrow, the most dangerous trader has no face, no conscience and no loyalty; as it is an algorithm—a machine that reads, learns and acts in microseconds that is long before any human eye can even blink.

INTRODUCTION

This article explores unsettling questions: when artificial intelligence trades on patterns and traces it has uncovered from the corporate data, the individual customer history on its own platform or the information available publicly, does this amount to Insider Trading? And if it does, who do we blame? Our laws were written for men, not for machines. Existing securities laws are principally structured around legally cognizable people and entities whose duties, knowledge, conduct and economic interests can be identified and attributed. The emergence of autonomous or semi-autonomous trading systems therefore raises an attribution problem rather than automatically creating a new legal actor. Current laws rest upon concepts of duty, intent and personal benefit, such notions make no sense when applied to a computer. Yet the problem is no longer theoretical. Platforms like BlackRock's Aladdin already manage trillions of dollars through algorithmic decision-making, whereas retail investors increasingly rely on AI-driven recommendations without understanding their inner workings.

The gap between technological speed and legal response has become a crisis of fairness. We must ask ourselves: have we built a market where the machine always wins? And what happens to confidence when the trail of evidence ends not with a guilty party, but with a silicon chip? To sculpt the law for this new age, we must first learn its weight—not to break it, but to turn it, as a sculptor turns clay. It is not a weapon for the angry, but a tool for the wise. And like any fine tool, it works best in the hands of those who have taken the time to understand its nature.

Where the Trail Runs Cold

In the modern era of High-Frequency Trading ("HFT"), capital flows in microseconds. It takes place in the space between one heartbeat and the next—a thousand transactions, a million shares, a verdict on market confidence, all determined before a single human eye has had time to blink. There is no rogue trader on a dealing floor, no hushed phone call, no envelope stuffed with confidential documents. There is only a silent, relentless churn of an algorithm, that is a machine who reads, learns and acts with a speed, that makes the very concept of "intent" seem archaic. This is the new frontier of finance which poses a question that strikes at the very heart of our legal system: when a machine trades on patterns it has discerned from the data exhaust of corporate executives, is that insider trading? And if it is then against whom do we bring a charge? Who bears liability when the algorithm autonomously executes trades based on market-sensitive information—the developer, the institution or the algorithm itself?

A Law Written for Men, Not Machines

The legal definition of Insider Trading is deceptively crisp- it is trading in securities while in possession of Unpublished Price-Sensitive Information (UPSI) in breach of the duty of confidentiality. However, this definition was built around human actions with fiduciary duties.

The trail of an algorithm cannot be similar to the trail of a man. Traditional insider trading law hinges only on human narrative: a tip-off, a breach of trust, a conscious decision to profit from secret knowledge. The 1961 case of *In re Cady, Roberts & Co.* established the "disclose or abstain" rule, obliging those with a special relationship of trust to a company to either share material non-public information (MNPI) or refrain from trading on it. Also, the law articulated in case like *SEC v. Texas Gulf Sulphur*, was built upon the principle that "all investors should have equal access to the rewards of participation in securities transactions". It was a noble aim, forged to combat the inequities of boardroom secrets being used for personal gain.

Yet, this framework presumes a clear chain of causation: a person obtains information, knows it is confidential and thus decides to use it. An algorithm does not "know" anything. As it observes correlations and then works on it. The challenge is magnified by cases like *SEC v. Panuwat (2024)*, which broadened this concept of insider trading to include "**shadow trading**"— using material non-public information about one company to trade in the securities of an economically connected one. The case involved an

employee of Medivation, a pharmaceutical company, who learned that his company was about to be acquired. Instead of trading in Medivation's securities, he purchased call options in Incyte Corporation, another pharmaceutical company. The SEC argued that Panuwat used the confidential information about Medivation because he understood that the acquisition could also increase the value of shares of a similar company in the same industry. In simple terms, he did not trade in the company about which he had inside information; he traded in another company that he believed could benefit from the same event. The court allowed the SEC's insider-trading claim to proceed, giving rise to what is commonly described as “**shadow trading**.” The important point is that the case doesn’t establish that every trade based on a relationship between two companies is insider trading. Rather, it shows that confidential information about one company may, in appropriate circumstances, be relevant to trading in another company's securities. This reasoning raises an interesting question for AI-driven markets. While Panuwat itself involved a human trader, its core logic directly maps onto AI-driven markets: an AI system, fed with confidential data; for instance- about Company A, AI agents can instantly identify a web of economically linked firms—competitors, suppliers or peers—and execute trades in Company B based on predictive correlations that no human could process at speed. The AI does not “intend” to shadow trade; it simply detects that Company A's confidential event historically and correlates it with Company B's price movement and acts on that same pattern. This expansion exponentially increases the complexity for regulators, as the web of potential cross-firm linkages becomes far too vast for traditional surveillance methods to track effectively. It is in this complex, data-rich environment that AI operates, identifying correlations that a human analyst might miss or that might be entirely coincidental. This is where the trail of evidence dissipates. As legal scholars have pointed out, the law was built around a human decision-maker and the requirement of scienter—an intent to deceive, manipulate or defraud. In the automated world, as scholars Dörte Poelzig and Paul Dittrich note, the core challenge lies in attributing ‘knowledge’ & ‘conduct’ when trade-decisions aren’t made by humans but by autonomous AI.

Here, the very concept of a “culpable mind” is rendered meaningless when the actor has no mind at all. Generative AI has shattered this paradigm. It creates what scholar J. Scott Colesanti calls the ‘**Black Box**’ problem, where the budding AI models already permeating the market—correctly infiltrate firewalls and procedures to provide unlawful trading asymmetry to the investors. The law, with its focus on human duty and personal benefit, is not equipped for this new reality. The 1980 Supreme Court case *Chiarella v. United States* famously held that there could be no liability without a duty. This creates a profound legal vacuum.

Further, the analysis from NYU's Journal of Law & Business puts it, “*due to the ‘black-box’ nature of AI and its deep reinforcement learning systems that evolve without human intervention, large language models can autonomously generate outputs that resembles insider trading, market manipulation and other deceptive conduct, in contravention of their programming*”. The opacity of the algorithm makes it virtually impossible to reconstruct the “decision-making path” to determine whether it acted “on the basis of” MNPI, a prerequisite for liability under Rule 10b5-1. Understanding it deeply, directs us that the opacity of certain AI systems may complicate the evidentiary inquiry into the question that, whether a human trader or the relevant institution possessed or was aware of the Material Non-Public Information when a transaction occurred. However, the difficulty, is absolutely evidentiary and attributional rather than a categorical exemption from the existing legal standard. The relevant question remains whether the statutory and regulatory elements of insider trading can be established against an identifiable legal person or entity, including the applicable requirements concerning awareness, possession, use and trading ‘on the basis of’ Material Non-Public Information.

In the United States, this issue is primarily governed by Section 10(b) of the Securities Exchange Act of 1934 and SEC Rule 10b-5, which prohibit fraudulent or deceptive conduct in connection with the purchase or sale of securities. Rule 10b5-1 further clarifies when a person is considered to have traded “on the basis of” Material Non-Public Information (MNPI). Broadly, a person is considered to have traded on the basis of MNPI when they were aware of that information at the time of the trade, subject to the exceptions and affirmative defences provided under the Rule. Importantly, the law does not require regulators to prove that the trader consciously analysed the information, deliberately relied upon it, or made the trade because of it. Awareness at the relevant time can be sufficient for the “on the basis of” requirement, subject to the

applicable legal framework. At the same time, Rule 10b5-1 does not operate in isolation. The broader requirements of Section 10(b) and Rule 10b-5, including the applicable requirement of scienter, continue to govern the establishment of securities-fraud liability.

Rule 10b5-1(c) also provides an affirmative defence for certain trades made pursuant to qualifying trading arrangements adopted before the person became aware of MNPI. This becomes particularly relevant when trading is automated. The fact that a trade is executed automatically by an algorithm or AI system does not, by itself, bring the transaction within the protection of the defence; the underlying trading arrangement must satisfy the conditions prescribed by the Rule. The SEC's 2022 amendments further strengthened these requirements by introducing additional safeguards, including cooling-off periods and other conditions governing trading plans. Thus, AI may change the manner in which a trade is executed, but it does not change the legal requirements that determine whether the trade is permissible.

The real difficulty emerges when the trade is no longer the direct result of a human decision. Imagine an employee who possesses MNPI and feeds that information into an AI-driven trading system, which subsequently generates a trading signal and executes a transaction. The immediate question is not whether the AI itself can be treated as an insider, but whose knowledge and conduct should be attributed to the trade. The inquiry may therefore shift towards the employee who possessed the MNPI, the institution that deployed the system, the instructions given to the AI, and the extent to which the relevant information was known to or accessible by those legally responsible for the transaction. In contrast, if an AI system generates a trading signal solely from information that was lawfully available to the public, the fact that its prediction is unusually sophisticated, unexpected, or highly accurate would not, by itself, establish insider trading.

This is where AI exposes a difficulty that traditional insider-trading law was not designed to confront. The law is accustomed to asking what a person knew, when they knew it, and whether that knowledge was connected to the trade. AI introduces another layer between the information and the transaction, making it harder to identify how information travelled through the decision-making process and, more importantly, whose knowledge should legally matter. The problem, therefore, is not that AI should automatically escape insider-trading liability because it has no human mind. Rather, the challenge lies in connecting the AI-driven transaction to the human or institutional actors whose knowledge, instructions, or control may attract legal responsibility. The "faceless insider" is consequently not necessarily the AI itself, but the possibility that an algorithmic decision-making system can obscure the human knowledge and responsibility behind a trade.

Reading the Silence Between Disclosures

The most insidious challenge, may not be the theft of secrets but the analysis of the ordinary. A modern AI agent need not hack a corporate server to gain an edge; it simply watches and observes the momentum. As legal definitions have evolved, a core component of the prohibition is- trading "on the basis of" material non-public information, with a circuit split over whether this requires "use" or mere "knowing possession." But the AI bypasses this debate entirely. It constructs a mosaic from publicly available fragments that, in isolation, would mean nothing. This is the fundamental tension at the heart of the "faceless insider." The AI is not using non-public information; in furtherance it is creating new proprietary intelligence from the tracks and patterns in data that is theoretically public. The law of insider trading remains anchored to an older analog model of information asymmetry—the duty-based model established in *Dirks v. SEC*, focused on breach of duty in tipping confidential information. But the AI trader has no duty to the marketplace.

This system is peculiarly susceptible to manipulation. For instance, a company only needs to flood the digital space with fabricated posts of profits and deals—a form of SEO distortion—to train the AI Agents toward false conclusions. Such practices sit squarely alongside more recognised algorithmic abuses: **spoofing** (placing large orders without intent to execute), **layering** (multiple fake orders at varying prices to shift the market), **quote stuffing** (rapid-fire orders to overload systems), **wash trading** (simultaneous buy-sell orders creating false volume), **momentum ignition** (triggering price spikes to lure followers, then exiting), **pump-and-dump via bots** (coordinated buying and misleading information), **closing-price**

manipulation (late-day trades to skew settlement), and **cross-market manipulation** (leveraging one market to profit in another). In short, the AI's very strength—its pattern recognition becomes its vulnerability, as bad actors can game the public data it trusts, while the regulator's duty-based framework stands powerless against a machine that owes no allegiance to the market.

When Public Knowledge Becomes Private Capital

This brings us to a second, even more unsettling legal frontier: intellectual property. In situation where, an AI company scrapes the public transaction histories and publicly disclosed life events of corporate executives—their Form 4 filings, their charitable donations, their social media activity and repackages these disparate pieces of data into a predictive report sold to hedge funds, whose asset is the resulting pattern?

The data is public, but the insight is proprietary. This tension exists in a legal grey area. The SEC's recent regulatory efforts around AI have largely centered on disclosure, i.e., pursuing enforcement actions against AI-washing and recommending robust disclosures of AI's operational impact, notes one analysis but it has "not substantively addressed the more fundamental challenge AI poses as an autonomous actor to the liability regime underpinning securities regulation". The question is whether synthesizing public data at scale constitutes a new form of information asymmetry or a structural advantage that undermines the "fairness" principle on which the market is built, creating what regulators now term a "structural insider advantage."

The Machine Already at the Table

In relevance to the real-world operations, the machinery of AI-driven finance is already embedded in the system's core. The most tangible example of this algorithmic-first era is BlackRock's Aladdin platform. According to the current data, BlackRock alone manages over **\$15.3 trillion** as of the second quarter. Aladdin is not merely a tool; it is a decision-making infrastructure that processes the market and risk data to guide investment decisions for its own firm and hundreds of clients across the globe. It is now being enhanced with generative AI to provide "narrative" and commentary based on its proprietary analytics. With generative AI now integrated, the platform enhances decision speed.

When the world's largest asset manager runs on algorithmic pattern recognition, the line between "smart research" and a "structural information advantage" blurs at an industry level. An individual investor lacks this access. The existence of Aladdin highlights that AI does not simply augment human research rather it fundamentally changes the structure of market competition.

The concentration of similar quantitative risk models across major institutions may increase systemic risk by encouraging correlated market responses during periods of stress. The **2022 UK LDI crisis** has been cited by some commentators as illustrating how simultaneous model-driven portfolio adjustments can amplify liquidity pressures, although official investigations did not attribute the crisis to AI itself.

Inside The Algorithmic Veil: The Black Box Beneath The Recommendation

This opacity is not limited to institutional platforms like Aladdin. It filters down to retail investors through apps & sites like Groww, Angel One, Zerodha, Metabagg.ai, HDFC Sky, Upstock, etc. in Indian market and Streak, Trendlyne, Jarvis Invest, Alice Blue ANT etc. in the international market, which offer AI-driven stock screening, score-based markings to stocks, research tools, algorithmic strategy building, AI-driven portfolio optimization, wealth management and stock recommendations. The increasing reliance on AI in investment advisory demonstrates the technology's growing influence on retail and institutional securities markets The black box problem is most acute here:

What data are these bots actually trained on?

Is a user's own trading data being fed back into the model, creating a feedback loop that benefits the platform at the user's expense?

How much the population should rely on these AI agents?

The absence of a "training data nutrition label" is a regulatory disaster waiting to happen. The challenge is compounded by a related issue: the **threat of "AI-washing"** and outright fraud.

In September 2025, the SEC charged two Las Vegas residents with raising over \$4.7 million from more than 140 investors by falsely claiming they would profitably invest their money using automated trading bots, when in fact they were running a Ponzi scheme. This case underscores a core reality of the AI era: the promise of the 'faceless' trading algorithm is proving to be a powerful tool for deception and when the machine is a fraud, there is often a human face behind it. Also, this case showcases the lack of awareness in the existing population by portraying that, how much the population trusts the AI and invests into it without a prior research on it.

Where the Law Falls Behind

The disconnect between technological innovation and regulatory adaptation is no longer merely a policy concern; it has become a significant legal challenge. In India, for instance, SEBI's existing insider trading framework was developed for conventional trading environments and is not adequately tailored to address the complexities introduced by AI-driven trading systems. This regulatory gap enables the firms and the individuals to develop and deploy AI-powered trading tools, including subscription-based online platforms offering proprietary trading agents, without necessarily falling within the ambit of existing insider trading prohibitions or having proper, verified and legally valid registration of these sites and apps online.

Consequently, AI may be employed to devise trading strategies that closely resemble insider trading in their economic effect, while remaining outside the current legal definitions. Although regulators have acknowledged these emerging risks but the law continues to evolve slower than the technology it seeks to govern. This regulatory uncertainty has prompted proposals to expand the definition of algorithmic trading; however, the critics and caution that excessively broad definitions could inadvertently stifle innovation and potentially extend beyond the scope of the regulator's statutory authority (*ultra vires*). The core of the problem is that the law is not written for an actor that lacks intent but produces effects. The central issue, therefore, is attribution rather than the creation of a separate legal personality for AI. AI doesn't necessarily create a new category of insider-trading liability; rather, increasingly autonomous trading systems expose an attribution problem within existing securities law, requiring regulators and courts to determine how knowledge, control, duty, conduct and economic benefit should be attributed to the human or institutional actors responsible for deploying such systems.

While some legal academics propose a "**sliding-scale liability model**" that assigns responsibility based on the AI's design intent and the degree of human supervision, this is still a far cry from the clear-cut rules needed to guide the market. This opacity also presents a powerful defence for human traders: they can simply claim they relied on an AI's recommendation, thereby creating "*reasonable doubt*" as to whether they acted "on the basis of" inside information.

The Verdict Yet Unwritten

The story of the faceless insider is still being written. It is a test of whether our laws, so carefully crafted to catch the individual swindler and rogue trader, can protect the integrity of a market where the most powerful participants are no longer human. If an algorithm can read the public silence and disclosures, crunch the data and profit from patterns that humans never consented to share, is it a revolutionary tool for market efficiency, or is it a new kind of insider trading with no insider to arrest? The philosophy of "**algorithmic accountability**" must be forged into legal principle. The alternative is a market where confidence is eroded, not by a crisis of fraud, but by an existential crisis of fairness—a system where the

REFERENCE

1. BlackRock. (n.d.). *Corporate website*. <https://www.blackrock.com/>
2. Bloomberg Law Automation. (2025, September 5). *Vegas firm's options trading bots didn't really exist, SEC says*. Bloomberg Law. <https://news.bloomberglaw.com/health-law-and-business/vegas-firms-options-trading-bots-didnt-really-exist-sec-says>
3. Cady, Roberts & Co., 40 S.E.C. 907 (1961)
4. *Chiarella v. United States*, 445 U.S. 222 (1980).
5. Colesanti, J. S. (2024). Accountable AI and insider trading: The other “black box” problem. *California Western Law Review*, 61(1), Article 2. <https://scholarlycommons.law.cwsl.edu/cwlr/vol61/iss1/2>
6. Davis, G. (n.d.). *A literature review of the impact of artificial intelligence (AI) in banking, insurance, and investment* [Undergraduate thesis/research paper].
7. Groww. (2025, June 19). What is high-frequency trading? Everything you need to know. <https://groww.in/blog/what-is-high-frequency-trading>
8. LawFoyer International Journal of Doctrinal Legal Research 3, no. 1 (2025), ISSN 2583-7753, <https://www.lijdlr.com/>
9. Poelzig, D., & Dittrich, P. (2025, August). *Artificial intelligence and EU insider regulation* (ECGI Law Working Paper No. 866/2025). European Corporate Governance Institute. <https://www.ecgi.global/publications/working-papers/artificial-intelligence-and-eu-insider-regulation>
10. Practical Law Corporate & Securities. (2022, December 15). *SEC adopts amendments to Rule 10b5-1 and related disclosure requirements*. Practical Law Legal Update.
11. Securities and Exchange Board of India. (2015). Securities and Exchange Board of India (Prohibition of Insider Trading) Regulations, 2015 (as amended June 26, 2024). <https://www.sebi.gov.in/>
12. *Securities & Exchange Commission v. Panuwat*, No. 3:21-cv-06322-WHO (N.D. Cal. Apr. 5, 2024).
13. SEC v. Texas Gulf Sulphur Co., 401 F.2d 833 (2d Cir. 1968).
14. SEO distortion occurs when search rankings are influenced by optimization tactics or manipulation instead of genuine quality or relevance (Ningsih & Setiawan, 2024).
15. The dataset nutrition label: A framework to drive higher data quality standards. (2020). <https://doi.org/10.5040/9781509932771.ch-001>
16. Torkian, V., Bamdad, S., & Sarfaraz, A. H. (2026). Integrating AI and OR for investment decision-making in emerging digital lending businesses: A risk-return multi-objective optimization approach. *Journal of the Operational Research Society*, 77(3), 716–735. <https://doi.org/10.1080/01605682.2025.2498652>
17. UK House of Commons Work and Pensions Committee. (2023, June 23). *Defined benefit pensions with liability driven investments: Seventh report of session 2022–23*. House of Commons. <https://publications.parliament.uk/pa/cm5803/cmselect/cmworpen/826/report.html>

18. U.S. Securities and Exchange Commission. (2000). *Selective disclosure and insider trading* (17 C.F.R. § 240.10b5-1). <https://www.ecfr.gov/current/title-17/chapter-II/part-240/section-240.10b5-1>
19. Vieru, S. (2026). *The algorithmic illusion: Mathematical proof of the structural harm of the “Big Six” and the Aladdin system to global financial stability: Analysis through the lens of the SVWORLD 73.5 v1.9 standard* [Preprint].