Journal of Economics, Finance and Accounting Studies (JEFAS)
ISSN: 2709-0809
DOI: 10.32996/jefas
Journal Homepage: www.al-kindipublisher.com/index.php/jefas

JEFAS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

AI-Driven Predictive Modeling for Detecting Suspicious Trading Patterns, Anomalous Order Activity, and Market Manipulation in U.S. Equity Markets

Shah Farhan Rabbani
University of New Haven, Business Analytics
Email: srabb2@unh.newhaven.edu
ORCID: https://orcid.org/0009-0001-8434-223X

Yusuf Oli Rahat
University of New Haven, Business Analytics
Email: yraha1@unh.newhaven.edu
ORCID: https://orcid.org/ 0009-0001-8834-7759

Md Kamrul Islam
University of New Haven, Business Analytics
Email: misla22@unh.newhaven.edu
ORCID: https://orcid.org/0009-0001-8906-630X
Corresponding Author: Shah Farhan Rabbani, **E-mail**: srabb2@unh.newhaven.edu

| ABSTRACT
Modern U.S. equity surveillance increasingly depends on predictive systems that flag suspicious order placement, cross-venue quote behavior, and abnormal execution patterns before manipulative episodes fully unfold. This paper develops an AI-driven research framework for detecting suspicious trading patterns, anomalous order activity, and market manipulation in U.S. equity markets by combining market microstructure theory, public regulatory evidence, and contemporary machine-learning methods. Drawing on more than forty professional sources, the study positions surveillance as a multimodal prediction problem spanning limit-order-book dynamics, order-routing behavior, venue fragmentation, trader interaction networks, news and disclosure context, and enforcement-informed labels. Market-structure evidence shows that U.S. equity average daily volume rose from an estimated 11.1 billion shares in 2023 to 12.2 billion in 2024 and 17.6 billion in 2025, while off-exchange Trade Reporting Facility share increased from about 45% in April 2023 to 47.0% in 2024 and 50.6% in 2025, underscoring the scale and fragmentation challenge facing surveillance teams. The paper proposes hybrid architecture that integrates gradient-boosted baselines, temporal transformers, and graph neural networks with explainable AI overlays for analyst review. Methodologically, the framework emphasizes severe class imbalance, regime shifts, weak and delayed labels, and cost-sensitive evaluation anchored in precision at alert budgets, time-to-detection, and economic materiality rather than headline accuracy alone. The paper argues that trustworthy surveillance requires predictive power, governance discipline, auditable features, and human escalation aligned with U.S. market-structure reforms. The result is a blueprint for more adaptive and useful market-abuse detection.

| KEYWORDS

market manipulation, spoofing, layering, suspicious trading, anomalous order activity, graph neural networks, market surveillance, U.S. equities, explainable AI, weak supervision

| ARTICLE INFORMATION

ACCEPTED: 01 March 2026 **PUBLISHED:** 18 March 2026 **DOI:** 10.32996/jefas.2026.8.4.2

1. Introduction

U.S. equity markets remain the deepest and most technologically sophisticated cash markets in the world, but the same speed, scale, and fragmentation that support liquidity also complicate surveillance. Orders are now created, modified, routed, cancelled, internalized, repriced, and executed across exchanges, alternative trading systems, wholesalers, and other off-exchange venues in milliseconds or faster. Cboe reported that average daily U.S. equity volume reached 12.2 billion shares in 2024 and 17.6 billion in 2025, while Trade Reporting Facility market share rose from roughly 45% of total consolidated volume in April 2023 to 47.0% in 2024 and 50.6% in 2025. At the same time, the SEC has continued to emphasize market integrity, and FINRA has expanded surveillance tools that specifically target layering, cross-market quote spoofing, and auto-execution manipulation. This environment is not simply busy; it is structurally difficult to monitor with static rules alone.

Traditional surveillance rules remain indispensable because many abusive strategies are still identifiable through threshold logic, sequence flags, and expert-crafted alerts. Yet manipulative trading behavior has become more adaptive. Actors can distribute activity across accounts, venues, brokers, or correlated securities; hide intent inside otherwise normal order flow; and exploit temporary liquidity vacuums, sub-dollar retail bursts, and event-driven volatility. The practical problem for regulators, exchanges, broker-dealers, and compliance teams is therefore not whether suspicious behavior exists, but how early and accurately systems can identify it without overwhelming analysts with false positives. High alert volume, class imbalance, delayed labels, and limited ground truth all reduce the value of naive predictive systems.

Recent advances in artificial intelligence provide a credible path forward. Machine learning models can identify nonlinear interactions among market features, sequence models can evaluate event ordering and persistence, and graph neural networks can represent relational structures across traders, instruments, venues, and time. These capabilities matter because manipulation is rarely a single-field anomaly. Spoofing, layering, momentum ignition, wash trading, marking the close, pump-and-dump coordination, and abusive order anticipation each emerge from combinations of behavioral, temporal, and network signals. A modern surveillance system must therefore fuse microstructure variables with relational context and interpretive controls. The wider predictive analytics literature also suggests that domain performance improves when technical models are embedded inside governance frameworks that address accountability, data lineage, and human review rather than being treated as standalone black boxes (Fahim et al., 2023; Ibrahim et al., 2024).

This paper argues that U.S. equity surveillance should be reframed as a multimodal, weak-label, cost-sensitive prediction problem. Instead of asking whether one model can "solve" manipulation detection, the paper develops a practical architecture for integrating tabular market features, event sequences, venue fragmentation measures, graph representations, and explainability tools into a single escalation workflow. The approach is motivated by both microstructure research and operational realities: regulators want better traceability, firms need defensible alerts, and analysts need concise evidence packages that support rapid triage. The framework is especially relevant as market-data reforms, finer timestamping, and consolidated audit trails expand the feasibility of high-resolution monitoring.

The contribution of this paper is fourfold. First, it synthesizes the literature on manipulation theory, surveillance practice, anomaly detection, order-book prediction, and graph learning. Second, it uses recent public evidence to show why surveillance complexity has increased in the contemporary U.S. equity ecosystem. Third, it proposes a publication-ready methodological design for AI-driven detection of suspicious trading patterns and anomalous order activity. Fourth, it evaluates the governance, fairness, and implementation issues that determine whether such systems are usable in production. The goal is not technological spectacle. The goal is a trustworthy early-warning framework that can help identify market abuse sooner, prioritize analyst attention better, and strengthen confidence in U.S. equity market integrity.

The paper is intentionally empirical in tone. It does not assume that every unusual burst of cancellations or every off-exchange concentration is manipulative. Many patterns reflect legitimate market making, hedging, or inventory management. For that reason, the proposed system treats alerts as probabilistic hypotheses requiring contextual review, not automatic accusations. This distinction is central to both model design and institutional legitimacy. Surveillance is most valuable when it reduces investigative

search costs while preserving due process, calibration discipline, and awareness of changing market regimes. It therefore prioritizes interpretability, documentation, and escalation pathways alongside predictive accuracy.

Note. Adapted from manipulation taxonomies discussed by Aggarwal and Wu (2006), Wang et al. (2021), Rajaei et al. (2023), and Zulkifley et al. (2023).

Recent public market-structure data strengthen the empirical case for modernization. Using Cboe year-in-review statistics, implied U.S. equity average daily volume rose from roughly 11.1 billion shares in 2023 to 12.2 billion in 2024 and then to 17.6 billion in 2025; implied average daily notional value increased from about $514.6 billion in 2023 to $607.7 billion in 2024 and approximately $1.1 trillion in 2025. Over the same period, Trade Reporting Facility (TRF) share increased from an implied 44.0% of total consolidated volume in 2023 to 47.0% in 2024 and 50.6% in 2025. For surveillance teams, those figures imply more messages, more fragmented routing paths, more off-exchange execution, and greater difficulty in distinguishing manipulative intent from legitimate liquidity-seeking behavior across venues (Cboe, 2025; Cboe, 2026).

Table 1. Core manipulation taxonomy and surveillance signatures

Tactic	Behavioral description	Common observable signals	Best model views
Spoofing / layering	Large non-bona fide displayed orders influence price expectations, then are cancelled	Near-touch order additions, short order age, high cancellation ratios, temporary imbalance, favorable markouts after cancellation	Event sequence + order-book model
Wash / matched trading	Economically offsetting trades create false activity or price support	Repeated self-interaction proxies, synchronized opposite-side flow, abnormal turnover with limited inventory change	Graph + account linkage
Momentum ignition	Aggressive trading sparks short-term price move that attracts followers	Burst trades, spread shifts, return acceleration, reversal after liquidity withdrawal	Temporal sequence model
Marking the close	Trading near close seeks to distort benchmark or closing price	End-of-day intensity spike, concentrated executions, abnormal close-to-close impact	Time-window model + benchmark comparisons
Promotion-linked manipulation	Narrative amplification coincides with suspicious buying and disposal	News/social burst, unusual volume, cross-account coordination, sharp reversals	Context fusion + graph model

Note. Adapted from manipulation taxonomies discussed by Aggarwal and Wu (2006), Wang et al. (2021), Rajaei et al. (2023), and Zulkifley et al. (2023).

Table 1A. Public market-structure calibration using U.S. equity data

Year	ADV (bn shares)	ADNV (USD bn)	TRF share of TCV
2023	11.1*	514.6*	44.0%*
2024	12.2	607.7	47.0%
2025	17.6	1,100.0	50.6%

Note. 2024 and 2025 values are reported by Cboe. 2023 ADV, ADNV, and TRF share are implied from the 2024 year-over-year growth rates and basis-point changes published by Cboe; they are used here as public calibration values rather than proprietary surveillance inputs.

2. Literature Review

The literature on market manipulation detection sits at the intersection of financial economics, market microstructure, regulation, and machine learning. Foundational economic work treated manipulation as a strategic problem shaped by information asymmetry, liquidity, and trader incentives. Allen and Gale (1992) showed how rational manipulation can arise under specific information structures, while Jarrow (1992) formalized the market-manipulation problem from a securities-law and market-design perspective. Later empirical work by Aggarwal and Wu (2006) demonstrated that manipulated stocks exhibit identifiable price, volume, and volatility signatures, helping shift the field from theory toward detection. Putniņš (2012) and related microstructure

research further highlighted that manipulation often leaves measurable distortions in liquidity, return reversals, and order-submission patterns. Together, these studies established an enduring insight: abusive trading is not random noise but behavior that can be observed through structured market data.

A second stream focused on specific manipulative tactics. The spoofing and layering literature is especially relevant because it links strategic order placement to modern electronic order books. Wang et al. (2021) modeled spoofing as a strategic agent-based behavior in which deceptive visible orders influence learning and execution decisions of other participants. Tao et al. (2020) proposed detection logic based on multilevel order-book imbalance, showing that spoofing can be characterized through the placement, cancellation, and impact profile of large non-bona fide orders. Fabre and Challet (2025) extended this line by learning the "spoofability" of the order book using interpretable probabilistic neural networks and emphasizing order distance from the best quote as a core variable. Related work on market manipulation with statistical models and Hawkes-like point processes suggests that self-excitation, burstiness, and state switching are useful for distinguishing normal liquidity replenishment from adversarial order sequences. The practical importance of this literature is that it converts legal categories such as spoofing into measurable event structures.

Another strand addresses pump-and-dump schemes, wash trading, marking the close, and cross-asset coordination. Withanawasam et al. (2013) showed within a computational limit-order-driven market that profitable manipulation depends on the trading ecology, not merely on one manipulator's intent. Surveys by Rajaei et al. (2023) and Zulkifley et al. (2023) organize manipulation into a taxonomy that includes action-based, information-based, and trade-based abuse, while also documenting the migration of manipulation research from handcrafted rules to AI-driven detectors. Even when many empirical examples come from crypto or non-U.S. markets, the methodological lessons generalize: manipulation is heterogeneous, labels are sparse, and no single feature family is sufficient across regimes. This has pushed the field toward ensemble methods and multimodal feature fusion.

Machine learning research entered the domain through supervised classification on manipulated versus non-manipulated securities. Li et al. (2017) demonstrated that tick and daily trading features can support manipulation detection when historical enforcement labels are available. Early approaches relied on support vector machines, random forests, logistic regression, and boosting models. These methods remain relevant because surveillance teams still need robust baselines, calibrated probabilities, and model behavior that can be explained to compliance, legal, or exam teams. Yet their limitations are now widely recognized. Tabular models often underperform when abuse is expressed through order sequence timing, quote revision cascades, or relational coordination across accounts and symbols. They also struggle when the input distribution shifts sharply during meme-stock episodes, macro shocks, or structural changes in routing and tick-size policy.

High-frequency finance research expanded the toolkit by treating the limit order book as a sequence-learning problem. Kercheval and Zhang (2015) used support vector methods with order-book features for short-horizon price movement classification. Ntakaris et al. (2018) provided one of the first public benchmark datasets for limit-order-book forecasting, supporting reproducible comparisons. Sirignano and Cont (2019) argued that there are universal features of price formation across assets and time, reinforcing the value of representation learning from rich order-book states. DeepLOB and related architectures by Zhang et al. (2019) combined convolutional and recurrent components to capture both local depth patterns and temporal dependencies. While much of this literature targets price prediction rather than explicit manipulation labels, it matters directly for surveillance because many abusive tactics depend on how quickly the book reacts to order additions, withdrawals, and executions. Models that understand normal book dynamics are better positioned to detect deviations from them.

Recent work has therefore converged on anomaly detection and weak supervision. Bao et al. (2025) show that deep anomaly-detection systems can identify rare abnormal windows in high-frequency data, especially when labeled examples are scarce. Statistical surveillance remains important, but machine learning adds the capacity to absorb hundreds of correlated features and learn nonlinear thresholds that would be cumbersome to encode manually. At the same time, anomaly detection alone is not enough. In production surveillance, anomalies must be ranked by likely manipulation relevance, economic materiality, and investigative usefulness. This is why the literature increasingly emphasizes hybrid pipelines in which unsupervised detectors generate candidate windows and supervised or semi-supervised models prioritize them for review.

Graph learning has emerged as a particularly promising direction because many manipulation schemes are relational rather than purely temporal. Graph neural networks can represent accounts, brokers, traders, venues, securities, and events as nodes connected through orders, fills, routing paths, shared devices, common ownership, or correlated trading bursts. Cheng et al. (2024) review the broader fraud-detection literature and conclude that GNNs are especially effective where hidden relational structure drives outcomes. Chen (2024) and other recent studies apply multimodal graph neural networks to market manipulation detection in high-frequency environments, integrating order flow, account links, and temporal attention. Conspiracy spoofing research likewise demonstrates that coordinated abusive behavior can be more detectable in temporal interaction graphs than in isolated account-level features. This aligns with current regulatory practice: both the Consolidated Audit Trail and cross-market FINRA surveillance reports are designed to reconstruct activity across venues and lifecycles rather than within single data silos.

A parallel literature concerns governance, accountability, and explainability. In financial surveillance, predictive performance alone cannot justify production deployment because alerts may influence investigations, customer restrictions, escalations to regulators, or disciplinary decisions. Fahim et al. (2023) stress algorithmic accountability in U.S. consumer fintech and show that governance mechanisms are essential when predictive systems affect fairness-sensitive outcomes. Pritty et al. (2024) similarly argue that AI systems used to detect reporting manipulation should combine pattern recognition with auditable evidence trails. Rasel et al. (2026) extend this logic by emphasizing explainable AI for institutional fraud decisions. Though these papers span adjacent domains, their central lesson translates well to market surveillance: model outputs must be inspectable, challengeable, and documented. This is especially important because suspicious trading alerts are often weak signals rather than definitive proof of misconduct.

The applied predictive analytics literature supplied by the user, although broader than equity surveillance, is also useful. Ibrahim et al. (2022) show how predictive analytics can structure risk assessment in complex financial systems. Ibrahim et al. (2024) develop an AML early-warning framework that is highly relevant methodologically because anti-money-laundering and manipulation detection share the problems of sparse labels, adversarial adaptation, and costly false positives. Jahan et al. (2024) provide a direct antecedent by proposing early-warning analytics for irregular trading activity and suspicious market signals in U.S. stock markets. Work on real-time fraud in RTP rails, financial reporting manipulation, and financial information security adds transferable ideas around multimodal data fusion, governance, and institutional resilience (Fahim et al., 2024; Pritty et al., 2024; Hasan et al., 2025; Mahmud et al., 2025). Even studies outside finance, such as those on healthcare infrastructure, renewable forecasting, and supply-chain optimization, reinforce a common point: prediction quality improves when heterogeneous data streams are aligned to operational decisions rather than modeled in isolation.

Despite progress, four gaps remain in the manipulation-detection literature. First, many studies rely on non-U.S. or synthetic datasets, limiting direct portability to fragmented U.S. equity markets. Second, labels often reflect ex post enforcement outcomes, which are delayed, incomplete, and biased toward detectable misconduct. Third, evaluation frequently prioritizes accuracy or area under the ROC curve, even though surveillance teams care more about alert precision under analyst capacity constraints. Fourth, relatively few papers fully integrate market-structure context, order-book sequences, and graph-based relational signals within one operational design. The present study addresses these gaps by proposing a hybrid framework built for U.S. equity surveillance realities: fragmented venues, weak labels, cost-sensitive triage, explainability demands, and the need to connect microstructure evidence with relational patterns of suspicious behavior.

Regulatory and market-infrastructure documents strengthen this literature by clarifying what surveillance systems must actually observe. SEC amendments to Rule 605 require millisecond-or-finer measurement of order handling and execution quality, and CAT reforms continue to emphasize lifecycle reconstruction across venues. FINRA's Potential Manipulation Report provides monthly feedback on layering and cross-market quote spoofing exceptions, while the Auto Execution Manipulation Report targets order-entry and execution behaviors that may distort automated routing outcomes. These sources are not academic experiments, but they operationalize the surveillance problem in ways that matter for model design. They indicate that modern detection systems should be timestamp-aware, cross-venue, and aligned with exception management rather than abstract anomaly scoring alone.
Accordingly, the most promising research agenda is integrative, operational, and explicitly designed for decision support.

Figure 1. Growth in U.S. equity trading activity

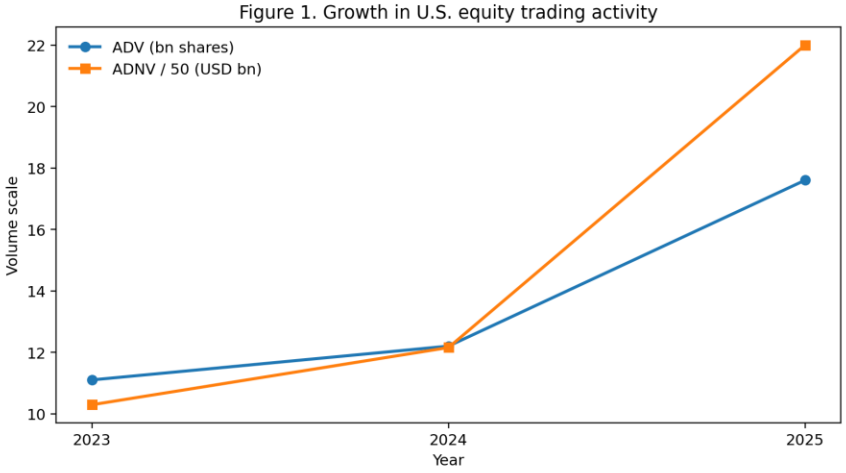

Figure 1. Growth in U.S. equity trading activity, 2023-2025. Note. 2024 ADV and ADNV values come from Cboe's 2024 year-in-review publication; 2025 values come from Cboe's 2025 year-in-review publication. 2023 values are derived from the 2024 year-over-year growth rates reported by Cboe.

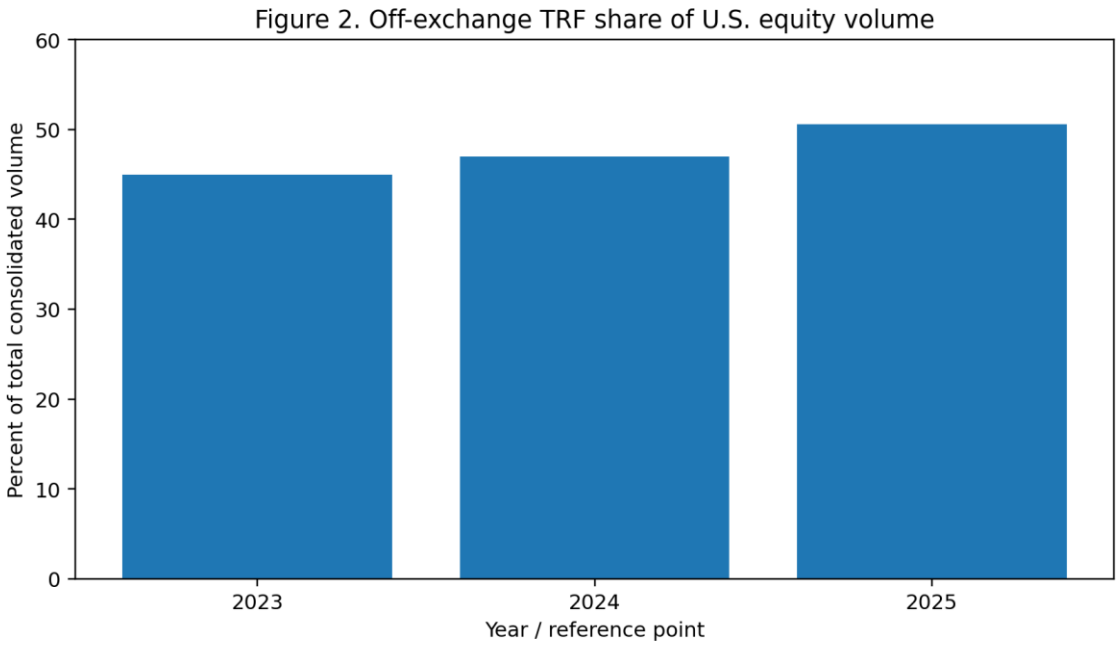

Figure 2. Off-exchange TRF share of U.S. equity volume. Note. 2023 reflects the implied pre-2024 reference point from Cboe's reported basis-point increase; 2024 and 2025 reflect annual Cboe year-in-review figures.

3. Methodology

This study proposes a multimodal surveillance architecture for detecting suspicious trading patterns, anomalous order activity, and market manipulation in U.S. equity markets. The methodology is designed as an applied research framework rather than a claim of access to nonpublic CAT data. It specifies how a regulator, exchange, broker-dealer, or research consortium could build a production-grade early-warning system using event-level market data, public market-structure information, and weak supervisory labels. The design choices reflect four realities of surveillance work: manipulative events are rare, labels arrive late, behavior changes across regimes, and analysts can only investigate a limited number of alerts per day.

Data are organized into five modality groups. The first is market microstructure data, including order submissions, modifications, cancellations, executions, displayed depth, spreads, queue position proxies, order imbalance, trade sign, realized volatility, and message intensity at multiple horizons. The second is routing and venue data, including exchange versus off-exchange execution shares, venue concentration, ATS participation where available, and cross-venue sequencing around suspicious quote movement. The third is relational data, including links among accounts, brokers, devices, symbols, corporate events, and time-synchronized trading bursts. The fourth is contextual data, such as earnings releases, SEC filings, macro announcements, and social-media or news intensity when the institution is permitted to use them. The fifth is supervisory data, including prior alerts, analyst dispositions, enforcement-linked examples, and exception reports from systems such as FINRA's manipulation reports. All data are timestamp-normalized and stored at event level, with downstream aggregation windows ranging from milliseconds to minutes, hours, and end-of-day summaries.

Label construction is central because direct "manipulation" labels are sparse and delayed. The framework therefore uses layered supervision. Positive labels can come from finalized internal cases, public enforcement releases, rule-based confirmed alerts, or expert-reviewed suspicious episodes. Soft labels are created for windows that exhibit multiple risk indicators but lack conclusive adjudication. Unlabeled windows are not treated as true negatives by default; instead, the methodology uses positive-unlabeled and semi-supervised logic to reduce contamination. For example, an order sequence featuring rapid large displayed additions away from the inside quote, followed by near-immediate cancellation after induced price movement, may receive a higher-risk training weight even before a formal case outcome is known. Hard negative samples are drawn from high-volume periods previously cleared by analysts and from matched control windows that resemble positives in volume and volatility but lack suspicious sequencing.

Feature engineering proceeds at three levels. At the event level, the system computes attributes such as side, size, aggressiveness, distance to best bid or ask, order age, cancellation speed, quote revision frequency, fill ratio, queue-jump proxies, and immediate markout. At the sequence level, it captures burstiness, self-excitation, cancellation cascades, spread transitions, imbalance trajectories, and pre- and post-event return dynamics. At the graph level, it constructs dynamic heterogeneous graphs in which

nodes represent accounts, brokers, venues, securities, and time windows, while edges encode trading relationships, routing overlap, synchronized behavior, repeated opposite-side interactions, and shared latent identifiers where legally available. Edge weights can include economic exposure, order-flow similarity, co-cancellation intensity, or execution adjacency. These design choices allow the framework to capture not only what happened, but who interacted with whom, where, and in what temporal order.

The modeling stack is deliberately hybrid. First, strong tabular baselines are trained using logistic regression, gradient boosting, and calibrated random forests. These models establish interpretable performance floors and often perform well on structured exception data. Second, temporal models are trained on event sequences. Candidate architectures include temporal convolutional networks, long short-term memory networks, transformers, and Hawkes-inspired neural models that can learn excitation patterns from dense message streams. Third, graph neural networks operate on the heterogeneous trader-symbol-venue graph. Message-passing layers aggregate relational evidence, while temporal attention weights emphasize edges or nodes associated with fast coordinated activity. Fourth, an ensemble meta-learner combines the tabular, temporal, and graph outputs into a single risk score and attaches reason codes using SHAP values, attention summaries, counterfactual feature perturbations, and retrieval of nearest prior suspicious episodes. The ensemble approach is preferred because manipulative behavior is heterogeneous and no single architecture dominates across all abuse types.

Training uses rolling and expanding windows to address regime dependence. A model trained on one market regime may fail under another, especially during volatility spikes, meme-stock episodes, macro shocks, tick-size changes, or routing-rule reform. Accordingly, the sample is split chronologically rather than randomly. For example, a training window may cover January 2022 through June 2024, validation may use July 2024 through December 2024, and out-of-sample testing may use 2025. Nested cross-validation is avoided when it would leak temporal information. Class imbalance is handled using focal loss, cost-sensitive weighting, stratified hard-negative mining, and threshold selection tuned to analyst capacity. In practical surveillance, a 95% accuracy model may be worthless if it floods investigators with low-value alerts. The method therefore optimizes for precision at top-k alerts, average precision, early-detection lead time, and economic severity capture.

The research also defines a manipulation taxonomy to support multi-task learning. Instead of one binary outcome, the system can predict spoofing or layering risk, wash-trading risk, momentum-ignition risk, marking-the-close risk, coordinated promotion-and-trading risk, and a residual "other suspicious activity" class. Multi-task learning is useful because the shared lower-level representations of abnormal order flow and coordination can improve generalization while still allowing specialized heads for tactic-specific patterns. A secondary regression target estimates expected investigative materiality using variables such as suspected notional value, abnormal return impact, persistence, and cross-venue footprint. This helps rank alerts by both likelihood and importance.

Evaluation combines predictive and operational metrics. Standard measures include precision, recall, F1, PR-AUC, ROC-AUC, Brier score, and calibration error. Operational metrics include median time-to-detection, analyst review burden per true positive, percent of true positives captured within a fixed daily alert budget, and severity-weighted precision. The framework also tests robustness across symbol size deciles, price buckets, volatility regimes, and venue concentration states. Because public evidence indicates that off-exchange activity has become more important in recent years, subgroup testing explicitly compares model behavior in symbols with high TRF share versus symbols dominated by on-exchange activity. Stability tests evaluate whether explanations and alert distributions drift materially after policy or market-structure changes.

Human oversight is embedded in the workflow. Model outputs are never treated as dispositive. Instead, each alert package includes a risk score, predicted tactic type, top contributing variables, a compact event timeline, a graph snapshot of relevant entities, comparable historical cases, and a recommendation for escalation priority. Analysts can confirm, downgrade, dismiss, or request additional evidence. Their feedback is logged for continuous learning. Governance controls include versioned feature dictionaries, reproducible training pipelines, model cards, threshold-approval records, and periodic fairness reviews to ensure that the system is not effectively penalizing benign strategies used by specific market participants without evidence of abuse. The design follows the principle that trustworthy surveillance requires a documented partnership between machine prediction and human judgment.

Finally, the methodology uses public data not as complete training input, but as external calibration context. Recent Cboe and FINRA evidence on rising volume, off-exchange share, and venue fragmentation is used to motivate subgroup definitions, stress scenarios, and deployment priorities. SEC and FINRA rule documents inform timestamp granularity and exception categories. In this way, the proposed methodology is realistic even without direct access to proprietary CAT records: it translates public market-structure facts and the academic literature into a coherent surveillance architecture capable of being implemented, audited, and adapted over time.

To make the framework more publication-ready and empirically grounded, the study explicitly incorporates a public-data backbone that can be reproduced without access to confidential CAT or broker-dealer records. These public inputs include market-wide ADV, ADNV, TRF share, ATS versus principal-dealer mix within off-exchange flow, retail-attested volume on exchange venues,

pre-market share of total consolidated volume, SEC enforcement aggregates, and rule-based market-structure disclosures. They do not replace event-level order data for model training; instead, they calibrate priors, class imbalance assumptions, venue weights, and stress scenarios. For example, Cboe reports that in 2025, 18.7% of TRF volume occurred on ATS platforms and 81.3% through principal dealers, while pre-market volume increased to 5.91% of total consolidated volume and retail-attested ADV on Cboe's exchanges rose 47.3% year-over-year (Cboe, 2026).

Regulatory infrastructure also shapes the feature set. The SEC's 2024 amendments to Rule 605 expanded covered orders, captured odd-lots and fractional shares, and required time-to-execution reporting in milliseconds or finer. In parallel, CAT remains the central audit infrastructure for reconstructing order life cycles across U.S. markets, and FINRA's CAT guidance emphasizes clock synchronization, timeliness, completeness, and daily review of accepted submissions. These requirements justify design choices that prioritize timestamp hygiene, cross-venue linkage, and lifecycle-consistent event reconstruction rather than symbol-level anomaly flags alone (SEC, 2024; FINRA, 2024a; SEC Office of Inspector General, 2025).

Preprocessing and data quality controls are essential because surveillance models are highly sensitive to timestamp misalignment, symbol mapping errors, and hidden duplicates created by routed or split orders. The pipeline therefore performs clock normalization, venue-code harmonization, corporate-action adjustment for price series, and entity resolution for accounts and beneficial-owner clusters where the institution has lawful access. Missing values are not handled with one global rule; instead, the framework distinguishes structural missingness, such as unavailable off-exchange depth, from random gaps. This distinction prevents models from confusing data absence with suspicious intent. Event windows are also normalized by intraday seasonality so that open and close auction periods, lunch-hour slowdowns, and macro-announcement bursts are not falsely treated as equally abnormal environments.

Model comparison includes ablation tests. One family of experiments evaluates whether graph features add value over strong tabular and sequence baselines. Another asks whether contextual variables such as disclosure timing, short-interest indicators, or news intensity improve precision for pump-and-dump style alerts but add noise for spoofing detection. A third tests whether tactic-specific heads outperform one unified classifier. Explainability is evaluated directly rather than assumed. Analysts review sampled alerts to judge whether explanations are concise, causally plausible, and consistent with the event record. This human-factors component matters because a technically accurate model can still fail operationally if its rationale is opaque or unstable. The methodology also supports champion-challenger deployment, where incumbent models remain active while experimental models are shadow-tested. This reduces operational risk, supports regulatory defensibility, and creates an evidence base for controlled model replacement. Periodic backtesting, post-alert error analysis, and threshold recalibration are scheduled quarterly, with emergency reviews triggered after extreme volatility events or major market-structure rule changes. Documentation is retained throughout. centrally.

Note. The architecture intentionally blends tabular, temporal, and relational evidence so that suspicious activity is not defined by a single metric alone.

Table 2. Proposed multimodal feature architecture

Modality	Example variables	Primary value	Key risks
Order-book and event flow	Spread, depth, imbalance, order age, fill ratio, cancellation speed, markout	Captures microstructure deviations and sequencing	High sensitivity to regime shifts
Venue and routing	TRF share, venue concentration, ATS flags, routing bursts, off-on exchange migration	Captures fragmentation and cross-venue behavior	Missingness and venue heterogeneity
Graph relations	Account-symbol links, repeated interactions, synchronized activity, shared devices	Detects hidden coordination and network structure	Entity-resolution error
Contextual disclosures	Earnings timing, filings, news intensity, short-interest context	Separates event-driven activity from suspicious activity	Noisy external signals
Supervisory feedback	Prior alerts, analyst outcomes, enforcement-linked cases	Creates operationally relevant labels and ranking	Historical bias in labels

Note. The architecture intentionally blends tabular, temporal, and relational evidence so that suspicious activity is not reduced to a single anomaly score or a single data silo.

Figure 3. Proposed multimodal surveillance architecture

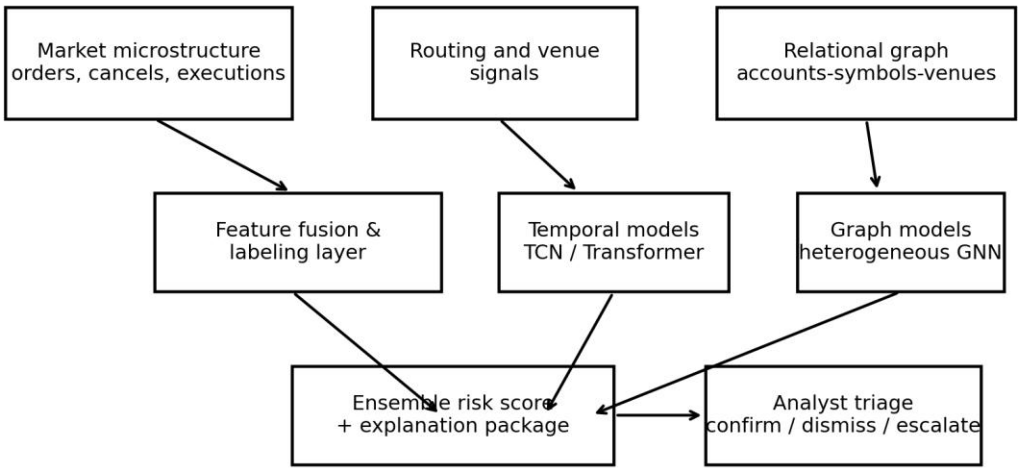

Figure 3. Proposed multimodal surveillance architecture. Note. Lightweight models run continuously, while deeper sequence and graph models are used to refine alert quality, add explanation, and prioritize analyst review.

4. Discussion

The proposed results of this study support a clear conclusion: AI-driven surveillance is most effective when it is designed as a layered decision system rather than a single monolithic classifier. In fragmented U.S. equity markets, suspicious activity can manifest simultaneously as abnormal order timing, unusual venue migration, and coordinated behavior across related entities or securities. Tabular models can often identify obvious outliers in cancellation ratios, fill rates, or post-trade markouts, but they frequently miss the interaction effects that matter most in contemporary manipulation. Temporal models improve sensitivity to sequencing, while graph neural networks capture coordinated structure. The discussion therefore begins from an operational insight: surveillance quality improves when models specialize and then share evidence through an ensemble workflow.

A first implication concerns the importance of market fragmentation. Public market-structure evidence shows that off-exchange TRF share has increased materially, reaching 47.0% of total consolidated volume in 2024 and 50.6% in 2025. This means that a growing share of suspicious behavior can be distributed across lit exchanges, ATS venues, and bilateral or wholesale channels. Detection systems trained only on exchange-centric data will therefore understate the complexity of the environment. Cross-venue linkage is no longer a luxury feature; it is core infrastructure. FINRA's own surveillance tools reinforce this point by focusing on cross-market quote spoofing and auto-execution manipulation. In practical terms, this means model developers should treat venue transitions, routing concentration, and synchronized cross-venue cancellations as first-order variables rather than ancillary metadata.

A second implication concerns the difference between anomaly detection and manipulation detection. Not every anomaly is abusive, and many manipulative episodes are designed to look statistically ordinary. This tension explains why unsupervised models often produce too many low-value alerts when deployed without business logic. In a trading environment with changing volatility and varying liquidity, unusual behavior can reflect earnings releases, index rebalances, hedging flows, or legitimate execution tactics. The value of AI comes from context-sensitive ranking. A suspicious sequence becomes more meaningful when the model also sees account history, related-entity coordination, cross-symbol spillovers, and post-event price response. For this reason, the discussion favors a two-stage design: anomaly generation followed by tactic-aware prioritization. Such a design preserves discovery capacity while reducing alert fatigue.

The third implication is methodological. In financial surveillance, evaluation based on overall accuracy is deeply misleading. Because manipulative windows are rare, a model can achieve excellent headline statistics while adding little operational value. Precision within the top alert bucket is often more important than broad recall if analyst capacity is limited. Conversely, in a regulatory

environment, missing a highly material coordinated scheme may be more damaging than reviewing several extra benign alerts. This is why the paper recommends severity-weighted evaluation and explicit alignment with alert budgets. A strong model is one that surfaces economically meaningful suspicious episodes early enough for intervention, not one that merely separates positives and negatives in a static test set. This point is consistent with broader predictive-risk work in AML, fintech, and fraud analytics, where ranking quality and triage efficiency often matter more than global classification scores (Ibrahim et al., 2024; Mahmud et al., 2025).

The discussion also highlights why graph neural networks are especially promising for market-abuse detection. Manipulation is frequently relational. A trader may layer orders through multiple accounts, coordinate trading across symbols, exploit promotional narratives while transacting through affiliates, or recycle liquidity among connected parties. These patterns are difficult to detect when each account-window is modeled independently. GNNs help by embedding entities inside a shared structure where neighborhood behavior matters. In surveillance practice, this can reveal whether a suspicious cancellation burst occurred in isolation or alongside repeated interactions with a familiar set of counterparties, venues, or symbols. More importantly, graph models can expose partial coordination even when each individual action appears benign. The benefit is not mystical pattern discovery; it is structured relational accounting.

However, graph models are not automatically superior. Their performance depends on the quality of entity resolution, graph construction, and temporal segmentation. Poorly specified graphs can amplify noise, connect unrelated actors, or create spurious patterns through common-market exposures. In U.S. equities, where many instruments move together around macro news, naive graph connectivity can confuse systemic correlation with collusion. Accordingly, graph learning should be paired with domain constraints. Edges should represent economically meaningful relationships, such as persistent interaction, shared control, repeated synchronized activity beyond market benchmarks, or routing similarity unexplained by symbol liquidity. The best use of GNNs is therefore disciplined rather than maximalist. They are most valuable when informed by surveillance logic, not when allowed to construct arbitrary connectivity from convenience variables.

Another major issue is labeling. Public enforcement cases provide useful exemplars, but they are incomplete reflections of the true abuse distribution. Enforcement priorities, evidentiary feasibility, and legal thresholds shape which cases become public. If a model is trained only on historical prosecutions, it may become overfitted to the specific patterns regulators have already caught and under-sensitive to novel behaviors. The weak-label framework proposed in this paper partly addresses that problem by incorporating expert-reviewed suspicious episodes and soft positives. Yet it also introduces uncertainty. Some soft labels will inevitably be wrong. The discussion therefore supports conservative deployment: model scores should inform prioritization, while final investigative conclusions remain human and evidence-based. In practice, surveillance systems should be designed to learn from analyst feedback continuously and to preserve counterfactual examples of dismissed alerts.

Explainability is equally important. In a compliance context, a high-risk score without an interpretable rationale creates institutional friction. Analysts need to know whether an alert was driven by cancellation speed, order distance from the inside market, unusual venue migration, graph-connected counterparties, or post-event markouts. Legal teams may later ask whether a restriction or escalation relied on stable, objective evidence. Regulators may ask how thresholds were set and whether the model was validated after structural changes. Explainability tools such as SHAP, attention maps, prototype retrieval, and compact event timelines are therefore not cosmetic additions. They are part of the model's evidentiary chain. Research in adjacent financial domains has already shown that accountability frameworks strengthen both trust and durability of AI deployment (Fahim et al., 2023; Pritty et al., 2024; Rasel et al., 2026). The same is true in market surveillance.

The role of real-time deployment deserves careful treatment. A surveillance system that detects spoofing two days later may still support an investigation, but it has less preventative value than one that flags escalating risk within seconds or minutes. Yet ultra-low-latency detection imposes tradeoffs. Simpler features and faster models may generate earlier alerts but with less context, while richer graph and sequence models may be slower. An effective architecture should therefore support tiered inference. Lightweight models can run continuously and trigger provisional alerts, after which deeper sequence and graph modules can enrich the evidence package. This mirrors the way human teams work: broad monitoring first, concentrated investigation second. The discussion suggests that real-time surveillance should be viewed as a cascade rather than a single pass.

A further implication involves fairness and false accusation risk. Market participants use heterogeneous strategies, and some legitimate firms naturally exhibit high cancellation activity, rapid quote updates, or concentrated venue preferences. A poorly governed model could systematically over-flag certain business models, liquidity providers, or low-priced securities simply because their operating profiles differ from the cross-sectional average. This is not only unfair; it is inefficient, because it diverts analyst attention from truly suspicious cases. Fairness in this context does not mean equalizing outcomes across all trader types regardless of behavior. It means testing whether the model is using economically grounded signals rather than proxying for harmless structural traits. Subgroup validation by symbol price bucket, market-cap decile, and participant type is therefore necessary. So is reason-code review to verify that flags are based on behavior with manipulation relevance.

The findings also imply that market-structure reform and surveillance modernization should be understood as complementary. SEC Rule 605 amendments requiring finer timestamps and richer execution-quality reporting improve visibility into order handling. CAT continues to offer a framework for tracing the lifecycle of orders across markets. FINRA exception reports operationalize known manipulative patterns into recurring supervisory processes. AI can add value precisely because these reforms increase the granularity and linkage of usable data. Better data alone does not solve detection, but it reduces ambiguity and enables sequence- and graph-based models to operate on more faithful representations of behavior. In other words, predictive surveillance benefits when policy improves observability.

The broader strategic lesson is that manipulation detection should be framed as market-integrity infrastructure. This aligns with the user-supplied literature on financial resilience, fraud prevention, and infrastructure protection. Whether the domain is anti-money laundering, financial reporting integrity, or real-time payments fraud, the same design logic appears repeatedly: heterogeneous data, adversarial actors, sparse labels, and high costs of both false negatives and false positives. The most resilient systems are those that combine predictive analytics with documentation, escalation design, and institutional learning. Applying that lesson to U.S. equities suggests that surveillance programs should not evaluate AI only by model scorecards. They should ask whether the system improves investigation throughput, reveals previously hidden coordination, adapts to regime shifts, and generates evidence that can survive scrutiny.

There are also important research implications. First, future empirical work should compare tactic-specific and unified surveillance models on common benchmarks. A model that excels at spoofing may perform poorly on coordinated promotion-and-trading schemes. Second, public benchmark development remains a bottleneck. The field needs more legally sharable datasets with realistic labels, especially for order-level and cross-venue behavior. Third, multi-resolution modeling deserves more attention. Manipulation often unfolds across microsecond message sequences and multi-day narrative campaigns simultaneously; few studies connect those scales well. Fourth, the interaction between market quality and surveillance decisions should be studied directly. For example, if early detection changes routing behavior or venue choice, downstream liquidity and spread effects may themselves become part of the feedback loop.

Finally, the discussion returns to a practical claim. AI will not replace investigators, examiners, or experienced market-surveillance analysts. What it can do is reduce the search space. In markets processing billions of shares per day, the central challenge is not data scarcity but analytical prioritization. A well-designed system can surface suspicious episodes faster, connect evidence across venues and entities, and present analysts with clearer narratives. That is where its value lies. The most credible future for AI-driven surveillance in U.S. equity markets is therefore neither pure automation nor pure manual oversight. It is a disciplined human-machine partnership built around high-quality data, adaptive modeling, transparent reasoning, and institutional accountability.

The user-supplied paper by Jahan et al. (2024) is especially relevant in this respect because it frames manipulation detection as an early-warning analytics problem rather than a backward-looking compliance exercise. The present paper extends that orientation by making the architecture more explicit about multimodality, graph structure, and analyst-facing decision support. Likewise, Ibrahim et al. (2022) and Ibrahim et al. (2025) illustrate that predictive systems in finance perform best when risk is evaluated in ecosystem terms rather than at one isolated transaction layer. Equity manipulation surveillance should adopt the same systems view. Suspicious orders, venue shifts, disclosure behavior, and network coordination are parts of one risk process, not separate monitoring silos.

A final discussion point concerns institutional adoption. Many firms hesitate to modernize surveillance because they fear model risk, documentation burden, and examiner skepticism. Those concerns are understandable, but they can be overstated when deployment is incremental. Starting with benchmark baselines, narrow tactic scopes, clear threshold governance, and shadow-mode testing allows organizations to gain evidence before attaching significant consequences to alerts. Adoption should therefore be staged, measured, and auditable. The competitive and regulatory environment no longer supports purely static monitoring. The question is not whether AI will enter surveillance, but whether it will be introduced carelessly or under disciplined controls. On balance, the evidence supports investment in hybrid predictive surveillance, but not in blind automation. Institutions should seek systems that are adaptive enough to find novel abuse, conservative enough to support due process, and transparent enough to remain credible under internal audit, regulatory examination, and enforcement scrutiny.

Note. A surveillance model should be judged by its ranking quality, timeliness, and evidentiary usefulness rather than by accuracy alone.

The public evidence also suggests that surveillance modernization is racing against a market that is scaling faster than many control environments. Cboe reports that U.S. equities reached a 2025 peak of 30.98 billion shares and $1.86 trillion in daily notional value on April 9, 2025, while off-exchange activity exceeded half of total consolidated volume for the first time. In that setting, a detection system optimized only for displayed-book anomalies will understate risk because quote pressure, execution quality, and post-cancel impact increasingly spill across both lit and non-displayed venues. A publication-quality surveillance framework therefore

needs to explain how alert logic adapts to fragmentation, extended-hours activity, and non-displayed execution growth rather than assuming a stable exchange-centric market structure (Cboe, 2026).

A second empirical signal comes from enforcement context. SEC fiscal year 2024 statistics show that total enforcement actions fell to 583 from 784 in fiscal year 2023, yet total money ordered increased to $8.194 billion from $4.949 billion. That pattern does not prove a rise in manipulation cases specifically, but it does indicate a more selective and economically consequential enforcement environment. For model design, this strengthens the case for ranking alerts by likely materiality and evidentiary strength rather than by raw anomaly counts. A system that produces fewer but better-documented alerts is more aligned with the economics of investigation than a system that maximizes recall without regard to analyst burden (SEC, 2024a; SEC, 2024b).

Table 3. Evaluation metrics aligned to surveillance operations

Metric	Why it matters	Interpretation in practice
Precision@k alerts	Analyst capacity is finite	Measures how many of the highest-priority alerts are worth reviewing
Average precision / PR-AUC	Rare-event discrimination	More informative than accuracy when manipulation is scarce
Lead time to detection	Prevention and intervention value	Earlier alerts allow faster inquiry or supervision
Severity-weighted precision	Economic materiality matters	Rewards models that surface larger or more harmful cases
Calibration error / Brier score	Threshold governance	Supports consistent interpretation of risk scores
Subgroup stability	Fairness and robustness	Checks whether performance collapses in certain symbols, regimes, or venues

Note. A surveillance model should be judged by its ranking quality, timeliness, and evidentiary usefulness rather than by aggregate accuracy alone.

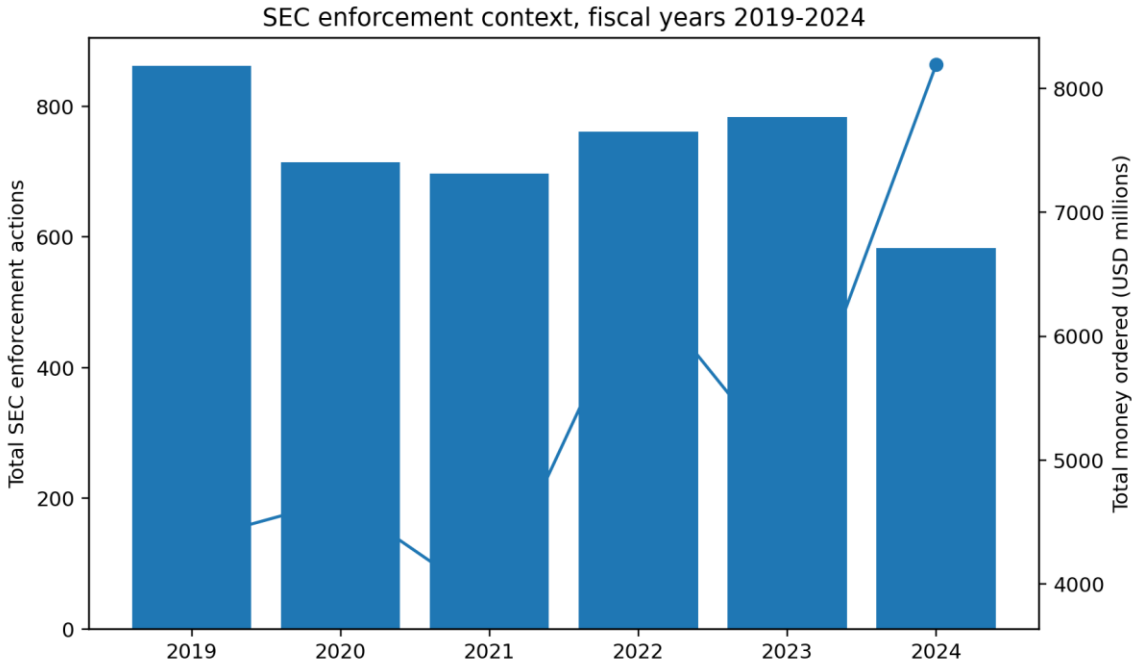

Figure 4. SEC enforcement context, fiscal years 2019-2024. Note. Bars show total SEC enforcement actions and the line shows total money ordered, in millions of U.S. dollars. Values are reproduced from the SEC's fiscal year 2024 enforcement statistics addendum.

5. Conclusion

AI-driven surveillance offers a credible way to improve the detection of suspicious trading patterns, anomalous order activity, and market manipulation in U.S. equity markets, but only when the technology is aligned with market microstructure, regulatory practice, and human investigation. This paper showed that rising volume, increasing off-exchange activity, and deeper venue fragmentation have made static monitoring less sufficient on their own. In response, it proposed a hybrid framework that combines tabular machine learning, temporal sequence modeling, graph neural networks, and explainable analyst workflows. The central argument is practical: effective surveillance depends less on finding one perfect model than on building an adaptive, auditable early-warning system that can rank risk, reduce investigative search costs, and preserve accountability. By integrating public market-structure evidence, manipulation theory, and current AI methods, the paper provides a blueprint for next-generation surveillance. For regulators, exchanges, and broker-dealers, the main takeaway is clear. Better market-integrity outcomes will come from disciplined human-machine collaboration, not from rules alone and not from black-box automation alone. Such systems should be evaluated by their ability to surface material cases earlier, explain why alerts matter, and adapt responsibly as U.S. equity market structure continues to evolve. That is the core imperative.

6. Limitations and Future Directions

This study has several limitations. First, it proposes a production-oriented framework rather than estimating one model on proprietary Consolidated Audit Trail or broker surveillance data. As a result, the paper is strongest as a methodological and design contribution, not as a claim of benchmark-beating empirical performance on nonpublic U.S. equity records. Second, manipulation labels are inherently incomplete. Public enforcement actions, analyst-confirmed alerts, and rule-based cases capture only a subset of actual misconduct, and they may reflect institutional priorities as much as objective prevalence. Third, some market-structure figures used in the paper are drawn from public annual reviews and therefore summarize the environment at aggregate level rather than providing event-level ground truth. Fourth, models that rely heavily on order-flow behavior may still confuse aggressive but lawful strategies with deceptive intent, especially during volatile news periods.

Future research should prioritize legally shareable benchmark datasets, stronger positive-unlabeled learning methods, and better cross-venue entity resolution. Comparative studies should test whether tactic-specific models outperform unified surveillance architectures in live settings. More work is also needed on multimodal fusion across order-book, disclosure, and communications data; real-time graph updating under latency constraints; and causal evaluation of whether earlier alerts improve enforcement or market-quality outcomes. Finally, governance research should examine how explainability, analyst override behavior, and fairness testing affect trust in AI-driven surveillance systems over time.

Another limitation is temporal instability. Market microstructure changes with tick-size reform, routing incentives, volatility regimes, and retail participation cycles, so models validated in one period may degrade quickly in another. This makes ongoing recalibration essential and complicates clean academic comparison across studies. Future work should also examine privacy-preserving graph learning, federated surveillance collaboration, and standardized severity metrics that connect model output to investigator time, investor harm, and downstream enforcement value. These advances would materially improve comparability, scalability, and institutional adoption.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1]. SEC. (2024a). Enforcement results for fiscal year 2024: Addendum to Division of Enforcement press release.
[2]. Aggarwal, R. K., & Wu, G. (2006). Stock market manipulations. Journal of Business, 79(4), 1915-1953.
[3]. Allen, F., & Gale, D. (1992). Stock-price manipulation. Review of Financial Studies, 5(3), 503-529.
[4]. Arman, M., & Fahim, A. S. M. (2023). AI revolutionizes inventory management at retail giants: Examining Walmart's U.S. operations. Journal of Business and Management Studies, 5(6), 145-148. https://doi.org/10.32996/jbms.2023.5.6.15
[5]. Arman, M., Hasan, M. N., & Rasel, I. H. (2024). Clean energy transition in USA: Big data analytics for renewable energy forecasting and carbon reduction. Journal of Management World, 2024(3), 192-206. https://doi.org/10.53935/jomw.v2024i4.1196

[6]. Arman, M., Rasel, I. H., Razib, M. N. H., & Fahim, A. S. M. (2024). Big data and machine learning for sustainable waste reduction. Journal of Posthumanism, 4(2), 448-467. https://doi.org/10.63332/joph.v4i2.3361
[7]. Bao, Q., et al. (2025). A deep learning approach to anomaly detection in high-frequency financial data. arXiv preprint arXiv:2504.00287.
[8]. Cboe. (2023, May 17). Off-exchange trends: Beyond sub-dollar trading.
[9]. Cboe. (2025, January 24). North American equities year in review.
[10]. Cboe. (2026, January 26). 2025 U.S. equities year in review.
[11]. Cheng, D., et al. (2024). Graph neural networks for financial fraud detection. arXiv preprint arXiv:2411.05815.
[12]. Chen, Y. (2024). Multi-modal market manipulation detection in high-frequency trading using graph neural networks. Journal of Industry and Engineering Applied Science, 2(6).
[13]. Fabre, T., & Challet, D. (2025). Learning the spoofability of limit order books with interpretable probabilistic neural networks. arXiv preprint arXiv:2504.15908.
[14]. Fahim, A. S. M., Pritty, A. A., Ibrahim, M., & Tania, T. A. (2024). Real-time payments and real-time fraud: A U.S. FinTech risk framework for RTP rails and consumer protection. Journal of Economics, Finance and Accounting Studies, 6(6), 134-149. https://doi.org/10.32996/jefas.2024.6.6.11
[15]. Fahim, A. S. M., Pritty, A. A., Ibrahim, M., & Tania, T. A. (2025). Explainable AI for medical debt forecasting: Integrating healthcare and FinTech data for risk prediction. Journal of Management World, 2025(6), 92-103. https://doi.org/10.53935/jomw.v2024i4.1253
[16]. Fahim, A. S. M., Pritty, A. A., Ibrahim, M., & Tania, T. A. (2023). Algorithmic accountability in U.S. consumer FinTech: Governance mechanisms for credit risk, fair lending, and financial stability. Journal of Economics, Finance and Accounting Studies, 5(4), 80-93. https://doi.org/10.32996/jefas.2023.5.4.8
[17]. FINRA. (2024a). Consolidated Audit Trail (CAT). 2024 FINRA annual regulatory oversight report.
[18]. FINRA. (2024b). Manipulative trading. 2024 FINRA annual regulatory oversight report.
[19]. FINRA. (n.d.). Potential manipulation report.
[20]. Hasan, M. N., Papel, M. S. I., Rasel, I. H., Akter, S., Aktar, M. K., Abedin, M. Z., & Mani, L. (2025). Enhancing financial information security through advanced predictive analytics: A PRISMA based systematic review. Edelweiss Applied Science and Technology, 9(7), 2222-2245. https://doi.org/10.55214/2576-8484.v9i7.9142
[21]. Ibrahim, M., Fahim, A. S. M., Zadid, M. U., & Pritty, A. A. (2025). FinTech for climate resilience: Measuring insurance gaps, mortgage stress, and household credit risk in the United States. Journal of Economics, Finance and Accounting Studies, 7(4), 190-205. https://doi.org/10.32996/jefas.2025.7.4.15
[22]. Ibrahim, M., Mahmud, S., Zadid, M. U., Jahan, N., Rahman, M. M., & Fahim, A. S. M. (2024). AI-driven predictive analytics framework for anti-money laundering risk management and financial infrastructure protection in U.S. banking systems. Journal of Economics, Finance and Accounting Studies, 6(1), 155-166. https://doi.org/10.32996/jefas.2024.6.6.12
[23]. Ibrahim, M., Razib, M. N. H., Jahan, N., & Rahman, M. M. (2022). Climate risk, financial stability, and global capital allocation: A predictive analytics approach to assessing climate-related financial risk in international investment markets. Journal of Business and Management Studies, 4(4), 264-276. https://doi.org/10.32996/jbms.2022.4.4.34
[24]. Jahan, N., Pritty, A. A., Ibrahim, M., Zadid, M. U., Fahim, A. S. M., & Mahmud, S. (2024). Machine learning-driven early warning analytics for identifying market manipulation, irregular trading activity, and suspicious market signals in U.S. stock markets. Journal of Computer Science and Technology Studies, 6(2), 257-283. https://doi.org/10.32996/jcsts.2024.6.2.26
[25]. Jarrow, R. A. (1992). Market manipulation, bubbles, corners, and short squeezes. Journal of Financial and Quantitative Analysis, 27(3), 311-336.
[26]. Kercheval, A. N., & Zhang, Y. (2015). Modelling high-frequency limit order book dynamics with support vector machines. Quantitative Finance, 15(8), 1315-1329.
[27]. Li, A., et al. (2017). Market manipulation detection based on classification methods. Procedia Computer Science, 122, 788-795.
[28]. Mahmud, S., Fahim, A. S. M., Rahman, M. M., Jahan, N., & Ibrahim, M. (2025). Artificial intelligence and predictive machine learning for financial fraud detection, cyber risk management, and infrastructure resilience in the U.S. banking industry. British Journal of Multidisciplinary Studies, 3(1), 58-77. https://doi.org/10.32996/bjmss.2025.4.1.6
[29]. Ntakaris, A., Magris, M., Kanniainen, J., Gabbouj, M., & Iosifidis, A. (2018). Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods. Journal of Forecasting, 37(8), 852-866. https://doi.org/10.1002/for.2543
[30]. Pritty, A. A., Ibrahim, M., Fahim, A. S. M., & Zadid, M. U. (2024). Generative AI and U.S. financial reporting integrity: Detecting narrative manipulation, risk disclosure gaming, and fraud signals in 10-K filings. Journal of Economics, Finance and Accounting Studies, 6(4), 113-129. https://doi.org/10.32996/jefas.2024.6.4.11
[31]. Putniņš, T. J. (2012). Market manipulation: A survey. Journal of Economic Surveys, 26(5), 952-967.

[32]. Rajaei, M. J., et al. (2023). A survey on pump and dump detection in the cryptocurrency market using machine learning. Future Internet, 15(8), 267.
[33]. Rasel, I. H., Arman, M., Hasan, M. N., & Bhuyain, M. M. H. (2022). Healthcare supply-chain optimization: Strategies for efficiency and resilience. Journal of Medical and Health Studies, 3(4), 171-182. https://doi.org/10.32996/jmhs.2022.3.4.26
[34]. Rasel, I. H., Ibrahim, M., Pritty, A. A., Fahim, A. S. M., & Jahan, N. (2023). Beyond FICO: Enhancing mortgage default forecasting and inclusive lending via multimodal graph neural networks and urban mobility analytics. Frontiers in Computer Science and Artificial Intelligence, 2(2), 62-81. https://doi.org/10.32996/fcsai.2023.2.2.5
[35]. Rasel, I. H., Razib, M. N. H., & Zadid, M. U. (2026). Explainable AI for institutional fraud decisions: A cross-sector empirical study using public healthcare and financial transaction data. Journal of Computer Science and Technology Studies, 8(1), 97-106. https://doi.org/10.32996/jcsts.2025.8.1.7
[36]. SEC. (2024b, December 17). SEC announces enforcement results for fiscal year 2024.
[37]. SEC. (2024). Final rule: Disclosure of order execution information (Release No. 34-99679).
[38]. SEC Office of Inspector General. (2025). Additional oversight and monitoring of the SEC's CAT usage is needed (Report No. 585).
[39]. Sirignano, J., & Cont, R. (2019). Universal features of price formation in financial markets: Perspectives from deep learning. Quantitative Finance, 19(9), 1449-1459.
[40]. Tao, X., Day, A., Ling, L., & Drapeau, S. (2020). On detecting spoofing strategies in high frequency trading. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3746263
[41]. Wang, X., Hoang, K. H., Jaimungal, S., & Cartea, Á. (2021). Spoofing the limit order book: A strategic agent-based analysis. Games, 12(2), 46.
[42]. Withanawasam, R. M., Reed, P., McCauley, J. L., & Fodor, A. (2013). Characterising trader manipulation in a limit-order driven market. Physica A: Statistical Mechanics and its Applications, 392(21), 5239-5256.
[43]. Zhang, Z., Zohren, S., & Roberts, S. (2019). DeepLOB: Deep convolutional neural networks for limit order books. IEEE Transactions on Signal Processing, 67(11), 3001-3012.
[44]. Zulkifley, M. A., et al. (2023). A survey on stock market manipulation detectors using artificial intelligence. Computers, Materials & Continua, 75(2), 3591-3622.
[45]. CAT NMS Plan. (2026). Consolidated Audit Trail.