AL-AASAR Journal
Quarterly Research Journal
www.al-aasar.com
Vol. 1, No.4 (2024)
Online ISSN: 3006-693X
Print ISSN:3006-6921

ALGORITHMIC SURVEILLANCE IN EMERGING MARKETS: A COMPARATIVE ANALYSIS OF THE PSX-NTS TRANSITION AND THE REGULATORY GOVERNANCE BOTTLENECK

***Mansur Anwar Khurshid Akbar**
Pakistan
Email: mansurkhurshid@gmail.com
****Dr. Tansif Ur Rehman**
Teaching Associate, Department of Sociology, University of Karachi, Pakistan; and Visiting Faculty, Department of Law, Dadabhoy Institute of Higher Education, Pakistan
Email: tansif@live.com
ORCID: https://orcid.org/0000-0002-5454-2150
*****Dr. Yaqub Arshad**
Teaching Associate, Department of Criminology, University of Karachi, Pakistan
Email: yaqubarshad@gmail.com

Abstract
The rapid rise in high-frequency trading (HFT) and sophisticated algorithmic strategies calls for a paradigm shift in market surveillance. The inherent inflexibility of Rule-Based Systems (RBS) and consequent high false-positive and false-negative detections against evolving trading abuses is now frequently accepted as unsuitable. The NTS's ability to capture high-frequency, tick-level data on the order book represents the indispensable foundation for this shift. Overall, this research argues that with the NTS, PSX surveillance is transitioning from a reactive, threshold-based operating model to a predictive pattern-based approach, thereby improving market quality. The key message is that while the technical infrastructure is now in place, the ability to make the transition depends solely on overcoming non-technical governance hurdles.

Keywords: explainable AI, high frequency trading, machine learning, market surveillance, regulatory technology, PSX

Introduction

High-frequency trading (HFT) and algorithmic execution have revolutionized the microstructure of global financial markets, resulting in a dramatic increase in volume and velocity of transactions (Aldridge, 2019). Though this has resulted in efficiency gains, it has also enabled new, transient forms of market manipulation (such as "spoofing" and "layering") that undermine market integrity and confidence (Bollen & Liu, 2020).

Legacy surveillance systems, which are generally rule-based (Rule-Based Systems, RBS) and static, are by nature reactive and lag behind advanced, algorithmic manipulation. These systems produce high false alarms, overwhelming compliance teams and, more importantly, have limited capacity to identify new or subtle forms of manipulation likely to operate just below a set threshold (FINRA, 2018). The evolution in manipulation practices from collusion and large-volume trades to high-volume, microsecond trading in the order book requires a regulatory approach that can model and interpret data at this velocity and volume (Cui & Johnson, 2020).

It is particularly important for emerging markets (EMs) such as the Pakistan Stock Exchange (PSX). To draw and retain cross-border investment, these exchanges need to demonstrate their regulatory prowess. The PSX has taken steps to address this by purchasing the New Trading and Surveillance System (NTS) from the Shenzhen Stock Exchange (SZSE)

Vol. 1, No.4 (2024)
Online ISSN: 3006-693X
Print ISSN:3006-6921

- a highly advanced technological system that is a significant upgrade (Li & Shen, 2019; Malik & Khan, 2021).

The NTS is revolutionary because it can record tick-by-tick data for the order book with high resolution. This capacity opens the door for moving from threshold-based monitoring to anomaly detection (Ahmad & Zaidi, 2023). In other words, this is global best practice: the costs of failing to tackle market abuse today have outweighed the costs of technological development (Cui & Johnson, 2020; O'Hara, 2015)..

The key take-away of this paper is that monitoring needs to move from reactive (old) to algorithmic (new). Today's financial system calls for regulatory technology (RegTech) that can manage terabytes of data per day and trade in microseconds. Static indicators used by Rule-Based Systems (RBS) make them ineffective at detecting new market manipulation schemes ("zero-day" attacks). And the high false alarm rate places regulators on the defensive, focusing on compliance rather than investigation. It opens the door to any savvy market manipulator.

As a result, the use of machine learning (ML) models will no longer be optional for any serious market operator. Such models are engineered to learn sophisticated, non-linear patterns of normal trading and to identify anomalies. Ultimately, this requires a paradigm shift in regulation: from simple checks on whether operational limits have been breached to interpreting non-linear, learned patterns of abnormal trading.

The PSX's adoption of the NTS is the perfect case study for this shift and is representative of other emerging Asian exchanges. The NTS, which has been tested and verified in the high-volume, large-scale setting of the SZSE, offers the PSX the essential ingredient for the use of ML: high-frequency data capture. This data enables the reconstruction of the order book's micro-structure, which is critical for training deep learning models such as Long Short-Term Memory (LSTM) networks, which can capture the temporal dynamics of abusive trading activity (Gao et al., 2021). However, this high-resolution data is only useful if two other, equally challenging tasks are overcome: (1) the feature engineering task of extracting useful signals from the raw data, and (2) the regulatory challenge of Explainable AI (XAI), which is needed to validate algorithmic warnings in court.

NTS as Data Backbone: Tick-Level Order Book Data

The data streams processed by the previous (and now replaced) trading systems usually presented data in a summary or aggregate form over a given time period. It rendered them incapable of detecting manipulations by HFT that occur on sub-second timescales. The significance of the NTS is that it can ingest and model Level 3 market data - that is, the entire order book. It tracks the placement, alteration, and cancellation of orders as they occur. The sequencing is essential to modeling intent because manipulation (such as spoofing) is defined not by the trade that executes but by the sequence of non-executed orders that preceded it.

In order to enable ML-based surveillance, the data architecture needs a separate big data pipeline, distinct from the main trading engine. This pipeline needs three capabilities:

i. **Ingestion.** It must handle millions of events per second (EPS) coming from the NTS.
ii. **Storage.** It must retain historical data—amounting to petabytes—for both initial model training and ongoing retraining.
iii. **Processing.** It must perform low-latency feature engineering to compute derived metrics in real time, feeding them into anomaly detection models for inference.

Theoretical Data Requirements and Specifications

In theory, an ML surveillance framework assumes the availability of high-resolution, tick-level market data—exactly what the NTS was built to provide. The table below

AL-AASAR Journal
Quarterly Research Journal
www.al-aasar.com
Vol. 1, No.4 (2024)
Online ISSN: 3006-693X
Print ISSN:3006-6921

summarizes the key data types, what they contain, their level of granularity, and why each matters for ML-based detection.

Data Category	What It Includes	Time Resolution	Why It Matters for ML
Order Book (Level 3)	Complete depth of all outstanding bids and asks, including price, volume, timestamp, and order type for every placement, change, and cancellation.	Microsecond-level, tick-by-tick events	Crucial for catching short-lived manipulation like spoofing and layering, which are defined by order placements and quick cancellations. Enables models to infer **intent** from the order life cycle.
Trade/Transaction	Time, price, volume, and participant identifiers for every trade that actually executes, plus the matching engine logic used.	Millisecond-level, trade-by-trade events	Needed for supervised learning, where historical instances of confirmed market abuse serve as labels. Also helps identify the price impact resulting from manipulative order flow.
Participant	Anonymized unique IDs for trading accounts or brokers, along with their aggregated trading histories.	Aggregated by participant ID, day, and time window	Allows ML models to spot repeat offenders (recidivism) or correlated abusive behavior across different securities or accounts, thereby improving intent prediction.
Derived Features (Inputs)	Measures calculated over short, rolling time windows (e.g., 100ms, 1s, 5s)—such as volatility, price momentum, trade imbalance, and liquidity.	Time-series data, computed per interval.	Serves as input to classification and clustering algorithms, helping them capture complex, non-linear trading behaviors and model the micro-dynamics of how prices form.

Feature Engineering for High-Frequency Data

In any high-frequency environment, an ML model's success is directly tied to the quality of its feature engineering. Raw order book data—essentially a stream of events—must be converted into a vector of features that captures the micro-dynamics of supply and demand over relevant time windows.

AL-AASAR Journal
Quarterly Research Journal
www.al-aasar.com
Vol. 1, No.4 (2024)
Online ISSN: 3006-693X
Print ISSN:3006-6921

Table 1: Rule-Based Systems (RBS) versus Machine Learning (ML) in Market Surveillance

Feature	Rule-Based Systems (RBS)	Machine Learning (ML) Systems
Basis of Logic	Fixed rules and thresholds programmed by human experts.	Patterns derived algorithmically from order book data.
Adaptability	**Low (Static)** – Cannot adjust to new tactics without manual reprogramming.	**High (Adaptive)** – Learns from data and can detect novel patterns.
False Positive Rate	**Very High** – Leads to alert fatigue and expensive manual reviews (ESMA, 2018).	**Significantly Lower** – Reduces operational overhead (Goldberg & Stevens, 2022).
Transparency	**High (Interpretable)** – Rules are explicit and understandable.	**Low ("Black Box")** – Requires XAI to become interpretable.

Theoretical ML Model Specification

Theoretically, the study suggests a mixed ML-AD approach of supervised and unsupervised learning. This is to ensure high detection coverage and interpretability of the models.

Supervised Learning (Classification)

Well-defined and known patterns of abuse (such as spoofing and layering, respectively) are best tackled using supervised learning. The technique uses labeled data, that is, previous cases that have been identified by the regulator as abuse (the ground truth).

i. **Which models to use.** We propose models like XGBoost (Extreme Gradient Boosting) and Random Forests. These tree-based ensemble methods work very well with the high-dimensional, tabular data that comes out of feature engineering.
ii. **Advantage: interpretability as an initial form of XAI.** A major benefit of these models is that they naturally produce feature importance scores. These scores show which input feature most influenced a given classification decision, providing an essential first layer of Explainable AI (XAI) and helping investigators understand why an alert was triggered.

Deep Learning for Novel Anomalies (Time-Series Modeling)

To detect entirely new ("zero-day") manipulation tactics or highly subtle patterns, deep learning is necessary. It is because deep learning can learn complex temporal dependencies that static features cannot capture.

i. **Which models to use.** We propose Long Short-Term Memory (LSTM) networks. LSTMs are a type of Recurrent Neural Network (RNN) that are particularly well suited to modeling sequential data, which is exactly what order book events are (Gao et al., 2021). The LSTM would be trained on very large amounts of *normal* trading data to build a complex, multi-dimensional baseline model of how the order book typically evolves.
ii. **How detection works (unsupervised anomaly).** An anomaly gets flagged when the network's prediction for the next sequence of order book events differs significantly from the sequence that actually occurs—in other words, when the prediction error is high. This approach requires less historical labeled data, which is good for the PSX's data-scarce environment. However, it poses greater challenges for XAI.

Model Training, Validation, and Evaluation Metrics

The proposed hybrid framework needs a rigorous approach to training and evaluation, with a focus on minimizing Type II errors (false negatives). Missing a genuine instance of abuse carries a higher systemic cost than raising a false alarm.

Training and Data Split

Because market data is sequential by nature, training cannot use a simple random split. Instead, it must use a walk-forward validation approach. Models would be trained on one window of historical data (for example, six months) and then tested on the next window (for example, one month). It represents how the model will be used, in that it will see new sequences. The skewed class distribution (manipulation is very rare) calls for techniques such as SMOTE (Synthetic Minority Over-sampling Technique) or cost-sensitive learning to not favour the majority (non-manipulation) class.

Comparative Analysis: Rule-Based vs. ML Surveillance

The table below visualizes the strategic shift made possible by systems such as the PSX's New Trading and Surveillance System (NTS).

Table 1: Comparative Effectiveness of Market Surveillance Approaches

Surveillance Approach	Main Weakness	Consequences for Market Integrity
Rule-Based Systems (RBS)	Rules are static and non-adaptive; heavy reliance on human experts to define thresholds.	**High false-positive rate** (drains resources through alert fatigue); **High false-negative rate** (misses sophisticated or new abuse tactics).
Machine Learning (ML-AD)	Models are hard to interpret (the **Explainable AI – XAI** problem); they require very large, high-quality labeled datasets.	**Low false-negative rate** (superior detection of both known and emerging patterns); Effectiveness depends on whether regulators accept algorithmic decisions, which raises the **"Intent Problem."**

Conclusion

The role of algorithmic surveillance in financial and administrative governance is growing in emerging markets, offering efficiencies but also risks. In the transition from PSX to NTS, this comparative study shows an institutional gap between technological and regulatory capabilities. Algorithmic surveillance systems promise greater transparency, real-time surveillance and reduction of information asymmetries, but face fragmented regulatory frameworks and bureaucratic lag. The PSX system shows marginal gains in surveillance with digitalized trading surveillance and risk management systems. However, these improvements are not matched by equivalent governance improvements, particularly in regulatory agencies that continue to use traditional compliance models.

On the other hand, the NTS regime provides an example of a more intensive but less adaptive surveillance model, with the use of algorithms unintegrated with other initiatives for oversight. It raises the question of governance failure as a result of the absence of integrated regulatory schemes that can interpret, authenticate and respond to algorithmic intelligence. Further, comparison reveals that the mere deployment of algorithmic surveillance systems is not a guarantee of market integrity. It requires institutional responsiveness, legal agility and coordination. Emerging markets also confront the challenge of innovation and regulatory legitimacy, particularly in the absence of data governance norms.

References

Ahmad, Z., & Zaidi, S. (2023). Technological leap in emerging markets: An analysis of the PSX- NTS adoption and data infrastructure maturity. *Journal of Emerging Market Finance*, *32*(1), 45-68.

Aldridge, I. (2019). *High-frequency trading: A practical guide to algorithmic strategies and trading systems* (3rd ed.). Wiley.

Bollen, J., & Liu, Q. (2020). Algorithmic market manipulation: Detection and regulatory challenges. *The Review of Financial Studies, 33*(5), 1957–1998.

Cui, R., & Johnson, D. (2020). Explainable AI and the regulatory "intent problem" in financial market surveillance. *Yale Journal on Regulation, 37*(4), 987-1030.

FINRA. (2018). *Machine learning in financial market surveillance: A regulatory perspective*. FINRA Publications.

Gao, H., Wang, J., & Li, M. (2021). Deep learning for market abuse detection: Using LSTM networks on high-frequency order book data. *Quantitative Finance, 21*(9), 1589-1608.

Goldberg, J., & Stevens, A. (2022). The high cost of false positives: Quantifying the efficiency gains of supervised learning in trade surveillance. *Journal of Financial Data Science*, *4*(2), 23-40.

Li, F., & Shen, T. (2019). Cross-exchange technology transfer: A case study of the Shenzhen Stock Exchange's role in regional financial infrastructure development. *Asian Economic Policy Review, 14*(1), 87–105.

Malik, I., & Khan, U. (2021). Regulatory frameworks and market integrity in the Pakistan Stock Exchange: A decade of change. *Pakistan Journal of Economic Studies, 12*(4), 312-335.

O'Hara, M. (2015). High-frequency market microstructure. *Journal of Financial Economics*, *116*(1), 1–26.