

Event-Driven Financial Crime Intelligence

Ethan Williams
Asst Professor

Abstract : Existing financial crime detection systems largely depend on unsupervised anomaly or novelty detection techniques grounded in data science. A central challenge is ensuring these models remain sensitive and informative as data distributions evolve over time. Performance degradation can go undetected, and models often require supervised retraining using ground-truth labels drawn from confirmed financial crime cases—events that are inherently rare and non-recurring. To address this gap, this study proposes an event-driven compliance intelligence framework that balances detection capacity with risk governance requirements. The framework is demonstrated through a practical implementation focused on Anti-Money Laundering (AML) risk detection and management for a digital wallet operator.

Recent advances point toward a promising direction: noise-filtering methods that support adaptive financial crime detection systems capable of generating insight-driven predictions to inform existing governance controls across a financial crime regime. In particular, this involves aligning known money-laundering typologies with the risk detection and management demands of an AML regime, consistent with the framework envisioned by regulatory bodies such as the Financial Action Task Force (FATF) through its 40 Recommendations. These developments broaden the conceptual link between signals, events, and monitoring through the lens of compliance intelligence, thereby strengthening regulatory science and regulatory-compliance theories of detection and governance.

Keywords: Event-Driven, Compliance Intelligence, Adaptive Detection, Governed Risk Analytics, Regulatory Science, Architectural Layers, Signal Extraction.

1. Introduction

Implementing rules and controls designed to deter financial crime is now a fundamental part of the corporate governance frameworks of the majority of large businesses and many small and medium enterprises. Enterprises also increasingly recognize the importance of understanding compliance and reputational risk exposure for their overall corporate governance. Expectations for a supportive and reliable KYC regime—along with proven processes for effective remediation of financial crime-related risks—are core to effective organizational reputation protection. Authority compliance officers expect institutions to be equipped with tools capable of rapid detection of noncompliance and operation outside prescribed risk parameters. The consequences of business failure, record-high penalties for proven noncompliance, and further limitations to company licenses and charters strongly support these expectations.

The novel framework enhances compliance intelligence by integrating regulatory

obligations, compliance risk monitoring signals, and broader external exchange and government intelligence flows into a single control process. Specifically, the work develops a proof of concept for an Event-Driven Compliance Intelligence Framework for Adaptive Financial Crime Detection and Governed Risk Analytics, providing an adaptive compliance intelligence engine capable of utilizing business-as-usual processes and supporting expert and corporate governance decisions. The approach enables compliance and governance operations to evolve with changing business demands and tackle new risk threats and priorities over time.

Mathematical Formulas:

$$L_i = t_i^{report} - t_i^{event}$$

Eq. (1) — Signal latency — delay between event occurrence and its report.

$$Lagg_f = \frac{1}{N} \sum_{i=1}^N (t_i^{report} - t_i^{event})$$

Eq. (2) — Flow-level lag (Lagg) — average reporting delay across N signals in a data flow.

$$R(s) = w_1 S_{ext}(s) + w_2 S_{int}(s) + w_3 D(s)$$

Eq. (3) — Composite risk score fusing external signal, internal signal, and drift-indicator components.

$$\text{Decision}(s) = \text{Enrich if } R(s) \geq \tau, \text{ else Merge/Discard}$$

Eq. (4) — Significance-threshold decision rule for routing a signal to enrichment vs. merge/discard.

$$PSI = \sum_{k=1}^K (P_k - Q_k) \ln \left(\frac{P_k}{Q_k} \right)$$

Eq. (5) — Population Stability Index (PSI) as a drift-detection statistic across K score/feature bins.

$$\theta_{t+1} = \theta_t + \eta(y_t - \hat{y}_t)$$

Eq. (6) — Online learning update rule for adaptive model parameters θ under closed-loop feedback.

A. Adaptive Detection and Learning

Adaptive mechanisms allow detection systems to respond promptly to changes in monitored environments, enhancing their ability to identify malicious events and activities. Three processes—online machine learning for classification, closed-loop feedback for detection, and drift management—facilitate adaptive detection. Validating the detection performance of adaptive models is nontrivial, as conventional metrics like true positive rate and F1-score generally assume stationary data-generating environments. The regulatory, ethical, and risk-management aspects governing data usage significantly impact adaptive detection performance. These elements shape the prospect of success for application scenarios crucial for a compliance-regulated organization, such as detecting financial crimes. Accordingly, the adoption of operationalizable detection objectives, the principles of study and data provenance, and the use of detection performance metrics are influenced by regulatory requirements, risk indicators, and compliance-intelligence workflows.

Adaptive detection strategies for financial crime demand a natural connection to government and risk analytics processes, which in practical term result from experience. Regulatory science emphasizes the need to separate noise from signal and mandates that noise may be treated in a business-as-usual fashion. In fact, noise management ensures that models recognize the lack of information and spontaneity. KYC, AML, and sanctions regimes determine the core structures and information sources that govern and control business activity. These aspects are formalized next with a mapping of risk controls to regulatory science pillars. Once identified, compliance intelligence allows a broader discussion of its relevance and underlying theory.




Figure A — Detection F1-score for a static vs. an adaptively retrained model under drift (synthetic illustrative data).





Figure B — Conceptual timeline contrasting single-signal latency with flow-level lag (synthetic illustrative data).

Note: Figures A and B use synthetic illustrative data to visualize the conceptual mechanisms described in the paper; they are not derived from empirical results reported in the manuscript.



2. Theoretical Foundations

Building on the concepts outlined in the preceding section, it is appropriate to explore the theoretical constructs that underpin the proposed framework. Financial criminals exploit the dynamic nature of markets, yet conventional models of criminal behaviour assume a constant distribution of signals across spatio-temporal contexts. Modelling frameworks that reflect criminals' motives, their use of risk indicators, and the nature of regulatory science can support approaches that detect the presence of such signals and are capable of adapting to changes in evolutionary behaviour. The following subsections consider these fundamental concepts and formally express them where appropriate.


The problem of adaptive financial crime detection is coherent with the concepts of Regulatory Science and Compliance Theory. Regulatory expectations require the detection of sanctions evasion, anti-money laundering failures, terrorist financing, fraud detection, and other financial crime regimes. These are typically detected ex-post using large and complex data sets. Financial criminals exploit the dynamic nature of markets and may use multiple indicators of risk and fraud. Detection frameworks that focus on change detection, alarm and event detection, and sensor placement learn from the environment and respond to change are capable of reacting to such signals more easily. The key objective of such frameworks is to detect drift and change both in the environment. In such a situation standard adaptive models break down. Online learning, through the use of feedback, explicit handling of drift, and the definition of appropriate loss functions become essential.

A. Compliance Intelligence and Regulatory Science

Integrating regulatory requirements, risk indicators, and intelligence generation results in a new contribution to the theory of compliance intelligence. Risk indicators vary and evolve over time, and this creates a window of opportunity for financial crime detection. However, such information is rarely reflected in risk management activities. The proposed alignment with traditional KYC, AML, and sanctions regimes generates a more comprehensive model that not only supports decision-making but embodies it. Key requirements from regulatory science allow presenting compliance intelligence as a structured taxonomy that captures the various lines of thought until now spread out through the literature. The alignment laid out contribute to define a path toward even clearer and explicit traditional models of compliance intelligence.

Table 1. Architectural Layers of the Event-Driven Compliance Intelligence Framework.

Layer	Primary Function	Representative Components
Data	Ingests external and internal signal sources	AML/KYC/sanctions feeds, VP & control-unit signals
Detection Engine	Extracts, filters, and enriches signals	Signal extraction, enrichment pipeline, model management
Processing	Validates data quality and timeliness	Drift monitoring, bias inspection, retraining triggers



Layer	Primary Function	Representative Components
Governance	Oversees decision execution and compliance	Contextual intelligence, protocol assurance
User Interaction	Surfaces decisions to compliance officers	Alert dashboards, case workflows
Application	Operationalizes outputs into business processes	Case management, regulatory reporting

Table 2. Signal Source Classification.

Source Type	Origin	Example
External	Published by regulatory/government authorities	Police alert, sanctions list update
Internal	Generated within the institution	AML incident report, fraud detection warning
Composite (fused)	Cross-validated across external + internal streams	Delta cross-validated counterparty update


3. Framework Overview

The Event-Driven Compliance Intelligence Framework encompasses a layered architecture for the governance and orchestration of elements required for adaptive financial crime detection and risk reporting systems. It lays the foundation for the development of a compliance intelligence centre that connects regulatory science, innovation in governed risk analytics, and event-driven detection engines. The architecture defines, complements, and completes both the detection and governance aspects of a single platform.

Research criteria for the design, construction, and utility of an event-driven compliance intelligence framework is proposed—namely proof-of-concept or implementation within either a single organization or exercises between multiple alignment partners in response to an intelligence request from Government authorities.

A. Architectural Layers

The Event-Driven Compliance Intelligence Framework for Adaptive Financial Crime Detection and Governed Risk Analytics comprises an architectural blueprint with six logical layers: data, detection engine, processing, governance, user interaction, and application. The data layer ingests external signal sources for anti-money laundering, know-your-customer, and sanctions regimes, together with internal bank signals from vice president and controlling unit decisions. The detection engine contains a comprehensive suite of components and modules that capture and process every data signal, feed the adaptive-detection engine, and manage the supervised-model assets for financial-crime detection.



The processing layer contains logic to ensure and validate data quality and timeliness, particularly for compliance with operational governance and regulatory requirements. Special focus is given to the detection of adaptive drift not only through detection-engine performance monitoring and retraining, but also through inspecting bias in KYC and sanction signal sourcing. A governance layer oversees and controls the proper execution and completion of decision recommendations, providing further contextual intelligence at different administrative levels, and ensuring that governance protocols for the financial institutions are satisfied.

Table 3. Enrichment Dimensions Applied to Extracted Signals.

Dimension	Direction	Purpose
Forward engineering	Prospective	Creates auxiliary data items enabling model response and detection
Backward engineering	Retrospective	Encodes empirical knowledge of prior model response behaviour
Sideways engineering	Lateral	Extracts hidden factors and associates them with data type


Table 4. Signal Timeliness Indicators.

Indicator	Definition	Applies to
Latency	Delay from event occurrence to its being reported	A single signal
Lagg	Latency generalized to a specific data flow	An aggregated flow of signals

4. Detection Engine Components

The detection engine consists of several modules, each using event streams and risk indicators from the architectural data layer. As shown in Figure 16, raw input is examined by signal extraction methods that identify expected patterns within event footprints. Based on the scope of regulatory risk, patterns below a certain significance level are either merged with similar footprints or discarded. The remaining patterns are passed to an enrichment process that offers further feature engineering and risk scoring, preparing them for model management and evaluation components.

The detection engine seeks to absorb and learn from a wide variety of signals to identify non-standard events or radar-like scans across the established risk thresholds. Novel, previously unobserved signal patterns are communicated to the model management processes and either associated with existing predictive models or examined for separate learning objectives. Detection capacity and performance are constantly monitored, and challenges such as data drift or deteriorating performance are automatically detected.



A. Signal Extraction and Enrichment


Enabling detection of financial crimes involves analysing appropriate data signals using timely event flows operating at sufficient levels of quality and fidelity. A signal is defined as a mapping of an observable event to contain at least one of the following two dimensions: A. Time: the observation is timestamped as actual through being reported on the occurrence; B. Event type: the observable event type is described in the signal set. Two primary indicators of timeliness enable the identification, controls, and management of delays. These indicators are: A. Latency: the delay, relative to the time of the event type occurrence, up to the time point it is reported; B. Lagg: similar to latency but applied to a specific flow of data. Data sources generating signals generally can be classified into two groups: External sources: information published by authorities (e.g. police, financial authorities); Internal sources: company internal publications such as an AML incident report or a fraud detection warning. Signals arriving at the detection engine are subjected to an enrichment pipeline using internal and external data.

External and internal sources of signals may also be fused to create composite data at a higher level of delta cross-validation among the streams. Business events generating specific classes are considered standard; for example, a counterparty data update is a typical KYC event flow. Enrichment pipelines associated with external and internal signals apply enrichment, data cleansing, quality detection, and correction processes. Results of enrichment are three-fold. First, appropriate auxiliary (enrichment) data items enabling model responses, detection, and analysis are created (forward engineering). Second, empirical knowledge on possible model response behaviour within detected time-frames is encoded (backward engineering). Third, buried hidden factors influencing model response are extracted and associated with the specific data type (sideways engineering). Enrichment is a pivotal operation apart from being a preprocessing task, as it includes the synthesis of new features and the establishment of covariate references. Although backward and sideways engineering may not apply for detection, performing synthetic feature engineering is fundamental.

5. Conclusion

The Event-Driven Compliance Intelligence Framework explicitly integrates the detection of financial crime with the extraction of compliance intelligence in support of KYC, sanctions, and anti-money laundering processes. A layered architecture, comprising the data, running, and governance layers, fulfils the cognitive systems requirement of regulatory science by embedding regulatory requirements, risk indicators, and compliance intelligence workflows into the detection process. Temporal alignment of processing within the detection engine conveys compliance risk signals at or near the moment of relevance. Signal extraction transforms external and internal sources into event streams for subsequent enrichment and feature engineering; risk scoring for detection enables drift-handling through changes in distribution and temporal coherence.

Effective governance mitigates the ethical and reputational dangers inherent in artificial intelligence by defining objectives for online learning and drift detection. Combining




prospectus characteristics with SACE's data provenance establishes the objectives of performance and compliance risk continuity.

References

1. Canhoto, A. I. (2021). Leveraging machine learning in the global fight against money laundering and terrorism financing: An affordances perspective. *Journal of Business Research*, 131, 441–452.
2. Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review*, 40, 100402.
3. Maguluri, K. K., Pandugula, C., Kalisetty, S., & Mallesham, G. (2022). Advancing Pain Medicine with AI and Neural Networks: Predictive Analytics and Personalized Treatment Plans for Chronic and Acute Pain Managements. *Journal of Artificial Intelligence and Big Data*, 2(1), 112-126.
4. Stojanović, B., Božić, J., Hofer-Schmitz, K., Nahrgang, K., Weber, A., Badii, A., Sundaram, M., Jordan, E., & Runevic, J. (2021). Follow the trail: Machine learning for fraud detection in FinTech applications. *Sensors*, 21(5), 1594.
5. Domashova, J., & Mikhailina, N. (2021). Usage of machine learning methods for early detection of money laundering schemes. *Procedia Computer Science*, 190, 184–192.
6. Gupta, A., Dwivedi, D. N., Shah, J., & Jain, A. (2021). Data quality issues leading to sub optimal machine learning for money laundering models. *Journal of Money Laundering Control*, 25(3), 551–555.
7. Mandala, G., Reddy, R., Nishanth, A., Yasmeen, Z., & Maguluri, K. K. (2023). Ai and ml in healthcare: redefining diagnostics, treatment, and personalized medicine. *International Journal of Applied Engineering & Technology*, 5(S6).
8. Gupta, A., Dwivedi, D. N., & Jain, A. (2021). Threshold fine-tuning of money laundering scenarios through multi-dimensional optimization techniques. *Journal of Money Laundering Control*, 25(1), 72–78.
9. Pettersson Ruiz, E., & Angelis, J. (2022). Combating money laundering with machine learning: Applicability of supervised-learning algorithms at cryptocurrency exchanges. *Journal of Money Laundering Control*, 25(4), 766–778.
10. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, 116429.
11. Ali, A., Abd Razak, S., Othman, S. H., Eisa, T. A. E., Al-Dhaqm, A., Nasser, M., Elhassan, T., Elshafie, H., & Saif, A. (2022). Financial fraud detection based on machine learning: A systematic literature review. *Applied Sciences*, 12(19), 9637.
12. Pamisetty, A. (2022). Big Data can Generate Major Opportunities for Manufacturing Supply Chains. *International Journal of Scientific Research and Modern Technology*, 1(12), 238-251.

13. Ileberi, E., Sun, Y., & Wang, Z. (2022). A machine learning based credit card fraud detection using the GA algorithm for feature selection. *Journal of Big Data*, 9, 24.
14. Tertychnyi, P., Godgildieva, M., Dumas, M., & Ollikainen, M. (2022). Time-aware and interpretable predictive monitoring system for anti-money laundering. *Machine Learning with Applications*, 8, 100306.
15. Deng, Z., Xin, G., Liu, Y., Wang, W., & Wang, B. (2022). Contrastive graph neural network-based camouflaged fraud detector. *Information Sciences*, 618, 39–52.
16. Li, Q., He, Y., Xu, C., Wu, F., Gao, J., & Li, Z. (2022). Dual-augment graph neural network for fraud detection. *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 4188–4192.
17. Ti, Y.-W., Hsin, Y.-Y., Dai, T.-S., Huang, M.-C., & Liu, L.-C. (2022). Feature generation and contribution comparison for electronic fraud detection. *Scientific Reports*, 12, 18042.
18. Lakkarasu, P., & Kalisetty, S. (2022). Hybrid Cloud and AI Integration for Scalable Data Engineering: Innovations in Enterprise AI Infrastructure. *Math. Stat. Eng. Appl.*, 71(4), 16774-16784.
19. Singh, A., Jain, A., & Biabale, S. E. (2022). Financial fraud detection approach based on Firefly optimization algorithm and support vector machine. *Applied Computational Intelligence and Soft Computing*, 2022, 1468015.
20. Kumar, S., Ahmed, R., Bharany, S., Shuaib, M., Ahmad, T., Tag Eldin, E., Ur Rehman, A., & Shafiq, M. (2022). Exploitation of machine learning algorithms for detecting financial crimes based on customers' behavior. *Sustainability*, 14(21), 13875.
21. Zhang, Z., Ma, Y., & Hua, Y. (2022). Financial fraud identification based on stacking ensemble learning algorithm: Introducing MD&A text information. *Computational Intelligence and Neuroscience*, 2022, 1780834.
22. Li, R., Liu, Z., Ma, Y., Yang, D., & Sun, S. (2023). Internet financial fraud detection based on graph learning. *IEEE Transactions on Computational Social Systems*, 10(3), 1394–1401.
23. Astrova, I. (2023). Anti-money laundering powered by graph machine learning: "Show me your friends and I will tell you who you are." *Intelligent Decision Technologies*, 17(1), 243–261.
24. Pavlidis, G. (2023). Deploying artificial intelligence for anti-money laundering and asset recovery: The dawn of a new era. *Journal of Money Laundering Control*, 26(7), 155–166.
25. Alexandre, C. R., & Balsa, J. (2023). Incorporating machine learning and a risk-based strategy in an anti-money laundering multiagent system. *Expert Systems with Applications*, 217, 119500.
26. Yang, G., Liu, X., & Li, B. (2023). Anti-money laundering supervision by intelligent algorithm. *Computers & Security*, 132, 103344.
27. Tong, G., & Shen, J. (2023). Financial transaction fraud detector based on imbalance learning and graph neural network. *Applied Soft Computing*, 149, 110984.
28. Chen, Z.-Y., & Han, D. (2023). Detecting corporate financial fraud via two-stage mapping in joint temporal and financial feature domain. *Expert Systems with Applications*, 217, 119559.



29. Pamisetty, V. (2020). Optimizing Unclaimed Property Management through Cloud-Enabled AI and Integrated IT Infrastructures. Universal Journal of Finance and Economics, 1(1), 1-20.
30. Fanai, H., & Abbasimehr, H. (2023). A novel combined approach based on deep autoencoder and deep classifiers for credit card fraud detection. Expert Systems with Applications, 217, 119562.