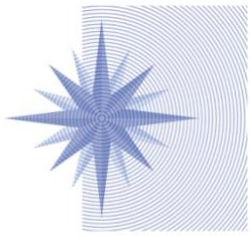

Banco Santander
Bank of America
Barclays
Citigroup
Deutsche Bank
Goldman Sachs
HSBC
JPMorgan Chase
MUFG Bank
Société Générale
Standard Chartered Bank
UBS

the Wolfsberg Group

**The Wolfsberg Group - Statement on Effective Monitoring for Suspicious Activity**

**Part II: Transitioning to Innovation**

## Introduction

In July 2024, the Wolfsberg Group (the Group) published the *Statement on Effective Monitoring for Suspicious Activity* (MSA), which outlined how financial institutions (FIs) can translate the Wolfsberg Factors (as follows) into an effective MSA programme:

1. Comply with Anti-Money Laundering (AML) and Counter Terrorist Financing (CTF) laws and regulations.
2. Provide highly useful information to relevant government agencies in defined priority areas.
3. Establish a reasonable and risk-based set of controls to mitigate the risks of an FI being used to facilitate illicit activity.

The 2024 MSA Statement encourages FIs to move beyond traditional transaction monitoring, introducing a vision for how FIs should target and develop effective outcomes as a priority across their monitoring programmes. It also stresses the need for innovation and explores those enabling factors that can maximise the identification of criminal activity. The Group believes that advancing these themes will aid FIs in transitioning from traditional rules-based programmes to more innovative approaches, including those that incorporate machine learning, artificial intelligence and automation.

The original *Statement on Effective Monitoring for Suspicious Activity* made it clear that the concept of MSA is wider than traditional automated transaction monitoring, as customer behaviour and customer attributes, when combined with the consideration of transactions, can provide broader insights into potentially suspicious activity.[1] As well as moving away from solely focusing on transactions, FIs must also recognise that their overall monitoring and investigations programme has likely evolved substantially since original automated TM systems were designed and implemented. Ensuring an FI's new approach to automated monitoring integrates and complements new capabilities is critical. Furthermore, as criminal networks continue to evolve

[1] With the publication of Part II of the MSA Statement, the original 2024 Statement will be renamed, *The Wolfsberg Group - Statement on Effective Monitoring for Suspicious Activity, Part I: Moving Beyond Automated Transaction Monitoring*.

at a rapid pace, as do national security priorities, FIs must embrace an innovation governance framework that is nimble and allows for prompt responses to evolving threats. This adaptability acknowledges that non-financial risk models do not have the same prudential impact on an FI's operations and therefore do not warrant a lengthy and invasive model risk management approach. Supervisory agencies in the US made it clear in an interagency statement on financial crime related models that they "support efforts by banks to innovate and update [their financial crime] systems and models to quickly adapt to an evolving threat environment".[2] Finally, as the economic barriers to leveraging automation, artificial intelligence and machine learning have relaxed, an FI's MSA strategy must be positioned to fold these capabilities into their financial crime risk management programme in a responsible, yet effective, way.

Since the 2024 publication, the Group has focused on building a responsible transition framework for innovation in MSA. The Group's framework, detailed within this statement, is based on three pillars:

1. Transition and validation
2. Model risk balanced with financial crime risk
3. Explainability to demonstrate transparency in coverage and effectiveness

The core aspects of the framework are drawn from the experiences of the member banks as they have advanced their own approach to monitoring for suspicious activity in their core jurisdictions, including learning from engagements with key government authorities when permission or notification was necessary. The transition framework has also benefited from a series of technical workshops, both virtual and in-person, with financial intelligence units, policymakers, and supervisory authorities across the world.

**A) Effective Transition and Validation**

The transition to a new MSA approach begins with acknowledging that the FI's legacy approach is nearing the limits of its effectiveness, and that the new approach must be distinct because the FI's overall financial crime risk management framework has evolved in line with the growing sophistication of criminal behaviour. Understanding why the legacy approach is less effective and more inefficient, and why the new approach should be different by design, is crucial to avoiding low-value comparisons between past and future performance.

The new approach should improve the FI's ability to manage financial crime risk more effectively and efficiently. It should work in concert with the broader financial crime control environment, which has evolved since the beginnings of automated transaction monitoring. This may result in an approach that does not detect everything historically considered suspicious but importantly improves the identification of quality leads for law enforcement, uncovering new risks that the legacy platform did not. This very well may represent an evolution in the FI's risk appetite, which the FI should be prepared to articulate to supervisory authorities. A stronger, more mature MSA programme is one that acknowledges the ineffectiveness and inefficiencies of historically imprecise scenarios (commonly referred to as a "drag net" approach) that led to the reporting of low quality, questionably suspicious filings. Continuing that legacy approach to monitoring and filing serves as a deterrent to innovation in a world where models only become more effective when they can learn from high quality examples, and, when possible, through law enforcement feedback.

Transitioning to a more effective MSA programme begins with re-establishing the programme's desired outcomes, tailoring the new approach to those desired outcomes, and then validating,

[2] *Interagency Statement on Model Risk Management for Bank Systems Supporting Bank Secrecy Act/Anti-Money Laundering Compliance.* Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and the Office of the Comptroller of the Currency (9-Apr-2021).

once implemented, that the new approach meets expectations on those revised measures of success.

*Establishing the desired outcomes of the MSA approach, in line with the FI's risk appetite and the evolving risk management framework*

Many elements viewed as critical today for a holistic suspicious activity monitoring strategy were not present when an FI's initial transaction monitoring platform was implemented. FIs now commonly rely on data analytics specialists to write targeted queries, focused on specific criminal typologies, or use off-the-shelf products from vendors to improve the accessibility of data analytics. The staffing of experienced investigators in the FI's internal financial intelligence unit has likely also increased substantially. Front office staff are more attuned to financial crime risk through tailored and more detailed training. This result is filing better quality internal reports of potentially suspicious activity and in higher volumes.

The FI's new approach to using an automated monitoring system for detecting suspicious activity must recognise these advancements, strengthening the larger monitoring strategy by integrating with, and complementing, these parallel activities. In practice, this means that the success criteria for the automated monitoring approach have evolved. Whereas previously the automated platform was the predominant, almost single pillar upon which suspicious activity identification was built, there is no longer the need for such a "catch-all" dependency.

In line with these advances, the following indicators – for consideration – will assist an FI in redefining the performance indicators associated with the new approach:

- **Priority risk coverage** – the ability of the new model to detect financial crime associated with high impact criminal activity not addressed by other controls, as determined by the FI's own threat assessment and the priorities set by national authorities. This may include recognising the value of controls outside of the traditional AML suite (e.g., in identity fraud or market abuse) that are nevertheless complementary in identifying financial crime.
- **Expanded risk indicator coverage** – the ability of the new model to consider high quality risk indicators in its design that draw on a variety of data points, not only transactions, including in instances where the underlying offence may be unclear but there is a strong suspicion of money laundering.
- **Precision rate** – the proportion of positive predictions (generated alerts) that are actually positive ("true positives" – instances of suspicious activity).
- **Recall rate** – the proportion of actual positive cases ("true positives" and "false negatives" – all known instances of suspicious activity) correctly identified by the model ("true positives") – i.e., the desired levels of risk or suspicion correctly identified by the model.
- **Internal and external feedback on Suspicious Activity Report (SAR)/ Suspicious Transaction Report (STR) quality** – the ability of the model to capture suspicion that will lead to higher quality SARs and/or escalations. In addition to feedback from national FIUs or law enforcement agencies (if available), this likely involves an internal exercise evaluating historical SAR filings to recognise true quality, as measured against those filings that may have been more defensive in nature and not based on sufficient confidence levels of crystallised financial crime. The findings from this evaluation would then be incorporated into the model training and acceptance criteria.
- **Downstream integration** – the ability of an analyst working an alert to have appropriate access or protocols in place to engage directly or indirectly with customers (e.g. through a relationship manager or contact centre) and clarify activities.

Underpinning these success criteria should be a constant reflection on how the new approach will be integrated into the FI's evolving risk management framework, minimising overlap in control coverage, optimising the use of new data points beyond transactional data and ensuring complementary processes are adequately designed to facilitate collaboration and coordination.

*Transitioning from an existing to a new approach should be tailored to the desired outcomes of the new approach*

Rigorous testing and validation are essential to ensure the right processes are in place and risks are managed appropriately. FIs must establish how the testing is performed, how they gain confidence in the end-to-end process and how they ensure that potential risks are being detected.

Given the divergence between the legacy approach and the new focus, particularly when transitioning to machine learning and AI-based models, classic transition methods may not be suitable. Testing environments have also improved substantially over the years. Running the old system and the new system in parallel for a set period (the "parallel run"), for example, is an artefact of a less advanced testing environment often paired with a high concern for execution risk as part of the FI's overall change management framework. The ability of the FI to run a proof-of-concept in a test environment has improved, with new model validation techniques achieved through assessing the new model's performance against historical data and balancing its recall with the identification of new areas of risk. Proof-of-concept evaluations should focus on whether the new model meets the revised performance criteria (as established in the previous section), in line with the FI's risk appetite.

There are also a series of additional areas that likely demand increased focus during the transition phase and may represent a break from the legacy approach. Examples include:

- **Product coverage** – ensuring there is clarity on those products, accounts, and transaction types that the new approach includes within its scope, and that there is clarity among all stakeholders on what is out of scope. It may be the case that the new proposed approach is not fit for purpose for some products or services, or alternatively, existing products and services not covered by the legacy platform may now benefit from the new approach.
- **Case investigation readiness** – FIs should address the risks of managing more sophisticated cases by ensuring investigators have the necessary domain knowledge. This can be achieved by investing in upskilling and training to prepare the workforce for advanced expectations in alert management with a focus on model capabilities, interpretability, and operational change. Under a rules-based approach, the catalyst for the alert is clear to the analyst but also blunt. In more advanced MSA programmes, a confluence of factors and trends will have prompted the alert. This complexity presents an opportunity: by leveraging advanced technologies such as large language models (LLMs), FIs can surface the most predictive features and contextual signals behind alerts. These models can generate investigator-ready summaries that clarify the underlying risk and suggest next-best investigative actions.

Of course, in addition to the focus areas captured in this statement (which are specific to MSA innovation), FIs should maintain an appropriate change management governance structure to oversee the transition. This includes technical stability and robust data availability to support the new model effectively, assessing the condition of the overall model portfolio before releasing new iterations, and in general, avoiding the deployment of new solutions during periods of heavy remediation.

*Validation and the role of historical comparisons*

As the focus shifts from transition to validation, it is critical to recall that the new approach is not aimed at replicating the legacy approach or "adding on" to what the legacy approach already does. Simulations and backwards-looking testing on historical transactional data are aimed at validating that the new automated MSA platform performs at an acceptable level against the newly defined measures of success.

This is not to say that comparing the outputs of the old and new approaches lacks any value during the validation phase. For instance, if the FI already rates SAR quality, and has used this method as a success indicator, understanding the quality levels of SARs identified by the old approach but missed by the new one is relevant. Equally, if the FI believes that other elements of the control environment would address the SARs not identified in the new approach, confirming that view will be an important aspect of the validation phase before fully shifting to the new approach. It is also possible that what was previously considered "suspicious" under the legacy approach may no longer be recognised as suspicious as the expansion of available data points (i.e. beyond transactions) brings new insights to the customer's behaviour that sufficiently clarifies what was previously unclear or unexplainable.

Improved productivity will also be a key component of the validation process, evaluating new alerts not generated in the old approach, to assess how well they align to revised objectives on areas like priority risk coverage and risk indicators. Subject matter experts will likely be required to assess a sample of alerts and rate their quality.

Overall, while initial model validation should confirm that the new approach is sufficiently fit-for-purpose, the process should be viewed as continuous, drawing on core concepts from model risk management and explainability as described in the sections that follow.

**B) Model risk vs financial crime risk**

Model risk occurs when a model, which measures quantitative information, potentially leads to adverse outcomes for the FI due to model errors or the inappropriate use of modelled outputs. Financial crime risk, alternately, is the risk that an FI's products and services will be exploited for criminal activity.

In an environment with rapidly evolving threats, FIs must be able to introduce new typologies at pace and drop ineffective typologies quickly, balancing both the model risk and the financial crime risk appropriately. Navigating model risk management practices (MRM), including supervisory expectations and the cycles of independent review by audit and assurance, continues to serve as a significant barrier for FIs as they seek to innovate. This is generally attributed to a "one-size-fits-all" mindset that equates financial crime risk with prudential/financial risk, leading to an overly robust model governance structure and preventing FIs from innovating in line with a more iterative (e.g., "fail fast") culture to detecting financial crime. As competent authorities continue to raise expectations of an FI's ability to demonstrate progress on law enforcement priorities, tailoring the MRM framework to embrace the FI's new approach to MSA should be considered a foundational element to a sustainable programme.

*Considering the financial crime risk when evaluating the impacts of model risk*

Independent validation teams typically use a one-size-fits-all approach, i.e. they test all models classified in the same model risk tier equally because of the expected regulatory scrutiny on models. Financial crime related models, however, should not be considered at the same risk tier as prudential/financial risk related models, which should be reflected in an FI's internal MRM policy. It is critical for independent validation teams to have a clear understanding and ability to balance model risk versus financial crime risk to inform business decisions. The decision of whether to operate a model should be based on an appropriate balance between model and

financial crime risk, i.e., the risk of using a model that is not fully optimised (model risk) versus the risk of not instituting the model to improve financial crime detection. If an imperfect model identifies financial crime over and above what the institution would have detected without the model, it may still be appropriate to run the model despite the existence of certain model risks. The risk of a single model failing due to a model deficiency does not necessarily result in the failure of the FI to detect financial crime. FIs have multiple methods and controls to monitor suspicious activity, and a single model is not the central point of failure for financial crime risk management programmes.

Most MRM frameworks commonly use materiality, reliance and complexity to determine model risk tiering; however, the weighting of the model risk tiering should also rely upon the materiality of financial crime risk. Materiality in model risk is where there is the greatest variance in the elements used by FIs to measure this component. Materiality in financial crime risk is a function of whether the FI is providing useful information to law enforcement. The materiality of model risk needs to consider a variety of factors to align better to the expected financial crime risk it may identify.

Although the Group does not recommend any specific materiality criteria, the Group encourages FIs to consider the following principles to differentiate the risks that the models pose:

- **Financial crime detection effectiveness** – some FIs use control "productivity" (i.e. number of SARs filed) to help measure the materiality of the model risk. Although productivity could be used as a component of model risk, because an unproductive model could lead to missed risk, it should not be the only element used to define materiality of the model risk. There are many reasonable explanations for models to have low productivity, as defined by model monitoring metrics, but provide a high degree of effective financial crime risk management outcomes for the FI.
- **Products and services** – considering the financial crime risk of the products and services that the model covers as defined by the institution's product risk assessment methods.
- **Coverage of transaction exposure** – considering the percentage of transactions analysed by the model against the total population. This aligns the impact of model failure to the business impact, as measured by transactions, and reflects a model's primary exposure to risk. The rating would be stable as transaction volumes are typically stable, with seasonality trends established within the calculation.
- **Concentration** – considering the concentration of the risk that is created when a model is solely relied upon to deliver coverage for a particular risk. A model with concentration risk has greater materiality because the potential for adverse outcomes for model failure is greater.
- **Size and scale of the business** – considering the size and scale of the business monitored by the MSA system is another known option for measuring model risk exposure. When a model's exposure is based on the business unit that it covers, the materiality risk is aligned with the business impact of model failure.

Overall, it is the FI's responsibility to ensure that internal MRM policies provide the appropriate governance framework to facilitate innovation in monitoring for suspicious activity. Supervisory teams will examine an FI against its own policies and procedures, and thus if the FI's own governance documents do not account for the unique situation of financial crime models, there can be no expectation that an incoming examiner will hold the FI to a different standard.

*Create streamlined, non-redundant independent oversight*

FIs use multiple functions to perform independent validation, oversight and testing of MSA programmes. These will primarily include audit, assurance and MRM. The Group acknowledges

that these reviews are essential for effective governance and oversight; however, overlapping and redundant processes by these functions also create significant delays, resulting in less effective ongoing monitoring controls.

Independent validation teams often assess the same MSA models from similar angles, and each of these functions requires their own independent and siloed process. As a result, implementing new MSA models covering new risks takes months instead of weeks.

Considering that financial crime typologies are constantly changing, a slow response due to inefficient oversight can weaken the FI's ability to detect new financial crime risks in a timely manner. This inefficiency is further compounded as the same resources who develop the model must engage with multiple, independent reviewers at the expense of enhancing detection models. This is exacerbated when different independent validation teams provide contradictory recommendations, forcing developers not only to respond to inconsistent expectations, but also to spend an inordinate amount of time addressing these findings. To reduce the impacts of these issues, FIs need a clear, structured and non-redundant validation framework across the three lines of defence, and closer alignment among all validation functions. For example:

- MRM, a second line of defence (2nd LoD) function, validates methodology, model assumptions and limitations, data quality and statistical soundness, including the technical and methodological soundness of the MSA model. This is performed through the validation of model design, inputs, assumptions, thresholds and outputs, as well as through back-testing and benchmarking to assess model performance and may make recommendations to improve model precision and recall.
- Assurance (sometimes referred to as compliance oversight, compliance monitoring, or similar), also a 2nd LoD function, performs targeted reviews as directed by risk stewards, focusing on the effectiveness of controls in respect to alignment with the FI's policies and standards.
- Audit, a 3rd LoD function, assesses governance, control frameworks, and compliance with regulatory requirements, including alert handling, escalation procedures, and compliance with policies. Furthermore, audit focuses on the identification of potential gaps in managing financial crime risk (such as typologies), procedural weaknesses and provides an independent view on the overall MSA system effectiveness.

*Balancing the effort on historical analysis with developing future capabilities*

FIs spend a disproportionate amount of time validating existing models using contemporary model validation techniques that are not commensurate with the level of financial crime risk coverage provided. These often require revisiting and revising long-established justifications for parameters set by subject matter experts, relitigating methodological approaches applied historically for simpler rules, and incurring administrative overhead on incremental adjustments that do not materially increase the effectiveness of financial crime detection. While this is not to suggest that existing and/or simpler coverage approaches do not warrant periodic review, proportionally allocating resource, budget and risk appetite to drive more innovative techniques would create a more agile approach to delivering future detection capabilities.

In order for FIs to be able to adapt to emerging risk, a paradigm shift is required with respect to prior model outcomes being used to validate newer models. Retroactively analysing outcomes from old models, especially those with different or obsolete risk targets, stifles the agility and speed with which FIs can address new risks. Another construct of current model risk governance that poses challenges to an FI's ability to maximise innovative MSA techniques is the expansion of model inventories. As an example, a model built for a certain typology can be rolled out to multiple business lines or jurisdictions, yet from a model inventory perspective, each application of the core model to these permutations is its own line item in the inventory. Thus, each line item

can require, depending on an FI's MRM practices, equal rigour of model risk governance, subjecting each iteration to the same standard as the main model.

A more agile approach would be to subject the core model to the requisite validation, with subsequent "sub"-iterations of the model (as they are rolled out to other business lines and/or jurisdictions) subject to a subset of testing and validation focused on the delta (e.g., data feed validations) between the new iteration and the approved core model. This has the effect of addressing overarching model risk expectations, while allowing the subsequent roll-out of the model across the FI's footprint to be more responsive and targeted to specific areas where the most risk can be mitigated.

**C) Explainability**

Ultimately, the FI's advanced approach to MSA must be accompanied by an equal level of transparency to ensure trust is established and maintained in the decision-making process as an FI identifies and reports suspicious activity. An FI must be prepared to explain the lifecycle of the MSA approach according to three core lenses:

- Risk coverage;
- Model design, development and calibration (i.e., how the model arrives at the decisions or predictions behind the model output); and
- Model usage (i.e., how the model users should understand the novel model output and control).

The traditional landscape begins with leveraging a country's relevant red flags, allowing FIs to assess product and customer risks. These red flags are linked to typologies and aligned with rules-based automated monitoring solutions. But as machine learning solutions are adopted, the explanation of the financial crime detection system changes significantly. The focus shifts to leveraging multiple machine learning outcomes to develop a holistic monitoring strategy, allowing FIs to consider customer behaviour across multiple typologies to detect criminal behaviour and risk, simultaneously. The advanced model is then built and trained through design and calibration, placing emphasis on understanding (and being able to explain) how the model operates and how to assess its performance. The third lens focuses on the analyst – how the consumer of the model's output interprets the findings and uses that information to advance the investigation further and assess suspicion.

*Risk coverage*

In a traditional rules-based monitoring solution, evidencing risk coverage involves mapping risk indicators to typologies and scenarios. In machine learning, typologies are mapped to data features which are fed into models. This process allows FIs to create and maintain a mapping inventory that consists of risk indicators or typologies and their relevant features to evidence risk coverage.

To develop an effective MSA process, it is essential to ground the design of the machine learning model in a comprehensive risk assessment exercise. This will involve identifying the specific risk indicators that the FI is required to monitor, as outlined by regulatory guidance and industry best practice.

The variables used in a machine learning model should align with such risk indicators, ensuring they are not only statistically significant but also contextually meaningful. By grounding variables in a specific risk assessment, FIs ensure that their models are not generic but tailored to the unique risk profile of their customers and national priorities, as communicated to the FI through the host country's national risk assessment, or more operational/tactical advisories published by the national FIU.

This activity should be complemented by performing model output analysis, where the outputs of a model are subject to a labelling exercise to understand the risk captured. This provides further comfort to risk owners and enhances the FI's ability to explain how the new approach covers different financial crime risk types. Labelling techniques used for training will provide insights for effective oversight, correlating feature importance back to typologies.

*Model design, development and calibration*

Rules-based methods are effective for well-defined risks, including certain reporting requirements that could be automated. Supervised methods are effective when there is clear and distinctive investigation output from the historical investigation, which can be used to train the model to achieve desired outcomes – reinforcing the need to focus on quality. Unsupervised methods are crucial for identifying unknown risks and emerging typologies. The Group recommends a hybrid approach to MSA in combining rules-based systems, supervised methods, and unsupervised methods. Such an approach leverages the strengths of each methodology, ensuring comprehensive risk coverage.

Traditional rules-based model design is based on the logical construct of rules and thresholds which are easier to interpret and explain. Machine learning models, on the other hand, consist of employing a more sophisticated data training set, feature selection analysis, and statistical algorithms that generate outputs based on their predictive power. There are a huge variety of different algorithms and machine learning models for design and deployment, including supervised and unsupervised methods.

The explainability for any model should cover the whole spectrum of model design considerations and processes. This includes data selection (with important reference to the previous section on risk coverage), the data features that were influential in the prediction of risk, as well as the data set used for training. The choice of machine learning algorithms should be able to be explained with clear articulation of the strengths and weaknesses of the methods chosen. FIs should be able to highlight the performance of machine learning models in capturing suspicious activity and ensuring appropriate risk coverage.

FIs should also be able to demonstrate that the models are subject to robust calibration and optimisation given that any model will deteriorate over time and there must be an effort to re-train and re-calibrate to ensure the model remains fit-for-purpose.

*Model usage*

Model users should understand and be able to interpret how the model can be used effectively, including in the context of other relevant, complementary controls. Traditional alerts are generated when certain activities breach pre-defined rule thresholds. A similar articulation can be used for machine learning models across the many variables considered. Since the output of the model identifies the probability of how likely it is that certain activities are suspicious, there is also a risk appetite consideration of how likely it is that the model will miss suspicious activities (equivalent to below-the-line testing). This risk appetite must be explained and agreed to allow the model design to have a clear objective and tolerance level.

When the outputs of models are sent for investigation, there is always a need to explain why an alert was generated in the first place. This explainability is crucial for the investigator to understand the rationale to assist in their investigation process. This can be illustrated, for example, by the features that influence the model output, the risk associated to those features, and/or the relevant risk indicators triggering behind those features. FIs should consider displaying such information through visualisation tools that facilitate an analyst's understanding of how various activities come together to seem suspicious.

If the output needs to be strengthened by another mechanism, including pre-defined conditions or guardrails, these controls should be well understood and continuously revised for enhancement and review. This will ensure that these models continue to operate within the FI's risk appetite.

**Conclusion**

The Group believes that innovation is an essential enabler for FIs to improve their ability to provide law enforcement with higher quality information in the fight against criminal activity. Embracing emerging technology presents a wealth of opportunity as highlighted in the July 2024 MSA statement, acknowledging that legacy methods and approaches are increasingly recognised as inefficient and ineffective at producing material outcomes that are useful to law enforcement.

Within an FI, there are a number of prerequisites to a successful transition to more innovative monitoring mechanisms. Firstly, senior management must be aware and accept any changes to risk appetite in reporting crystallised financial crime risk through SAR filings. Secondly, internal auditors, assurance teams, and model risk management functions must embrace the tailored governance and oversight approach to ensure the pace of production and approval for financial crime models aligns with the threat. Finally, financial crime compliance departments must be ready to explain the new approach, demonstrate and track how the approach aligns with the FI's revised measures of effectiveness, and train their investigative analysts to understand how to interpret the more complex results drawn from sophisticated monitoring features. As technology and associated outcomes become more effective, it becomes all the more important for an FI to be able to articulate how its MSA platform operates and how the FI will continue to iterate its approach through calibration and performance oversight. The Group's members are committed to improving the MSA process through responsible innovation – deploying advanced technology, focusing on effectiveness, and ensuring transparency in how the FI's control environment responds directly to the FI's established risk appetite.

## Glossary

**Accuracy:** Overall proportion of correct predictions across all classes.

**Artificial Intelligence (AI):** The ability of a computer or computer-controlled robot to perform tasks commonly associated with intelligent beings, such as reasoning, discovering meaning, generalising, or learning from past experience.

**Below-the-line Testing (BTL):** Involves conducting tests by lowering the thresholds or criteria below the baseline. This helps identify the point at which the system may generate false negatives, potentially missing potentially suspicious activity.

**Classification:** Predicting discrete categories (e.g., fraudulent/legitimate transactions).

**Crystallised Risk:** Realised risk events, i.e. risk events that have occurred rather than those that are theoretical.

**Explainability and Explainable AI (XAI):** The ability to understand how a model arrived at its predictions. XAI refers to the development and use of machine learning models that are understandable and transparent to humans. Many AI systems, particularly those using complex algorithms like deep neural networks, can be seen as "black boxes" where the internal workings and reasoning behind their outputs are unclear.

**False Negative:** An instance incorrectly classified as negative.

**False Positive:** An instance incorrectly classified as positive.

**Feature Engineering:** Transforming raw data into features suitable for machine learning models.

**Generative AI:** Algorithms that create new content based on existing data.

**Large Language Model:** An advanced computer programme capable of understanding and generating human-like text by learning from vast amounts of written language data.

**Machine Learning:** A subfield of artificial intelligence (AI) that uses algorithms trained on data sets to create self-learning models capable of predicting outcomes and classifying information without human intervention.

**Model:** A representation of information learned from data that can be used to make predictions.

**Model review / validation**: A model's accuracy, reliability pre- and post-deployment through evaluation and approval processes.

**Monitoring for Suspicious Activity:** Various control elements that identify the risk of customer behaviour.

**Parallel Run:** Running both the existing and new systems simultaneously during transition.

**Precision**: Proportion of positive predictions (generated alerts) that are actually positive ("true positives" – instances of suspicious activity).

**Recall:** Proportion of actual positive cases ("true positives" and "false negatives" – all known instances of suspicious activity) correctly identified by the model ("true positives").

**Supervised Learning:** Uses labelled data to train algorithms to make predictions or classifications.

**Transaction Monitoring:** The automated or manual process of monitoring transactions after their execution in order to identify unusual transactions, including monitoring single transactions as well as transaction flows, for subsequent review and, where appropriate, reporting to the authorities.

**True Negative:** An instance correctly classified as negative.

**True Positive:** An instance correctly classified as positive.

**Unsupervised Learning:** Learning from unlabelled data where the model identifies patterns on its own.