HONG KONG MONETARY AUTHORITY
香港金融管理局

# Supporting A.I. Adoption in Fighting Financial Crime

June 2026

# Contents

# 01 Foreword

*Raymond Chan, Executive Director (Enforcement and AML)*

Over the past year, I have spoken frequently about the rapid evolution of financial crime risks and the increasing importance of technology in reducing the harm which it causes while also safeguarding the integrity of our financial system.

That message bears repeating – standing still in the face of such rapid change is not an option.

Digitalisation has fundamentally altered the way in which financial services are delivered as well as the speed, scale and sophistication of criminal activity. Fraud networks operate across borders and leverage technology to adapt quickly. The rise and impact of Artificial Intelligence (A.I.) driven scams are making detection harder and more scalable. At the same time, synthetic A.I. powered identities have the potential to challenge the effectiveness of existing approaches to verification. In this environment, traditional approaches to transaction monitoring (TM) and static, rule-based controls are under growing strain. To remain effective, our defences must evolve at an equal pace.

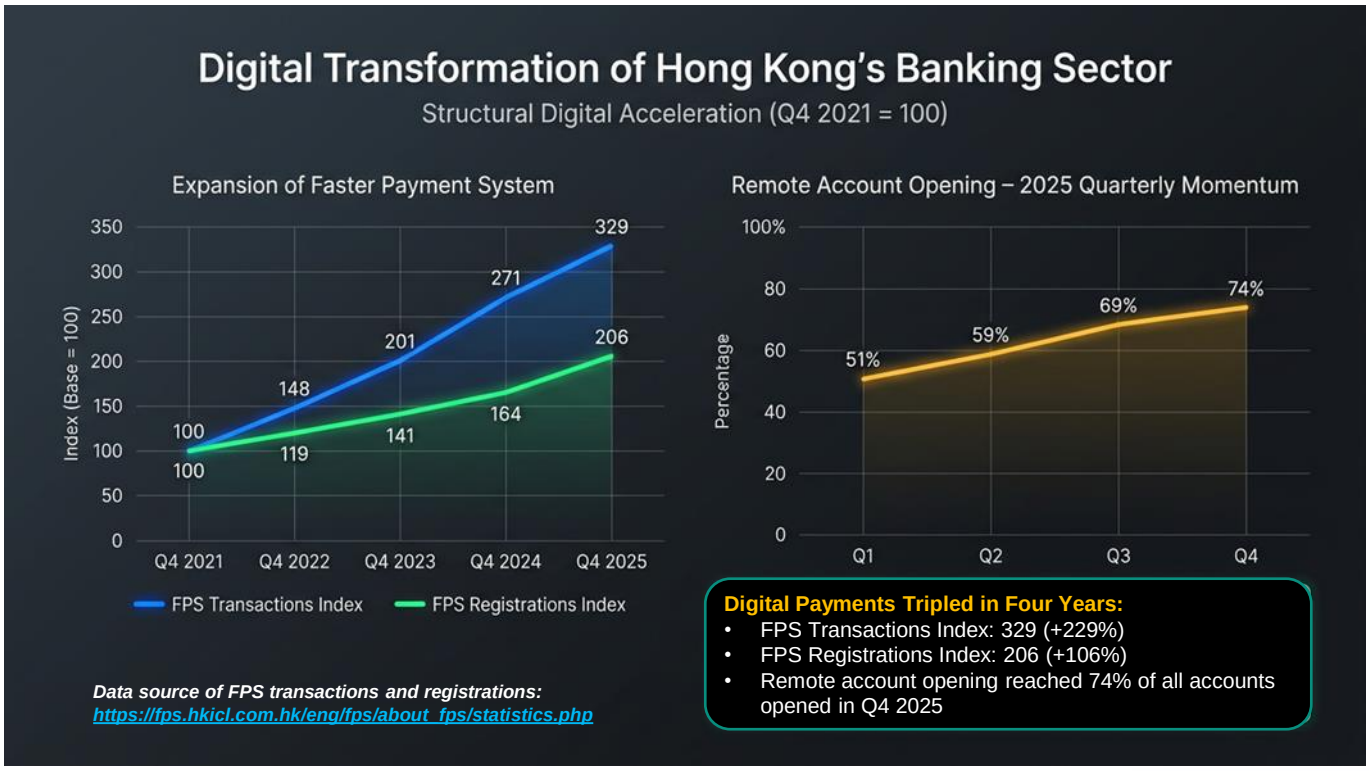

A.I. is no longer an emerging concept in anti-money laundering (AML) frameworks. It is already embedded in many institutions' monitoring, analytics and investigative processes. We have seen tangible benefits where A.I. and network analytics are used to identify suspicious accounts earlier, uncover hidden connections, and intervene before harm escalates. When this data is shared through Hong Kong's maturing framework of public-private partnerships (PPPs), significantly improved results and disruption are being achieved.

These are meaningful advances. However, innovation must be accompanied by accountability.

Where A.I. systems influence risk assessments or decision pathways, the responsibility remains firmly with senior management and Boards. Technology does not replace governance. Institutions must be able to explain how models operate, how outcomes are validated, and how risks such as bias, drift or over-reliance are managed. Supervisory tolerance for experimentation exists within well-defined guardrails — not outside them.

The next phase of development will see a more consistent sector-wide move beyond efficiency gains. A.I. should not be confined to accelerating documentation or optimising workflows. Its real potential lies in strengthening early detection, enhancing intelligence-led monitoring, strengthening complex decision making and enabling more proactive intervention against bad actors.

To achieve this, the ecosystem must ensure that innovation is underpinned by strong governance frameworks, supported by cross-line collaboration, and aligned with clearly articulated risk appetites.

Our objective is clear. To foster responsible innovation that materially improves the effectiveness of Hong Kong's anti-money laundering and counter-financing of terrorism (AML/CFT) regime. By combining technological capability with disciplined oversight, we can strengthen trust in our financial system and better protect the community we serve.

Our aim in this programme is to encourage the industry to set ambitious goals and move towards them collaboratively.


*Standing still is not an option. Banks need to adapt and constantly strive for greater efficiency and effectiveness, adopting proven technologies at scale and being prepared to explore and test new technologies as threats develop. Otherwise, they will be overwhelmed by threats and noise from systems that do not adequately mitigate those threats.*

***Raymond Chan***
*Executive Director (Enforcement & AML)*

# 02 Hong Kong – The Changing AML Landscape

## Industry A.I. Adoption

Detecting, deterring and disrupting financial crime risks remains a cornerstone of financial stability and market integrity. Advances in technology have seen financial crime risks evolve in scale, complexity and sophistication requiring banks to adapt their risk-management and control frameworks.

The innovation of new technologies has also presented banks with opportunities to strengthen their risk-management playbook by deploying network analytics, A.I. and advanced behavioural analytics to reshape how they tackle different financial crime threats.

A further catalyst for how banks are reshaping core financial crime prevention capabilities is the closer collaboration between regulators, banks and law enforcement (for example through PPP like FMLIT[1] and private-private platforms such as FINEST[2]), which is helping to improve financial crime risk mitigation effectiveness. The ability to prevent financial crime, driven by enhanced collaboration, prediction and detection capabilities, has never been stronger. This collaborative innovation has also supported the acceleration of A.I. adoption and is challenging the way the industry thinks about transforming how financial crime risk is managed.

## Number of new suspicious accounts associated with crimes under investigation by law enforcement agencies shared via FMLIT

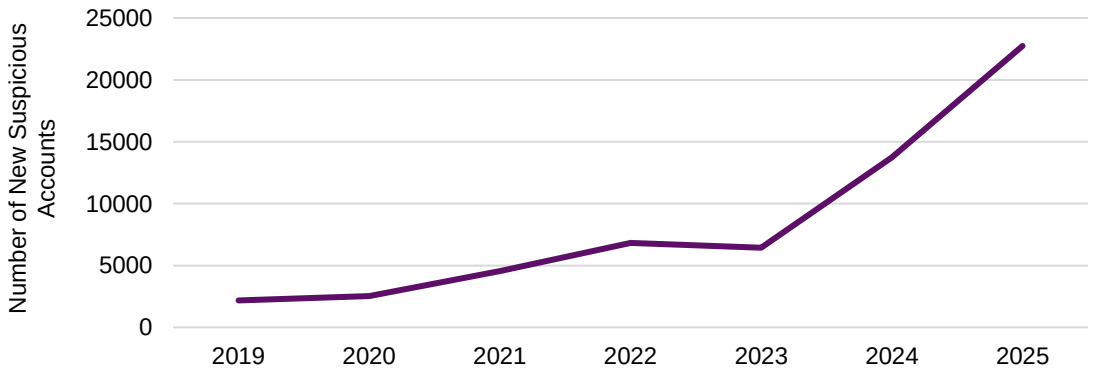


[1] Established in 2017, Fraud and Money Laundering Intelligence Taskforce (FMLIT) is a collaboration between Law Enforcement, the HKMA and the banking sector together with the Hong Kong Association of Banks (HKAB) under the leadership of the Police Force.

[2] Established in 2023, the Financial Intelligence Evaluation Sharing Tool (FINEST) is a bank-to-bank information sharing platform initiated by the HKAB with the guidance from the HKMA and operated by the Police Force to enable timely sharing of financial crime risk information among banks.

The financial services marketplace in Hong Kong has witnessed an expansion of banks (traditional and digital), payment providers and Stored Value Facility Licensees, and this in itself is fuelling greater investment in A.I. solutions – with varying levels of maturity – as these operators look to manage financial crime risk more effectively.

Against this backdrop, the HKMA has been promoting a multi-year industry initiative to drive the responsible adoption of A.I. technology. This has formed an integral part of the regulator's forward-looking vision for Hong Kong banks to continue their digital transformation journey under its Fintech 2030 strategy[3].

The HKMA asked banks with significant operations in Hong Kong to undertake a feasibility study[4] into adoption of A.I. in their AML/CFT monitoring systems and to formulate implementation plans based on the outcome of those reviews. The results showed broad alignment on the continued adoption of A.I. – and new technologies more broadly – to strengthen financial crime risk detection.

The outcome of the feasibility studies, alongside a number of live use cases in the Hong Kong market, were showcased at an industry workshop in November 2025. The key message from industry leaders at the event was that banks should not adopt A.I. for A.I.'s sake but deploy targeted solutions that are aligned with business strategy, operating models and levels of readiness.

A number of these solutions are outlined in the Case Study section of this publication.

## A.I. Adoption Trends in TM Workflows

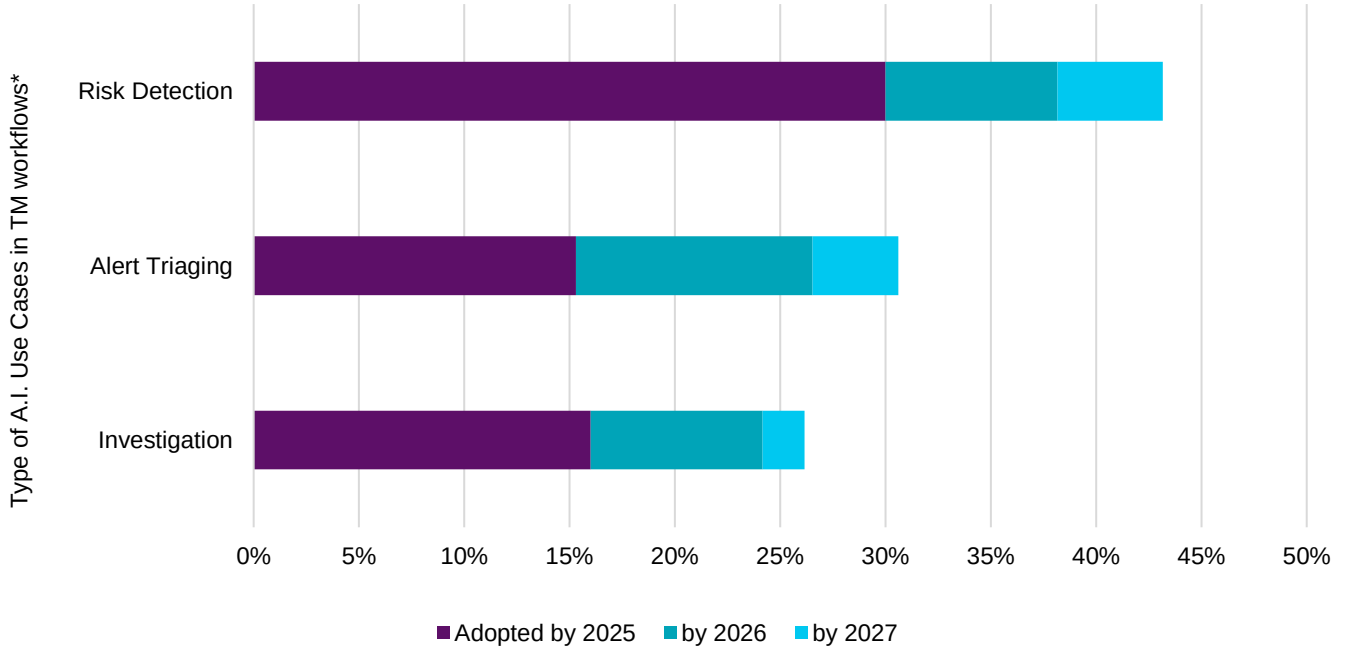


**One bank may have more than one A.I. use case*

[3] HKMA Press Release on "The HKMA Unveils "Fintech 2030" at the Hong Kong FinTech Week 2025" issued on 3 Nov 2025.
[4] HKMA Circular on "Use of Artificial Intelligence for Monitoring Suspicious Activities" issued on 9 Sep 2024.

# Fighting Financial Crime: A.I. Adoption – What To Expect

## Point of View

A.I. adoption in financial crime functions will accelerate from controlled experimentation to enterprise-wide integration. Institutions that succeed will not simply deploy new tools; they will adopt new mindsets, operating models, and governance frameworks that position A.I. as a strategic capability rather than a technical enhancement.

### 01 From control to curiosity

The prevailing mindset within many financial crime environments has historically been one of compliance and control: minimising regulatory exposure, limiting model risk, and constraining operational variance. While these disciplines remain essential, the next phase of A.I. adoption requires a shift towards structured curiosity. Institutions will need to create safe environments to test generative artificial intelligence (GenA.I.), advanced analytics, and network intelligence – accepting controlled failure as part of innovation. Regulatory expectations are evolving in parallel, with supervisors increasingly distinguishing between unmanaged experimentation and well-governed innovation.

### 02 Holistic architecture over building blocks

Early A.I. initiatives have often been deployed as point solutions: TM enhancements, name screening optimisation, or case summarisation tools. Over the coming two years, competitive advantage will emerge from holistic architectures that integrate data strategy, model governance, workflow management, and human oversight. A.I. adoption should be designed end-to-end — from data ingestion to alert disposition — rather than bolted onto legacy rule engines. This requires aligning data standards, model validation, explainability, and operational processes into a unified framework. Building-block experimentation remains valuable, but it must feed into an enterprise A.I. blueprint.

### 03 Risk resides in the business

Financial crime risk does not originate in the compliance function; it originates in products, customers, channels, and geographies. A.I. programmes must therefore be co-owned by business lines and control functions. Over the next 24 months, we expect to see more financial crime subject matter experts embedded in frontline teams, joint model governance committees, and shared accountability for A.I. performance metrics. Effective adoption will depend on breaking down silos between first and second lines of defence. When business leaders understand how A.I. improves customer onboarding, reduces friction, and enhances detection quality, secure and scalable business growth can be better supported.

### 04 Leveraging public–private partnerships and the wider ecosystem

A.I. models are only as strong as the data and intelligence that inform them. Financial institutions are increasingly looking beyond internal data sets, leveraging PPPs, law enforcement typologies, financial technology collaborations, and cross sector intelligence sharing. The expansion of information sharing frameworks and privacy enhancing technologies will enable richer network analytics and cross-institutional risk signals.

## 05 Controlling the narrative to the Board

Board engagement will be decisive in securing funding, visibility, and sustained momentum. A.I. in fighting financial crime must be positioned not merely as supporting compliance obligations but as a strategic enabler of improved customer experience and growth. Reporting to the Board should balance opportunity and risk: clear articulation of measurable outcomes (fewer false positives, more efficient investigation, better detection), transparent risk mitigation strategies, and forward-looking investment roadmaps supporting capabilities that are delivering the most improved outcomes. Institutions that control this narrative will unlock sustained capital allocation and cross-enterprise prioritisation.

## 06 From rule-based to risk-based intelligence

The most significant shift will be away from static rule-led frameworks toward dynamic, risk-based intelligence. traditional rule-based engines, while transparent, are reactive and threshold-driven. A.I. enabled systems will increasingly leverage behavioral analytics, network link analysis, anomaly detection, and adaptive learning to prioritise risk. The future state is not the elimination of rules, but their integration within layered intelligence models. Supervisors expect explainable, well governed A.I. that demonstrably improves effectiveness rather than merely efficiency.

# Call to Action

The next 24 months will define how transformative A.I. in financial crime can become in Hong Kong with many banks already driving innovation programmes forward to deliver material benefit and measurable outcomes. Institutions that continue to embrace curiosity, integrate holistically, collaborate across lines of defence, leverage ecosystem intelligence, and secure Board sponsorship will move beyond automation toward intelligence-led risk management.

In Hong Kong and in other jurisdictions, regulators are playing a key role in supporting and encouraging innovation to identify foundational and next-generation technology capabilities which will support strategic transformation to financial crime risk management.

Through this programme of supervisory support, the HKMA has formalised how it encourages innovation in identifying A.I. solutions to better manage the harm caused by financial crime. In addition to tracking progress of each bank's implementation plan, over the next 24 months, Boards and senior management are expected to demonstrate progress across these supervisory priorities, anchored to institutional plans.

## 01 Deliver measurable effectiveness, not experimental activity

Programmes must deliver measurable improvements in financial crime risk management outcomes. Boards should expect clear evidence that A.I. is:

- Improving detection while reducing unnecessary alerts
- Accelerating investigations without compromising judgement
- Making regulatory reporting clearer and more useful
- Prioritising genuinely higher-risk activity through intelligence-led models

The objective should not be automation for its own sake. It must demonstrate improvement in risk identification, prioritisation and escalation.

## 02 Embed cross-line governance and clear accountability

While more than half of financial institutions report A.I. deployment within risk functions, fewer than one third operate fully integrated model governance frameworks across lines of defence. Institutions should:

- Establish formal cross-line A.I. governance committees
- Define Board approved A.I. risk appetite statements
- Implement real-time monitoring for model drift, bias and explainability
- Conduct annual scenario testing aligned to enterprise test frameworks

## 03 Transition from rule-based monitoring to intelligence-led risk management

Static rule sets are increasingly insufficient against adaptive and networked financial crime threats. A.I. adoption should enable:

- Dynamic, behaviour-based risk scoring
- Integration of external intelligence and typologies
- Continuous model recalibration in response to emerging threats
- Transparent explainability to supervisors and internal audit

## 04 Collaboration across the AML/CFT ecosystem

In Hong Kong and in other jurisdictions, regulators are playing a key role in supporting and encouraging innovation to identify solutions to better manage the negative effects of financial crime.

- Hong Kong has made great progress in integrating PPPs into the AML/CFT eco-system. Technology, including A.I. is recognised as a key enabler that will allow PPPs to move from incremental level improvement to system level impact
- The behavioral shift will frame the next phase of AML/CFT transformation, enabling banks to foster collaborative learning and deliver sustained resilience

It is important to understand that the presence of A.I. tools is not enough. There needs to be sufficient governance, effectiveness and strategic oversight surrounding them. Board-level sponsorship, measurable outcomes, and cross-ecosystem intelligence integration will distinguish between compliant adoption and transformative risk management.

# 03 Internationally Aligned Perspectives

Internationally, efforts against financial crime are entering a structural phase in which technologies – including A.I. – are central to the debate around effectively implementing the risk-based approach which underpins the global standards on AML/CFT set by the Financial Action Task Force (FATF). The FATF has encouraged the use of innovative technology, which can improve effectiveness and proportionality in applying AML/CFT controls while also maintaining safeguards in the promotion of financial inclusion. The core challenge is no longer whether institutions have compliance frameworks in place to meet legal and regulatory requirements, but whether those frameworks can generate meaningful, network-level intelligence, with the objective of driving stronger outcomes against agreed priority threats in an increasingly complex environment.

Criminal networks operate across firms and borders. A.I. — particularly where it supports advanced analytics of shared data in Hong Kong's PPPs — enables detection that reflects this reality. It allows institutions to identify behavioural anomalies, hidden ownership links, and interconnected actors that would not be visible through rule-based monitoring based on individual entity datasets.

At the same time, global transparency reforms — including those relating to beneficial ownership — are producing vast datasets. A.I. is what makes these datasets operational. Without advanced analytics, transparency remains passive. With it, institutions can resolve entities, detect circular fund flows, and identify risk concentrations across sectors.

A.I. also offers a more proportionate alternative to blunt de-risking strategies. Instead of withdrawing from higher-risk categories wholesale, advanced modelling allows firms to differentiate genuine risk from contextual noise, supporting both financial integrity and inclusion.

However, the international consensus is clear: A.I. is not a substitute for governance. Its deployment must be in line with applicable laws and regulations, be explainable and lead to measurable effectiveness outcomes. The future of financial crime prevention will depend not simply on adopting A.I., but on embedding it within collaborative intelligence frameworks that strengthen trust as well as capability.

In this evolving landscape, effectiveness — not volume of reporting — will define success.

*PPPs have made great progress over recent years, but most PPP models around the world still rely primarily on regular meetings between individuals. The more advanced PPPs have shifted to leverage technology and "data fusion" to support advanced analytics of shared data. Technology, including A.I., is a key enabler that allows PPPs to move from incremental improvement to system level impact. We are seeing scam groups and organised crime use A.I. to attack financial institutions. We now need to see PPPs maximise the use of technology, data fusion and A.I. - under appropriate data protection safeguards - to improve detection of criminal networks, reduce duplication of compliance efforts, improve suspicious transaction report quality and deliver earlier, more impactful disruption of financial crime.*

***Nick Maxwell**, Head of the Future of Financial Intelligence Sharing (FFIS) Research Programme delivering "The new era of private sector collaboration to tackle economic crime" in the 2025 APAC Fighting Financial Crime Conference, hosted by the HKMA.*

# 04 Case Studies in A.I. Adoption in Fighting Financial Crime

The case studies shared in this paper are based on the adoption journeys of four banks in Hong Kong and illustrate innovation in three different areas to support their strategies to combat financial crime.

## 1 Bank A

**Dynamic Risk Assessment built holistically using external resources partnership**

## 2 Bank B

**Dynamic Risk Model built holistically using internal resources**

## 3 Bank C

**Mule monitoring and detection using a building block approach for Retail during client onboarding**

## 4 Bank D

**Mule monitoring and detection using a building block approach for Corporate during client onboarding**

# Case Study 1

# A Global Bank's Approach to Dynamic Risk Assessment (DRA)

## 1. Why it was put in place: re-engineering financial crime detection

Up to 2020, most global banks were operating large-scale, rules-based AML systems to identify suspicious activities which generated large volumes of false positives that required manual investigation, creating operational strain and limiting effectiveness. The challenge was two-fold:

- Escalating financial crime complexity, including networked laundering and rapid movement of funds
- Inefficiency in rules-based detection, which struggled to prioritise true risk

The bank's top leadership recognised that incremental tuning of rules would not be sufficient, so it pivoted towards an intelligence-led risk-based model aligned with regulatory expectation and best practice.

In partnership with a cloud provider, the bank co-developed a dynamic A.I.-driven platform designed to autonomously identify suspicious behavior patterns rather than relying solely on static rules.

The strategic objectives were to:

- Improve detection precision
- Reduce false positives
- Shorten investigative timelines
- Enhance identification of criminal networks

## 2. A summary of the construction and technical architecture

Every bank will have its own process of implementing such a programme. However, the core components are likely to include:

- A cloud-based A.I. platform that generates risk scores, uses transaction data, customer profiles, Know Your Customer (KYC) data, and prior suspicious activity that integrates into case-management systems for analyst review
- A Domain Training Layer, where the bank contributes historical financial crime data, customer and transactional datasets, alongside internal typologies and risk expertise
- Machine learning models that are trained to identify suspicious activity including the detection of patterns such as rapid fund movements and abnormal behavior changes
- Cloud-native deployment introducing A.I. design for AML purpose for secure implementation within the bank's own cloud tenant, addressing model governance and regulatory requirements
- Scale of operations to handle vast data volumes, which in this instance averaged monitoring over one billion transactions per month

## 3. Collaboration timeline: what worked and what did not

The bank had been exploring cloud analytics and machine learning use cases well before DRA's formal launch. By around 2021–2022, it partnered with the cloud provider to pilot A.I.-enhanced AML detection. The successful trial led to the formal launch of DRA as a product in 2023. As with all large-scale programmes and "first through the door" innovators, the bank faced its own unique set of challenges.

The solution worked in areas such as reducing false positives by 60%, increasing the detection of suspicious activities by 2 to 4 times, and faster execution of case investigation (c.50%); all coupled with strong model governance and secure cloud deployment.

Some of the challenges of such a large-scale transformation programme included transitioning from compliance-driven, rules-based frameworks to A.I.-first, risk-based models. It involved communicating with multiple regulators as to the value of the shift and demonstrating that risk detection would be at least as good as under the previous systems. Whilst this was daunting, it ultimately paved the way for a paradigm shift in how the rest of the regulated industry would come to view the inadequacies of relying on just rule-based monitoring.

Other challenges that were encountered and overcome included ensuring explainability and model validation in a highly regulated environment as well as cultural and operational adaptation within compliance, operations and business teams.

## 4. Measurable outcomes and what's next

The bank reported filing over 137,000 suspicious activity reports in 2025 globally, underscoring the scale of its financial crime operations. Strategically it has improved signal-to-noise ratios, enabled investigators to prioritise higher-risk cases, enhanced network-level detection and reduced operational compliance burden. It has also positioned the bank as a reference case for its A.I.-driven AML transformation.

The bank is currently expanding its deployment capability across more markets and business lines as well as investing further in cloud-based analytics and A.I. It is further strengthening its A.I. governance and validation structures and framework as well.

The dynamic risk model represents one of the most prominent large-scale A.I. deployments in financial crime compliance in Hong Kong and globally. It represents a move from rule-based compliance to predictive, intelligence-driven risk detection. It combines A.I. platforms for AML/CFT with domain expertise and data scale. Operationally it has delivered measurable performance gains, and it illustrates for the industry how A.I., when embedded within governance and regulatory frameworks, can materially transform risk management in global banking.

# Case Study 2

# A Regional Bank's Approach to Dynamic Risk Monitoring Model (DRM)

## 1. The strategic pivot

The bank became increasingly convinced that rule-based monitoring was not fit-for-purpose to handle complex and layered money laundering schemes which were becoming increasingly dominant. Instead of single suspicious transactions, the bank saw multi-step, multi-product cross-jurisdictional flow deliberately crafted to avoid detection.

These legacy rules-based systems struggled to see patterns that unfolded across accounts, products and time horizons. They also generated too many false positives, despite tuning efforts, creating operational strain and alert fatigue and increasing the risk of missing the signal that mattered.

Following the analysis of the response in one case that exposed how fragmented systems were, accelerating the change needed to shift from chasing alerts to understanding patterns, thereby moving from rule-based detection to risk-based intelligence.

That decision was not about technology or A.I. alone but a strategic pivot to reimagine how compliance, data and analytics could work together.

## 2. The holistic approach

This required being able to see the entire picture of customer behavior. It involved building integrated data-lakes, combining customer profiles, transactional activities, behavior patterns and external data into a single environment.

A dynamic risk-scoring model was introduced to continuously adjust risk ratings based on behaviors and context, not simply on static thresholds.

Network analytics played a role in allowing the identification of links between customers, counterparties and entities that would have otherwise gone unnoticed. This drove a shift from a siloed transaction-only perspective to an intelligence-led, behaviorally informed approach. A shift from "What is suspicious?" to "What is really happening?".

The core objective was to improve detection rather than just efficiency and centred on the predictive detection of bad customers by assigning scores. This differed from other banks' strategies, which use A.I. primarily to achieve efficiency.

## Dynamic Risk Monitoring Model

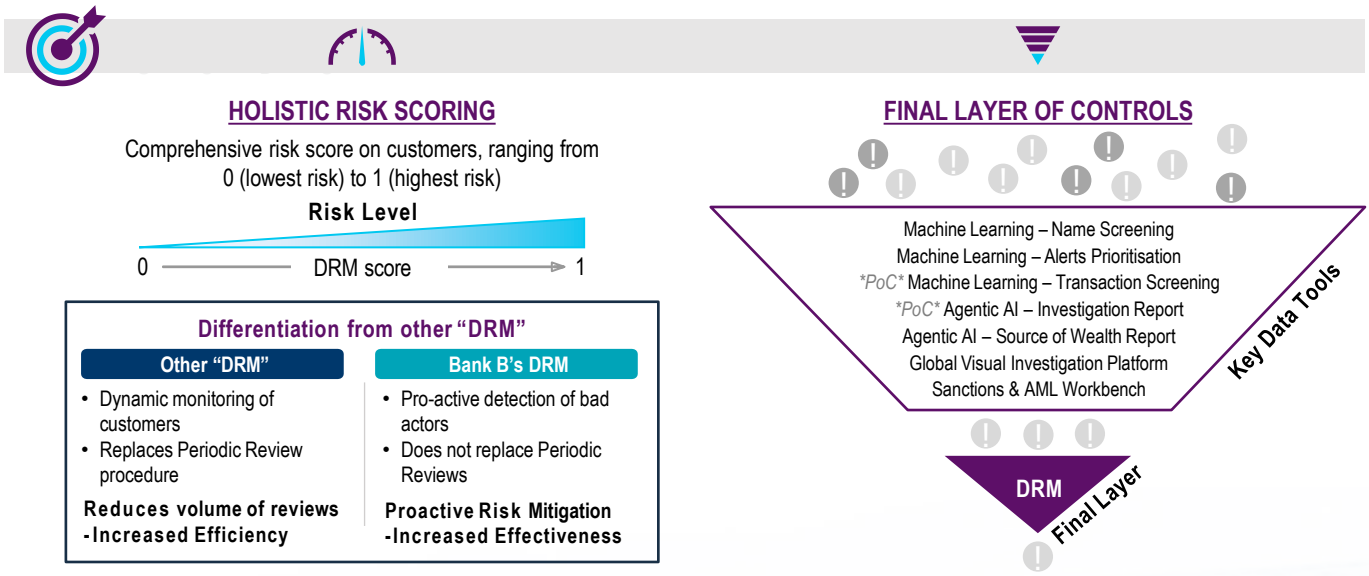

## 3. Deeper dive into the model

The system operates at a post-onboarding, layered process level and is part of the ongoing monitoring that is continuously scaling customer risk, rather than focusing on pre-boarding and focusing on activities.

It is not a binary choice (rules-based vs dynamic) but a two-layered process as it continues to use the transaction-monitoring system in parallel. The first layer of alert prioritisation is crucial to efficiently handling the volume generated by the underlying systems before applying the second predictive DRM layer.

This foundational alert prioritisation layer is machine learning based and uses transaction data to define a score and prioritise alerts. The advanced DRM layer then sits on top of the prioritisation model and connects with other models such as Network Risk Model to take into consideration the holistic activities of the customer.

This sequencing is vital, especially for smaller banks, as it adopts dynamic principles after it solves the foundational and structural issues of high alert volumes before tackling more advanced problem statements.

## 4. Navigating implementation challenges

Breaking the silos – getting old systems to communicate with new ones was the largest challenge, coupled with each business having its own standards of data, definitions and systems. This made integration complex and time-consuming requiring investments around data governance, standardisation and architecture.

- Shifting mindsets and culture – In addition to reform of systems, significant re-training of staff was required to enable them to operate as analysts and investigators rather than simply enforcing rules. The shift in mindset from control to curiosity was the key to successful adoption.

- Justification of the investment – the narrative to leadership needed to be re-framed from compliance spend to business value through better service for low-risk customers, improving efficiencies and building reputational resilience. It had to be a Board-level discussion and once the Board acknowledged that this wasn't just about catching criminals but to also enable better banking, support began to flow naturally.

- Impact on the customer – a significant benefit was the reduced friction for legitimate customers. With more accurate risk scoring, fewer false positives and unnecessary reviews. High-risk activities were being detected earlier and more precisely in the cycle especially prior to issues being escalated. In short; less noise, more focus and better service.

- Roadmap for A.I. deployment – the bank has put in place a roadmap for the ongoing innovation and development of A.I. and Data Analytics (DA). The roadmap sets out its plans for new A.I. and DA tools to be added to further strengthen its compliance risk management.

## Looking Ahead: The bank's A.I. / DA roadmap for 2026/2027

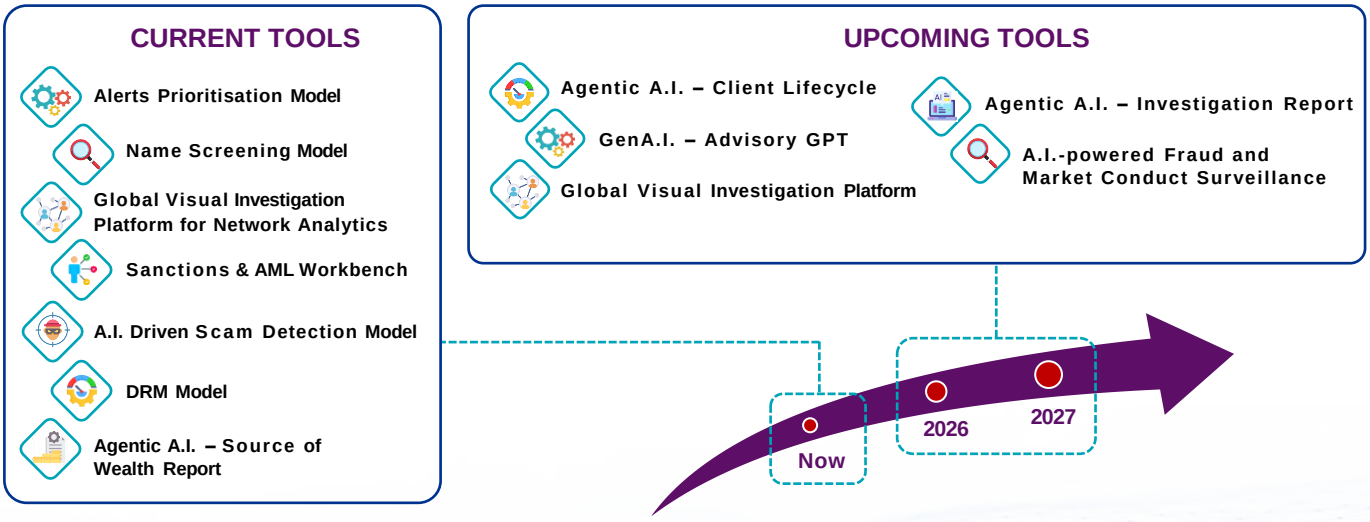

# Case Study 3

# Face Watch List in a Digital Bank

## 1. Background and context

A digital bank had to be more innovative in its pursuit of a target A.I solution to make optimum use of limited resources. Seeing a rise in mule activities industry-wide and individuals opening accounts for illicit fund movement, the bank recognised the needs to overcome these challenges by balancing seamless digital onboarding with stronger fraud detection.

Using a building-block approach, they started small, demonstrated value and prepared to scale thereafter. By leveraging the HKMA GenA.I. Sandbox for the initial development and testing of the proof of concept (PoC), the bank was able to build and evolved a number of practical use cases, through multiple iterations and an agile development approach, with limited internal resources (including data engineers and data scientists). By deploying a test-and-learn approach to develop small modular A.I. models, the bank went beyond rule-based transaction to holistic monitoring of customer behaviors.

The use case involved detecting third parties attempting to impersonate customers by identifying irregularities in the backdrops of their pictures submitted as part of the remote onboarding process.

The outcomes of the case were later used to help law enforcement further investigate and disrupt criminal syndicates using this modus operandi.

## 2. Why adoption was successful

**A. The innovation was practical and executable with strong results:**

- It has seen mule account detection increase by 30%
- Screening times have improved to within seconds.
- By identifying malicious actors at the initial stage, the solution has reduced end-to-end monitoring costs delivering benefits that cascade through the entire AML/CFT risk monitoring process

**B. It was not difficult to scale the A.I. used:**

- Implementation was fast. A PoC was initiated in June 2025 and the first version deployed into production within 3 months
- This included training the machine learning models to learn suspicious background scenarios

- The agile development approach, streamlined governance and industry infrastructure (Sandbox) enabled the team to secure executive sponsorship, drive forward the development in multiple sprints and accelerate delivery

**C. It involved more than financial crime professionals in the second line:**

- Incorporated skillsets across the bank and not solely reliant on financial crime compliance staff

- Demonstrated cross-function collaboration and agility between risk stewards, business lines, technology and data architecture as well as management

## 3. How it began

It was initiated by several team members analysing an actual mule-account case. It involved dissecting the case including the account and customer profile, transaction patterns as well as digital footprints including facial images. Several scenarios were identified that had common attributes.

## 4. How does it work?

It uncovered the same unknown individual appearing in facial recognition records of multiple fraudulent account holders. It identified that a controller was possibly managing multiple accounts simultaneously, and the images of the suspected controller would be added to the bank's internal watchlist.

It was also noticed that there were common backgrounds in some of the images. These included suspicious backgrounds involving alleged gambling or drug dens, illegal casinos and other unusual locations.

In addition, specific objects appeared in the same background of different account holders. In one instance, syndicated activity was indicated in over 200 accounts by a similar moving vehicle, an apartment with similar furniture or customers holding on to multiple sets of handphones while logging on to the bank's accounts.

## 5. Steps to using building blocks

With smaller agile banks, problem statements are often broken down into building blocks or bite-sized chunks, while identifying the key components needed to achieve the desired output.

Components can often include:

- Image database set up to store all onboarding and step up authentication images as well as an adjacent database to house all suspicious images. The next step is labelling them into different categories to train the A.I. models

- Getting the right A.I. model to address the problem statements that utilise both real-time and batch processing screening for fast, accurate and nimble deployment

- Continuous re-calibration including training and validation of the A.I. models to further enhance accuracy and to build out the library of suspicious images

# 6. Hurdles and Outcomes

The adoption of new technology brought some challenges. The key to approaching challenges is to break them down into smaller blocks and tackle them one at a time. In this use case there were two main hurdles:

### 1. Matching speed

The initial PoC took five to seven days to run a match against a million images; this evolved to between 1 to 3 seconds with results delivered in real-time within the onboarding process. Certain scenarios such as identification of suspicious backgrounds in images take slightly longer.

### 2. Training and accuracy

Initial results showed between 65% to 70% accuracy, which was later enhanced after studying the failure points. It was fixed by labelling the suspicious part of the image, which drastically improved the learning rate of the A.I. model.

One of the most measurable adoption outcomes involved reducing TM alerts by 25% largely because bad actors could be stopped at the door, so the bank did not waste precious time and resources detecting them after onboarding.

## The Human Element

The human-in-the-loop factor is critical for implementation and outcomes. Subject matter experts are constantly needed to teach and train the A.I. model on how to identify suspicious images, including translating those requirements for the technology professionals involved in the training and application. Last and most importantly, support from senior management was critical to deliver the outcome the business required.

## Leveraging Public-Sector Support

Seeking support from the public sector (i.e. through the HKMA GenA.I. Sandbox) to take an initial idea from baseline to “Minimum Viable Product” was invaluable in helping the bank generate the required internal support and funding to drive and scale up the A.I. solution.

# Case Study 4

## Detecting and Disrupting Money Mules throughout the Client Life Cycle

A large bank became aware of increasing payment fraud volumes within its Hong Kong commercial banking portfolio, where shell companies were used to facilitate transactions and flow funds from illicit scams through the bank's accounts.

The rise in the number of mule accounts and the increasing sophistication of mule networks left the bank with limited visibility into the sources of incoming funds and the risks present at account origination. Although the bank had strong payment fraud detection systems, it needed to strengthen its ability to detect shell companies and mule accounts used for money laundering, as traditional rules-based approaches were less effective in identifying and freezing these funds.

The business (First Line) implemented a new "Two Masters Model", by leveraging prevention tools to freeze payments in near real-time, with investigation teams providing strong support and collaboration.

The shared objective was to shift from reactive investigations of money mule transactions and shell companies to a proactive mode that would:

- Prevent the onboarding of money mules based on various DA (e.g. device biometrics, addresses etc.)
- Identify money mules based on transactional and behavioural analytics
- Use the banks fraud risk controls to freeze funds in near real-time

### Two New Detection Tools – Application Engine and In-Life Monitoring

To respond to the problem, the bank developed two new detection capabilities to monitor customers at onboarding (account application) and throughout the customer life cycle (In-Life).

**Prevent: Application Engine (Onboarding)**

An "Application Engine" was built with real-time application programming interface (API) to enable onboarding assessment and controls leveraging data sources including:

- Device intelligence
- Cross-application analytics (against historic onboarding applications)
- Business information (KYC data – internal and external).

The engine scored a customer at onboarding against a specific set of risk indicators and behavioral identifiers with the application receiving:

- Pass / Reject / Review (escalate for human intervention and assessment)

Review cases were used as part of the tuning and feedback loop to train and strengthen the solution.

By iteratively testing and tuning the analytics models, the bank trained its application engine to reject suspected money mules at onboarding, helping to reduce exposure earlier in the client life cycle.

**Disrupt: In-life Monitoring (Customer Life Cycle)**

A second "In-Life Monitoring" capability was established to monitor customer account behavior for suspicious activity across the whole client portfolio combining:

- Traditional human-explainable risk indicators
- A new machine learning prediction model

The solution leverages data (KYC, transactional and external) and machine learning models to:

- Assign a risk score to the customer, or
- Evaluate against predefined typologies

Where suspicious behaviour is identified, the bank can intervene quickly through targeted actions, for example:

- Initiate a Customer Due Diligence (CDD) review
- Raise an Unusual Activity Report (UAR)
- Implement payment controls on the account

### Delivery: Approach (Agile, Co-Creation, Pace)

Both capabilities were developed using an agile execution approach led by the business (First Line) with “co-creation” workshops incorporating representatives from multiple areas i.e.:

- AML / Fraud
- DA
- Technology

Critical to this was the ability to incorporate knowledge, experience and learnings from different functions and subject matter experts to build strong risk models.

In addition, to ensure pace of delivery and continued support from executive sponsors (budget owners), the business implemented an iterative development approach to:

- Drive innovation
- Test prototypes
- Fail fast
- Deliver quick wins

## A Solution Pioneered Through Data & Analytics

The bank adopted a transformation approach to the design and orchestration of the innovative solutions.

A cross-functional team led by the business incorporating colleagues from risk functions, data office, payment operations and Information Technology were brought together to collaboratively evaluate how new technology and risk-detection techniques could be used to solve the money mule detection challenge.

Fundamental to this was the support of DA colleagues and data scientists to:

- Identify the data sets required to develop the solution
- Establish thresholds to deliver appropriate risk scoring across the different data sets (transactional, behavioral, network, KYC and risk indicators)
- Build out the network analytics capability and define the machine learning models
- Ensure adequate levels of data quality

Business and risk colleagues were engaged to validate the model outcomes for explainability based on subject matter experts' knowledge and judgement. In addition, the models were subject to formal model governance.

A key takeaway was that while data quality is typically critical to the success of traditional rules-based detection, the machine learning models improved risk detection even when some data attributes were missing or data quality was suboptimal. The machine learning model rated 42% of customers as high risk whereas the rules-based detection model had previously rated these customers as medium risk.

Of these customers, 70% were escalated for investigation and 40% of these customers were exited for financial crime concerns. By leveraging:

- Network analytics for relationships (mule networks)
- Machine learning for detection
- GenA.I. for explainability (to summarise money laundering signals, contextualise findings and augment external research)

The team were able to design, test and tune an effective detection model for money-mule activity.

## Adoption by other markets

The early success in detecting and disrupting money mule activity in Hong Kong, together with the continued evolution of the solution, generated interest from other jurisdictions and the wider bank helping the team to sustain ongoing funding and executive support.

While initially designed to solve a growing financial crime issue in Hong Kong, the solution has subsequently been rolled out in several other markets with commercial banking portfolios. Its application is also being evaluated in other business lines e.g. retail banking.

The data-driven and collaborative approach has also been presented as a model of success internally to drive the development of other financial crime risk management projects within the bank; and the techniques established are being leveraged by the wider business.

By moving beyond a purely rules-based approach to aggregated risk dimensions, the solution enabled the bank to target shell company and money mule behaviour more effectively across the client life cycle.

## The Model

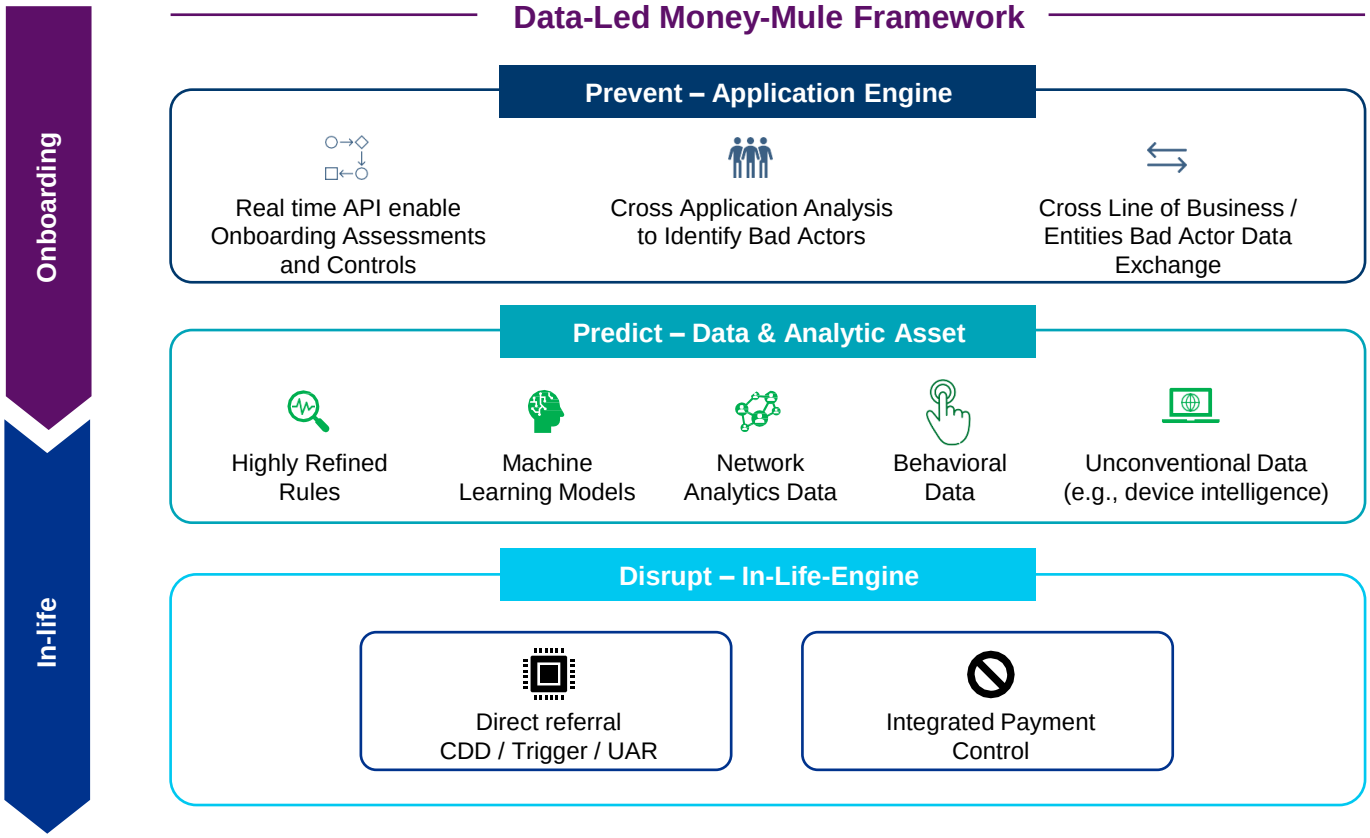

# 05 Where to From Here?

## From Automation to Symbiosis: The Emergence of Agentic A.I.

In any discussion of the strategic transformation of financial crime risk management, the adoption of agentic A.I. is already front and centre; it is no longer conceptual or theoretical.

However, deployment remains largely efficiency oriented as A.I. is being used to summarise cases, triage alerts, draft narratives and optimise workflows. It is accelerating processes, but it is not yet materially reshaping detection strategies.

A majority of Hong Kong banks reported active A.I. deployment in risk or compliance functions, with a few describing those capabilities as decision-supporting beyond workflow optimisation. Fewer still currently deploy autonomous analytical agents within core detection environments. Most programmes remain tightly constrained to productivity enhancements rather than intelligence-led amplification.

This begs a strategic question for the industry:

If agentic A.I. can analyse network relationships, test hypotheses, simulate behavioral scenarios and adapt to emerging typologies, why is it not yet being deployed as a supervised co-pilot in financial crime detection?

Financial crime threats are increasingly adaptive, networked, and technologically enabled. Agentic A.I. introduces the potential for:

- Dynamic hypotheses generation across linked accounts and counterparties
- Continuous and tireless anomaly detection beyond pre-defined rule parameters
- Real-time investigative assistance during complex case escalations
- Proactive identification of emerging typologies

The near future is not autonomous decision-making without oversight. It is supervised augmentation with A.I. operating as a co-pilot alongside investigators and risk officers.

In this model, the human would define the risk, the risk appetite, the judgment thresholds and escalation pathways. The agentic A.I. would continuously interrogate data, unveil unseen connections and challenge assumptions. All surrounded by governance frameworks to ensure transparency, explainability and auditability.

This represents a shift from A.I. as a productivity tool to A.I. as a strategic intelligence partner.

The use of agentic A.I. will be a future focus area of the programme.

# 06 Glossary of Terms

| | |
|---|---|
| **A.I.** | Artificial Intelligence |
| **AML** | Anti-Money Laundering |
| **APAC** | Asia-Pacific |
| **API** | Application Programming Interface |
| **CDD** | Customer Due Diligence |
| **CFT** | Counter-Financing of Terrorism |
| **DA** | Data Analytics |
| **DRA** | Dynamic Risk Assessment |
| **DRM** | Dynamic Risk Monitoring Model |
| **FATF** | Financial Action Task Force |
| **FFIS** | Future of Financial Intelligence Sharing |
| **FINEST** | Financial Intelligence Evaluation Sharing Tool |
| **FMLIT** | Fraud and Money Laundering Intelligence Taskforce |
| **GenA.I.** | Generative Artificial Intelligence |
| **HKAB** | Hong Kong Association of Banks |
| **HKMA** | Hong Kong Monetary Authority |
| **KYC** | Know Your Customer |
| **PoC** | Proof of Concept |
| **PPP** | Public-Private Partnership |
| **TM** | Transaction Monitoring |
| **UAR** | Unusual Activity Report |