september 2025

REGCOUNSEL FINANCIAL SERVICES WHITE PAPER SERIES #5

Using Agentic AI in Transaction Monitoring
Part I: EU Regulatory Framework, Design and Governance Considerations

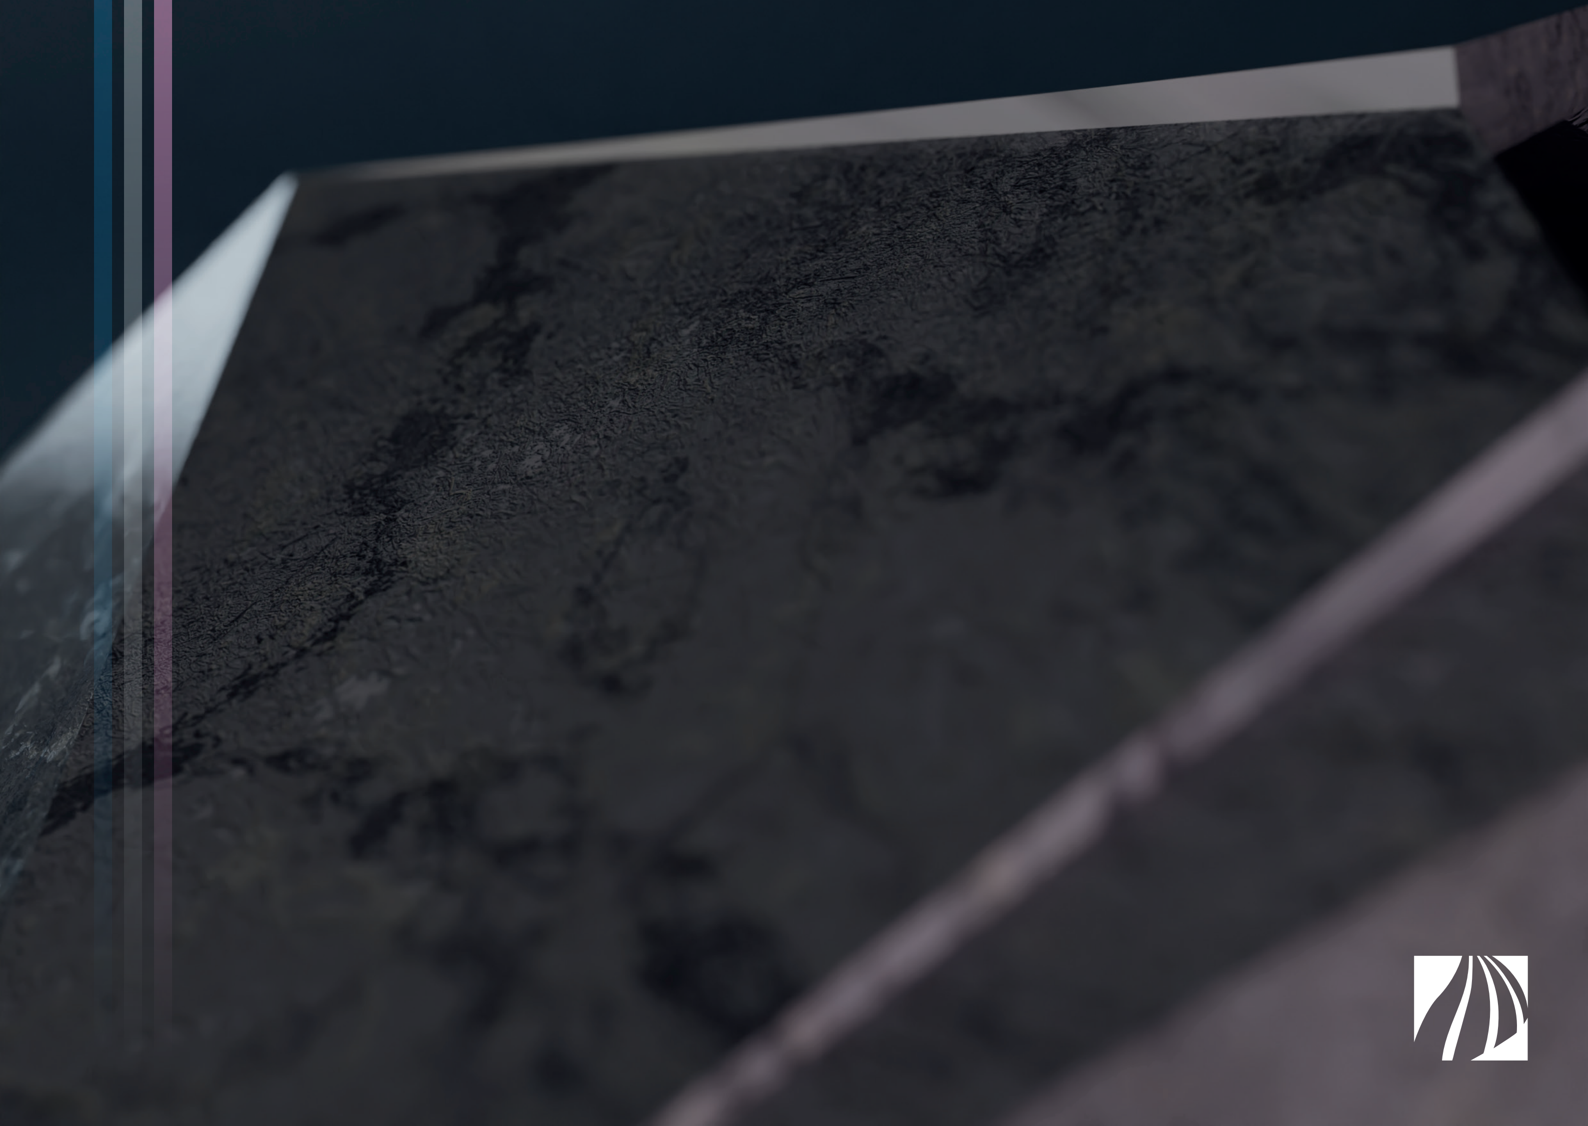

ABOUT THE AUTHORS

Mr. dr. drs. Jan Jans is a Dutch qualified lawyer specialized in financial regulatory law. He has a strong focus on the regulatory framework covering investment firms, banks, insurance companies, payment service providers, crypto-asset service providers and investment funds. Prior to joining Regulatory Counsel | Financial Services as Partner in 2022, Jan worked as a financial regulatory lawyer at various (international) law firms in the Netherlands.

Jan graduated from Nijmegen University in 2008. In addition to law, he studied financial economics and business administration. He completed the post-academic education Securities Law at the Grotius Academy in 2012. Jan publishes regularly in national law journals and successfully defended his PhD thesis on the level playing field between banks and non-banks in the European market for electronic payments in 2023.

Mr. drs. Jan-Jouke van der Meer is a Dutch qualified lawyer specialized in financial regulatory law. He advises clients on Dutch and EU financial regulatory laws (including MiFID, AIFMD, MiCAR, PSD2 and DORA), licensing issues, corporate governance (including integrity and suitability assessments), compliance and integrity-related issues (including AML and sanctions) and dealing with financial regulators. Prior to joining Regulatory Counsel | Financial Services as Managing Associate, Jan-Jouke worked as a financial regulatory lawyer at an international law firm in the Netherlands.

Jan-Jouke obtained a master's degree (LLM) in Corporate Law from VU Amsterdam in 2018, and holds a master's degree (MSc) in Financial Economics from Erasmus University Rotterdam (2018). He has also published various articles in the area of financial regulatory law.

Mr. Maarten Mol-Huging is a Dutch qualified lawyer specialized in financial regulatory law. He advises on (prudential) regulatory issues with regards to banks, investment firms, insurers, crypto-asset service providers and other financial sector entities.

Maarten graduated cum laude at the VU Amsterdam in International Business Law in 2022 on his thesis titled: The Biggest Challenge yet for EU Mortgage Lending: The Onset of Artificial Intelligence Tooling. In 2023, he furthermore completed the prudential regulation-oriented EBI Master in EU Banking and Financial Regulation in Frankfurt.

Recofise is a boutique law firm located in Amsterdam and specialises in financial regulation. Our clients consist of financial undertakings (e.g., banks, investment firms and funds, insurance companies, electronic money institutions, payment institutions and crypto-asset service providers) that are active on the Dutch and international financial markets. Recofise has extensive experience in advising and supporting clients with complex legal issues regarding prudential requirements, compliance, risk management, supervisory reporting and supervisory audits.

EXECUTIVE SUMMARY

European financial institutions face an evolving regulatory landscape where agentic AI – autonomous AI systems capable of perceiving, deciding, executing, and learning without continuous human intervention – presents both compliance opportunities and novel supervisory challenges in transaction monitoring.

Regulatory Framework and Compliance Position

Under current EU financial regulation, agentic AI deployment for AML/CTF transaction monitoring operates within existing principle-based frameworks. AMLD4 and the Dutch Wwft impose ongoing transaction monitoring obligations without prescribing specific methodologies, creating regulatory opportunities for AI innovation. However, the forthcoming AMLR introduces critical constraints: automated decision-making systems may only process CDD-obtained data, and decisions regarding customer onboarding, offboarding, or enhanced due diligence require "meaningful human intervention".

The regulatory tension is clear: while supervisors including DNB and EBA demonstrate openness to AI innovation in transaction monitoring, the autonomous nature of agentic AI poses a challenge within existing regulatory frameworks that assume direct human oversight.

Key Regulatory Compliance Issues

Human Intervention Requirements: The AMLR's meaningful human intervention standard remains undefined, creating regulatory uncertainty around the permissible degree of AI autonomy. Institutions must structure systems to ensure human oversight for high-impact decisions while leveraging AI's autonomous capabilities for routine processing.

Data Processing Limitations: Agentic AI systems face restrictions under the AMLR, limiting automated processing to CDD-obtained data. This constrains the system's analytical scope while requiring institutions to maintain clear data governance boundaries.

Explainability and Auditability Obligations: Transaction monitoring decisions must remain explainable to customers and auditable for supervisory review. Agentic AI's multi-step reasoning processes through chain-of-thought technology may satisfy explainability requirements, but institutions must ensure decision logic remains transparent and documentable.

Governance and Supervisory Expectations

This White Paper proposes a dedicated AML AI Governance Taskforce operating within the second line of defence, combining AI technical expertise with AML/CTF regulatory knowledge. This structure addresses supervisory expectations for:

- Clear accountability allocation to the management body;
- Continuous monitoring and control capabilities matching AI system autonomy;
- Rapid intervention protocols when AI exceeds predefined guardrails; and
- Comprehensive audit trails supporting regulatory examinations.

Regulatory Strategy Recommendations

Financial institutions should adopt a structured compliance approach prioritizing regulatory alignment over technological capability. Key regulatory preparedness measures include:

I **Governance Infrastructure**: Establish specialized oversight combining second-line independence with technical AI expertise;

II **Documentation Protocols**: Implement comprehensive audit trail systems supporting supervisory examinations and customer explanation rights;

III **Human Override Mechanisms**: Design systems ensuring meaningful human intervention for all material customer-impacting decisions; and

IV **Supervisory Engagement**: Maintain proactive dialogue with regulators regarding AI deployment within existing compliance frameworks.

Forward Regulatory Outlook

The regulatory environment for agentic AI remains developmental, with supervisors balancing innovation encouragement against consumer protection and financial stability concerns. Institutions must prepare for evolving supervisory expectations while building systems capable of adapting to regulatory clarifications, particularly around meaningful human intervention standards, fairness expectations and data processing requirements.

Part II of this White Paper Series will address the specific risk management frameworks required for responsible agentic AI deployment, including operational resilience, model governance, and data quality obligations under, *inter alia*, DORA and sectoral prudential requirements.

TABLE OF CONTENTS

1. INTRODUCTION

1.INTRODUCTION: WHAT IS AGENTIC AI?

Transaction monitoring for anti-money laundering and counter terrorism financing ("**AML/CTF**") purposes has long seen the application of artificial intelligence ("**AI**"). Up until now, AI used in transaction monitoring has largely been in the form of specifically trained AI models rather than the more generalist large language models ("**LLMs**"). Recently, advances in AI technology and IT capabilities have given rise to the powerful potential of agentic AI.

Definition of AI

Though AI in and of itself is a rather broad notion referring to a large range of technologies, the concept of an 'AI System' has been defined in the EU AI Act.[1] This definition is as follows:

"AI system' means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.''[2]

Though still in its infancy, agentic AI is poised to fundamentally alter the way in which analytical processes such as transaction monitoring are conducted.

[1] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (2024).

[2] Art. 3(1)(1) and Recital (12) EU AI Act.

1.1 What is agentic AI?

Agentic AI refers to AI systems that – very briefly put – pursue goals, make decisions, and adapt based on feedback from their environments **without human intervention**. In other words, agentic AI systems act autonomously of humans once they have been set up. This distinguishes agentic AI from, for example, generative AI ("**Gen AI**"), which requires human input to generate an output. In the case of Gen AI, once the output is produced, the AI system's role effectively ends, and it is up to the human user to act upon the result.

At a high level, agentic AI follows a four-step approach for solving problems:

I **Perceive** – Agentic AI can interact with a particular environment and gather data to adjust in real-time, constantly analysing its surroundings within that environment;

II **Decide** – Based on pre-defined (human) plans and objectives agentic AI can assess situations and decide without, or with minimal, human input;

III **Execute** – Agentic AI is uniquely able to autonomously operate on its own once it has reached a decision. For instance, if a transaction is deemed unusual, an agentic AI system can autonomously follow up on the transaction and may even conduct a review of the respective customer – all with minimal human intervention; and

IV **Learn** – After following-up on its decision, agentic AI autonomously learns from the completed process (feedback). In the context of transaction monitoring, this may entail the agentic AI system to identify new fraud typologies, or it recalibrating one of the indicators for unusual transactions.

The underlying technology of agentic AI is, despite its autonomy, not dissimilar to that of Gen AI. Both are (generally) based on LLMs, which allow these systems to process vast data inputs and reason through what the most logical output should be (at least according to the model). For example, agentic AI may use an LLM to interrogate itself regarding the various steps required in a complex process – a method known as *chain-of-thought reasoning* ("**CoT reasoning**").

An example of CoT reasoning by an agentic AI model, in the context of transaction monitoring for a transaction X, could (in simplified form) look like this:

Problem: Is transaction X unusual?

a. What is considered unusual?
b. What are the characteristics of transaction X?
c. Do these characteristics match what is considered unusual?
d. If the transaction is unusual, what should happen next?

Therefore, agentic AI is – at least theoretically – particularly suitable for *multi-step processes* such as transaction monitoring and analysis. In this domain, agentic AI can leverage output generated by predefined rule-based AI systems and then dynamically assess the veracity of a detected hit, escalate decisions, and/or adjust its detection strategies accordingly. Hence, agentic AI may provide opportunities to continuously reprioritise alerts, reconfigure detection parameters, or autonomously flag new typologies of suspicious behaviour.

Box: Agentic AI vs. AI agents

Agentic AI is not the same as an **AI agent**. Fundamentally, AI agents are the technical components within the broader conceptual framework of agentic AI – that is, a system capable of solving complex problems without human intervention. Accordingly, a single agentic AI system may consist of multiple AI agents. Each AI agent is designed to handle specific tasks and processes with a degree of autonomy.

1.2 Impact agentic AI

The increase in AI autonomy that agentic AI brings will logically be accompanied by a decrease in direct human involvement in transaction monitoring; at least for the part covered by the agentic AI system. This is the very essence of agentic AI use, reducing the need for human resources in labour intensive processes such as transaction monitoring. Agentic AI systems are designed not just to execute tasks but to operate with a level of independence, learning and adapting over time.

As such, the use of agentic AI raises fundamental questions about accountability, supervision, and compliance. Specifically in a heavily regulated industry such as the EU financial sector[3], ensuring compliance whilst leveraging the capabilities of agentic AI is key to not only long-term productivity growth and commercial success but also to effective AML/CTF compliance.[4]

1.3 Purpose and structure of this paper

This paper aims to support EU financial institutions in understanding the regulatory requirements and supervisory expectations surrounding the deployment of agentic AI for transaction monitoring purposes.

First, we take a deeper look at the drivers of (agentic) AI adoption in transaction monitoring. Second, we map the various regulatory requirements under EU financial regulation applicable to agentic AI use. Third, we outline a theoretical design for an agentic AI system used in transaction monitoring. Fourth, we take a look at the governance regulation and organisational embedding of agentic AI. The White Paper will be concluded with some key recommendations for future application.

[3] E.g., pursuant to the Dutch Anti-money Laundering and the Counter Terrorism Financing Act (*Wet ter voorkoming witwassen en financieren terrorisme*, "**Wwft**") or the Anti-Money Laundering Regulation ("**AMLR**") (Regulation (EU) 2024/1624 of the European Parliament and of the Council of 31 May 2024 on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing (2024)).

[4] For instance, art. 2a(2) Wwft requires the taking of adequate measures to combat increased AML/CTF risks related to new technologies. AI is already boosting fraud significantly; whereby financial institutions have to meet the challenge and take adequate measures themselves.

2. DRIVERS OF (AGENTIC) AI ADOPTION IN TRANSACTION MONITORING

2. DRIVERS OF (AGENTIC) AI ADOPTION IN TRANSACTION MONITORING

2.1 Evolution of AI in transaction monitoring

Financial institutions have used AI in transaction monitoring for some time. Its evolution has broadly followed three (**3**) phases:

I **First**, comparatively simple rule-based systems were deployed to scan large volumes of transactions against fixed business rules;

II **Second**, institutions increasingly adopted machine learning systems that could dynamically develop and refine transaction profiles for individual customers or groups of customers; and

III **Third**, the sector is now entering a new phase, marked by the early adoption of agentic AI in transaction monitoring processes.

2.2 Drivers

This transition is not solely driven by technological innovation but is also the product of a convergence of **economic**, **operational** and **regulatory** pressures. The following drivers underpin the gradual shift towards more autonomous and dynamic AI models in transaction monitoring.

Escalating transaction volumes and complexity – The volume and velocity of payment transactions continue to rise across traditional and emerging financial channels (e.g., due to the shortening of settlement periods from two (**2**) days to ten (**10**) seconds). Real-time processing and cross-border flows introduce additional layers of complexity, making manual or rule-based monitoring increasingly untenable. Agentic AI offers the potential to autonomously triage, prioritise, and act on alerts, reducing human bottlenecks.

Evolving typologies of financial crime – The same technological advances that empower financial institutions also equip malicious actors. Illicit financial activities – from fraud to money laundering – are becoming more sophisticated and adaptive. The use of AI by fraudsters – as recently highlighted by the European Banking Authority ("**EBA**")[5] – for purposes such as social engineering,

[5] EBA, Opinion and Report on money laundering and terrorist financing risks affecting the EU's financial sector, (2025), par. 3.10.

fake identification (documentation), or automated money mules raises the bar for effective transaction monitoring. This adversarial dynamic may further accelerate the necessary adoption of more capable, agentic AI solutions.[6] Agentic AI, through continuous learning and autonomous reasoning, may help identify novel patterns and typologies faster and more reliably. Moreover, AI-enhanced fraud may comprise patterns undetectable to humans, whereby agentic AI may be particularly suited to stay ahead of criminal developments.

Resource-intensive screening processes – Traditional transaction monitoring processes remain labour-intensive, particularly in respect of the alert handling and case review procedures. Agentic AI systems can automate multi-step decision chains – for example, validating a (false) hit, collecting contextual data, and drafting a first-line assessment – potentially reducing the burden on compliance teams. Financial institutions are expected to do more with less, and agentic AI adoption supports that imperative.

Shortcomings of static rule-based systems – While rule-based systems are reliable and transparent, they lack the adaptability required to handle edge cases and novel behaviour. They are also prone to producing large volumes of false positives if incorrectly calibrated. Agentic AI can contextualise individual transactions and weigh relevant factors dynamically, potentially reducing noise while increasing detection accuracy.

Advancements in machine learning and natural language processing – Recent developments in LLMs and multi-agent architectures have made it feasible to build agentic AI systems capable of complex reasoning, self-prompting (i.e., CoT), and task coordination. These developments serve as a potential technological enabler for use cases that were previously unfeasible or suboptimal, such as autonomous alert resolution or ongoing rule optimisation.

Increased regulatory expectations and reduced tolerance for failure – The regulatory landscape for AML/CTF continues to evolve, placing a heavier burden on institutions. The deployment of more intelligent AI systems may offer institutions the tools necessary to meet these heightened regulatory standards without having to expand their compliance teams exponentially.

[6] Seeing as the transaction monitoring mechanism must remain effective and appropriate, see EBA, *ML/TF Risk Factor Guidelines* (2023), par. 4.72.

Supervisory openness to innovation – Supervisors in the EU, including EBA and the Dutch Central Bank (*De Nederlandsche Bank*, "**DNB**")[7], have shown openness to innovation – if somewhat cautiously – particularly where AI enhances risk-based monitoring.[8] While conscious about explainability, accountability and control issues, supervisors increasingly appear to allot for proportional experimentation with AI systems under existing regulatory frameworks.[9]

[7] For instance, DNB mentions the use of AI in transaction monitoring as a Good Practice, see DNB, *Q&A and Good Practices Wwft*, (2024), p. 55-56.

[8] See, e.g., EBA, *Report on the use of AML/CTF SupTech tools,* (2025).

[9] As long as in line with supervisory expectations; see, e.g., DNB & AFM, The impact of AI on the financial sector and supervision, (2024).

3.

LEGAL AND REGULATORY FRAMEWORK IN THE EU

3. LEGAL AND REGULATORY FRAMEWORK IN THE EU

In principle there are no explicit regulatory requirements that target the use of *agentic* AI. Nevertheless, the use of agentic AI – and AI in the broadest sense – for transaction monitoring purposes is covered by a variety of existing financial regulatory instruments. The applicable regulations can generally be divided into three broad categories:

I **Purpose** – Covering the specific use of agentic AI for transaction monitoring purposes;

II **Governance** – Covering the prudent embedding of agentic AI in the organisation and governance model of financial institutions (see paragraph 5); and

III **Risk management** – Covering the specific risks related to AI technology and related data usage underlying agentic AI systems used for transaction monitoring. This topic, where in addition to EU financial regulation the EU AI Act plays a pivotal role, will be covered in Part II of the White Paper Series on agentic AI.

3.1 Purpose regulation – Transaction monitoring

3.1.1 Current: AMLD4, Wwft and DPR

Most financial institutions are subject to customer due diligence ("**CDD**") requirements under applicable AML/CTF regulations. A core component of CDD is the obligation to perform ongoing transaction monitoring.

Currently, the updated fourth Anti-Money Laundering Directive ("**AMLD4**")[10] sets out the relevant EU rules. Specifically, the following article contains the core transaction monitoring obligation:[11]

[10] Directive (EU) 2015/849 of the European Parliament and of the Council of 20 May 2015 on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing (2015) (OJ L141/73).

[11] Further modified by art. 15(3) AMLD4 in light of simplified due diligence ("**SDD**"), art. 18(2) AMLD4 in light of enhanced due diligence ("**EDD**"), art. 18a(1)(f) AMLD4 for high-risk third countries, art. 19a(1)(c) AMLD4 in respect of transactions involving crypto-assets sent or received from self-hosted wallets and art. 20(b)(iii) in respect of politically exposed persons ("**PEPs**").

Art. 13(1)(d) AMLD4

'[CDD] measures shall comprise: [...] (d) conducting ongoing monitoring of the business relationship including scrutiny of transactions undertaken throughout the course of that relationship to ensure that the transactions being conducted are consistent with the obliged entity's knowledge of the customer, the business and risk profile, including where necessary the source of funds and ensuring that the documents, data or information held are kept up-to-date.'[12]

In the Netherlands, AMLD4 has largely been implemented into the **Wwft**. The relevant provision relating to transaction monitoring reads as follows (freely translated):

Art. 3(2)(d) Wwft

'The [CDD] enables the obliged entity to: [...] (d) to conduct ongoing monitoring of the business relationship and the transactions carried out during the course of that relationship, in order to ensure that they are consistent with the institution's knowledge of the client and their risk profile, including, where necessary, an examination of the source of the funds used in the business relationship or the transaction.'

A further provision seeing to transaction monitoring closely related to CDD, is included in the Dutch Decree Prudential Rules FSA (*Besluit prudentiële regelingen Wft*, "**DPR**"):

Art. 14 DPR

'The financial undertaking [...], or, as the case may be, the branch, has procedures and measures in place for the analysis of client data, including in relation to the products or services obtained by the client, as well as for the detection of unusual transaction patterns. Based on these procedures and measures, the financial undertaking also determines the risks associated with certain clients, products, or services for the sound and ethical conduct of its business.'

[12] Which must be read in conjunction with art. 33(1) AMLD4, comprising the reporting obligation of financial institutions *vis-à-vis* the Financial Intelligence Unit ("**FIU**") of suspicious transactions.

Although the use of AI in transaction monitoring has been best practice for several years, neither AMLD4, Wwft nor the DPR explicitly reference such technologies. Instead, these instruments adopt a **technologically neutral**, **principles-based** approach. No specific methodology for transaction monitoring is prescribed; it is up to financial institutions to implement appropriate structures that enable them to adequately detect unusual transactions.

Traditionally, this has meant a mix of *manual* and *automated* monitoring processes. Over time, however, practice has increasingly shifted toward (semi-)automated systems, including the use of AI.

3.1.2 Forthcoming: AMLR

The AMLR introduces several provisions that are relevant to the use of AI – although not targeted specifically at agentic AI or transaction monitoring.[13] Nevertheless, these provisions impose limitations and safeguards that are directly relevant for institutions deploying agentic AI in this context.

Art. 76(5) AMLR permits the adoption of automated decision-making systems, including AI, under the following cumulative conditions:

I The data processed by the automated system must be limited to the data obtained during the CDD-process;

II Decisions regarding onboarding, offboarding, or the application of SDD or EDD must be subject to *meaningful human intervention*;

III Customers must be able to obtain an explanation of the automated decision and have the opportunity to challenge it (except in relation to a report to the FIU).

From this, it follows that decisions relating to transaction monitoring (and CDD in general) may, in principle, be automated – including those supported by agentic AI – provided they comply with these conditions. Accordingly, agentic AI may be used for transaction monitoring, assessing transactions on the basis of facts and data known to, or in the possession of, the financial institution.[14]

[13] Generally regulated similarly as under the AMLD4, see art. 20(1)(f) and 26 AMLR.

[14] Art. 69(2) AMLR.

Like AMLD4, the AMLR maintains a principle-based approach regarding the requirement to perform transaction monitoring. However, it introduces a more rule-based layer, such as the condition that all available data must be used. This creates potential tension with the limitation under point (**I**), which restricts automated systems to processing only data obtained during the CDD process.

In practice, this means that:

- Agentic AI may perform post-hit analysis, such as evaluating whether a transaction monitoring alert is a **true or false positive**. This would include reviewing the transaction against the relevant customer's CDD file;
- If the alert is deemed a true positive, the agentic AI system may forward a draft report for reporting to the FIU. However, per the AMLR, the act of reporting should still be performed by a human – specifically the compliance officer;[15]
- The agentic AI system may also autonomously decide to initiate a **review** of the customer's CDD file.[16] However, should the review result in the conclusion that EDD measures are required, or that the customer should be offboarded, meaningful human intervention would be necessary under point (**II**) above.

What constitutes *meaningful human intervention* remains undefined in the AMLR and will require further clarification. Nonetheless, it appears that fully automated application of EDD is not permissible under the current AMLR language.

Moreover, agentic AI systems may technically conclude – based on a completed review – that additional customer information is needed, or even that the customer should be offboarded due to an elevated risk. However:

- Offboarding remains subject to human review point (**II**) and must be explainable to the customer upon request point (**III**);
- A system-initiated request for additional data may fall outside the scope of point (**I**), as the data requested would not yet have been obtained during the CDD process. That said, arguably the request itself is part of the ongoing CDD process and therefore falls within the permitted use under point (**I**).

[15] Art. 69(6) and 11(2) AMLR.

[16] E.g., because there are reasons to suspect the risk-profile of the customer has changed, see art. 26(3) and Recital (69) AMLR.

3.2 Transaction monitoring for fraud-detection purposes

Another use of transaction monitoring can be found in the applicable payment regulation. Specifically, the second Payment Services Directive ("**PSD2**")[17] prescribes transaction monitoring in relation to the (dis)application of strong customer authentication ("**SCA**").[18] This transaction monitoring should enable payment service providers ("**PSPs**") to detect unauthorised or fraudulent payment transactions.[19] Akin to the CDD transaction monitoring, PSD2 only set principle-based rules that are agnostic of the actual technology being used (e.g., agentic AI).

Under the proposed Payment Services Regulation ("**PSR**")[20] (specifically the Council position)[21], the promise of AI for fraud related transaction monitoring is specifically mentioned.[22] This fits in the larger trend also visible in the AMLR, of the EU expressly embracing powerful technology such as AI for data intensive tasks like transaction monitoring.[23]

[17] Directive (EU) 2015/2366 of the European Parliament and of the Council of 25 November 2015 on payment services in the internal market (2015) (OJ L337/35).

[18] Art. 97 PSD2.

[19] Art. 2(1) Commission Delegated Regulation (EU) 2018/389 of 27 November 2017 supplementing Directive (EU) 2015/2366 of the European Parliament and of the Council with regard to regulatory technical standards for strong customer authentication and common and secure open standards of communication.

[20] Proposal for a Regulation of the European Parliament and of the Council on payment services in the internal market and amending Regulation (EU) No 1093/2010 (COM/2023/367 final).

[21] We refer to the Council position here as the PSR, though the PSR itself is still subject to interinstitutional negotiations and may change. The choice for the Council position is due to the recent nature of the document (June 2025), whereas the previous version of the PSR (the European Parliament Position) is more than a year older.

[22] Recital (103) PSR: *'[...] Therefore, to be able to prevent ever new types of fraud, transaction monitoring should be constantly improved, making full use of technology such as artificial intelligence. [...]'*

[23] Though we do not focus on fraud transaction monitoring in this paper, the considerations of this paper principally do apply to it and could be used to embed agentic AI powered fraud transaction monitoring in financial institutions' organisations.

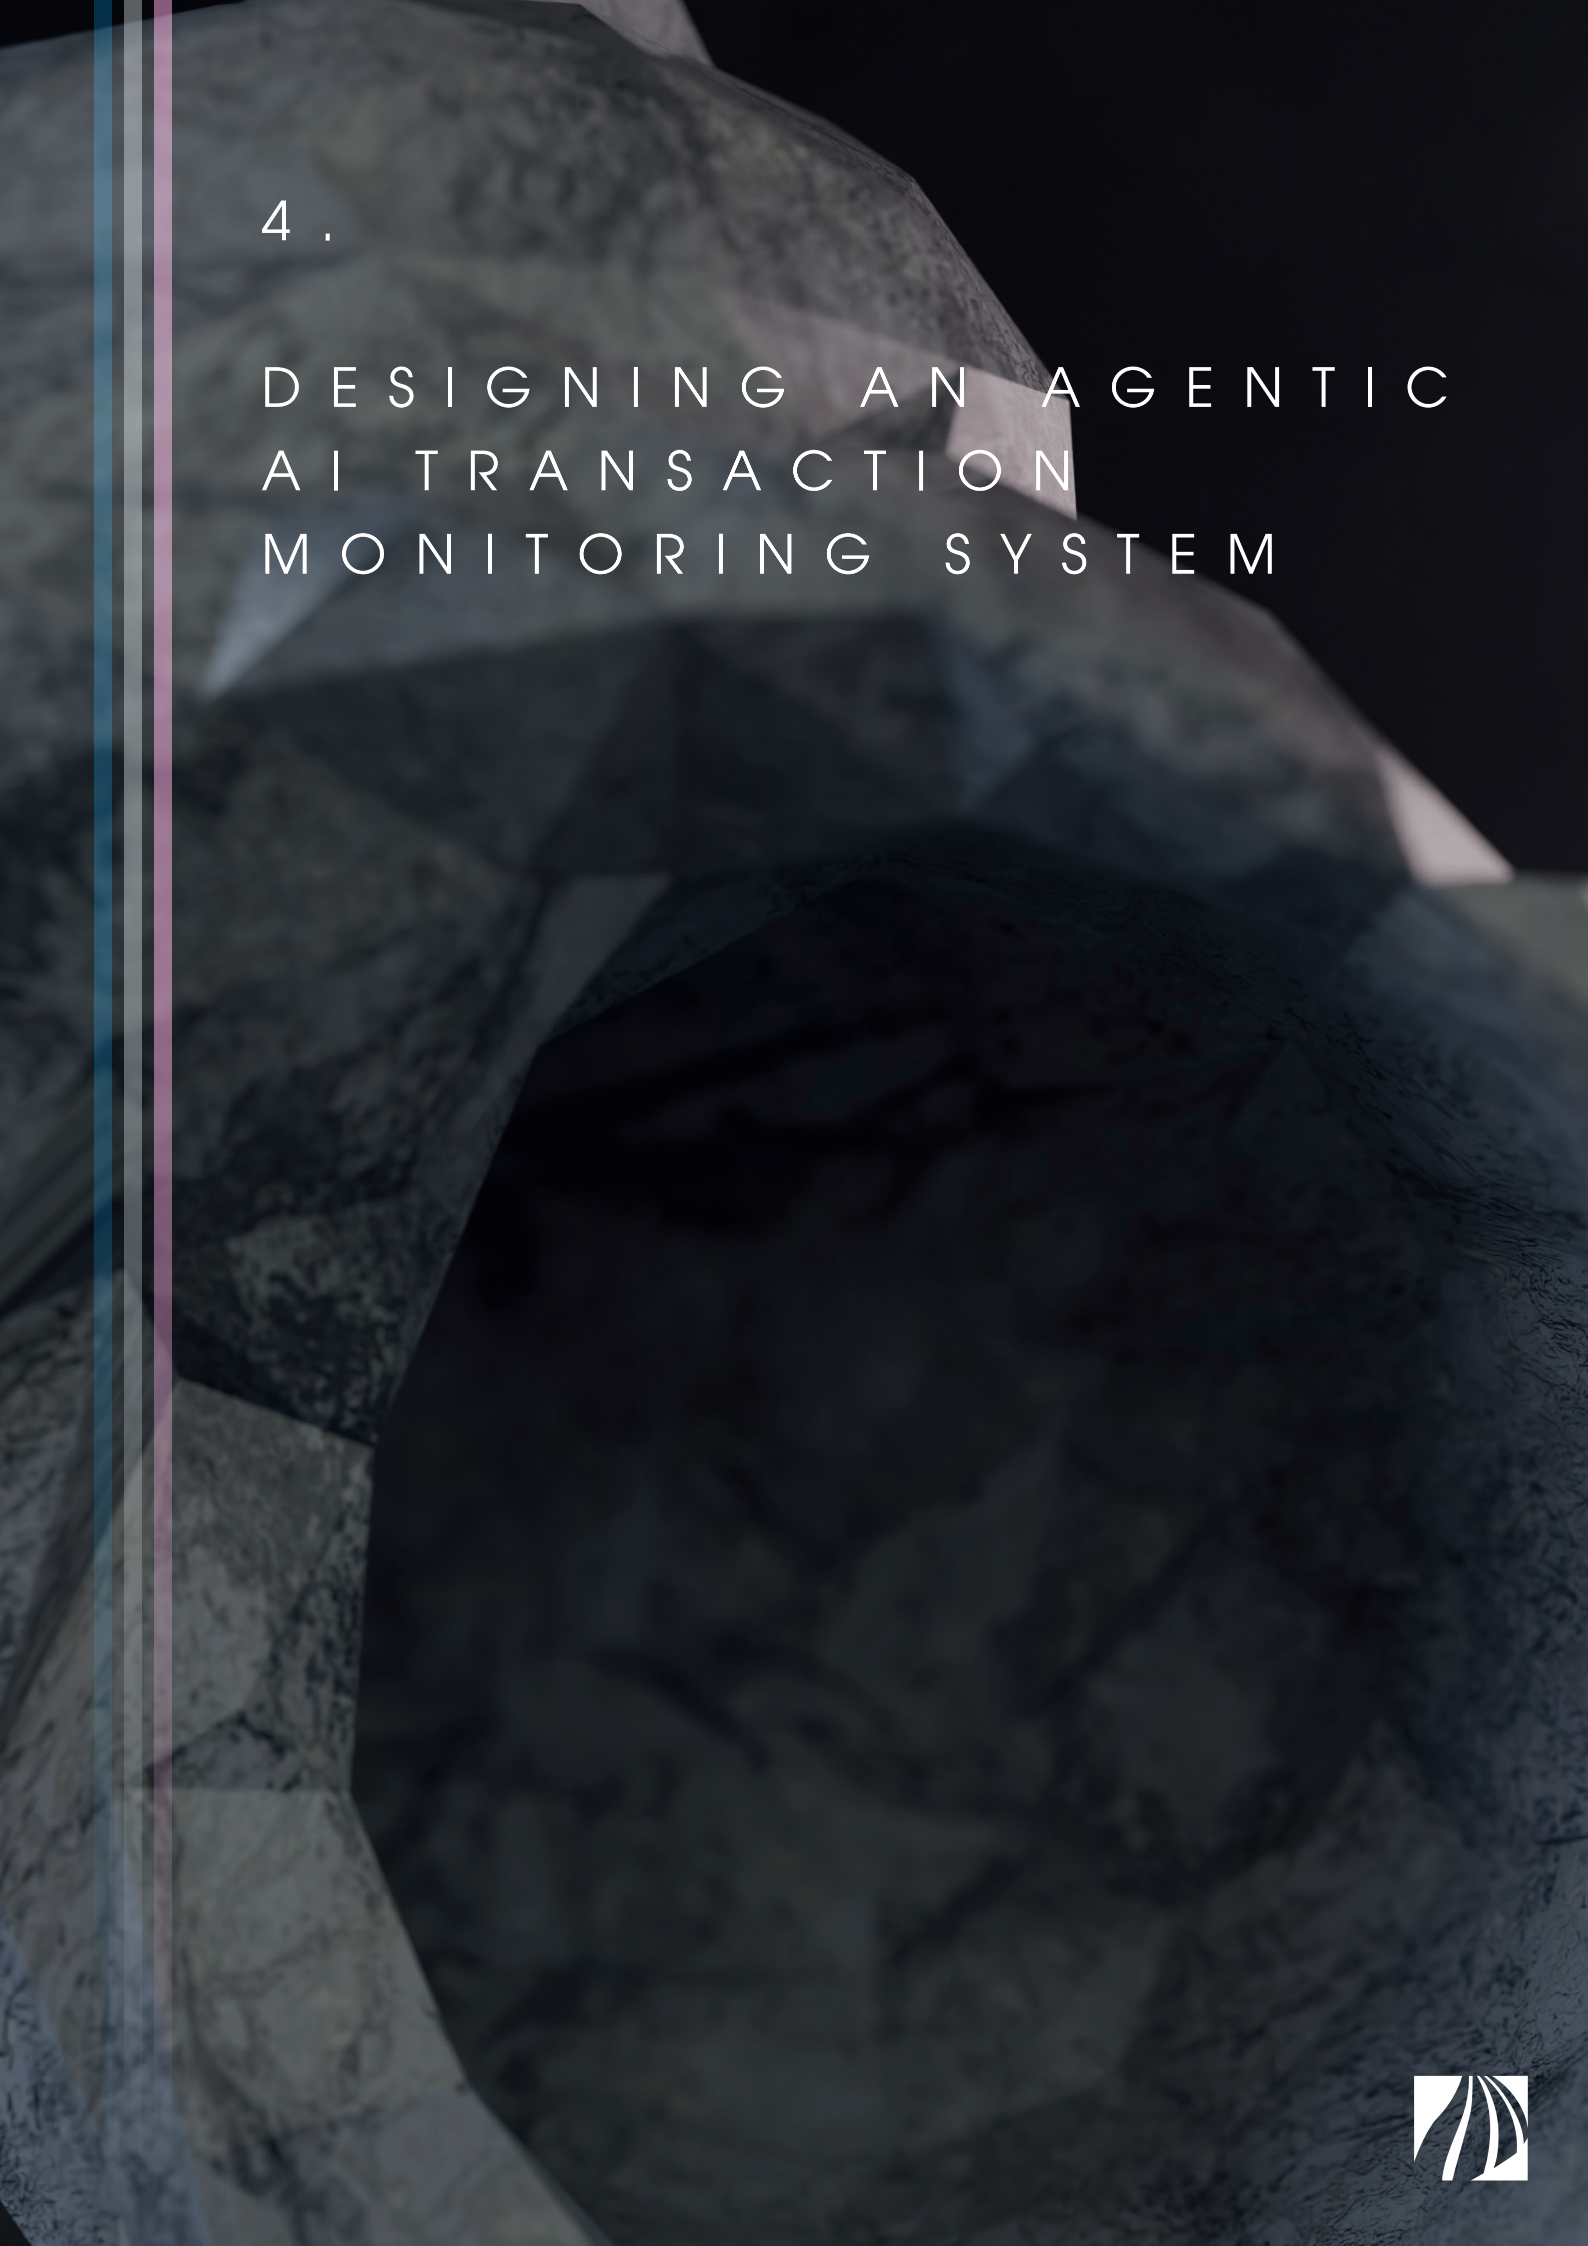

4.

DESIGNING AN AGENTIC AI TRANSACTION MONITORING SYSTEM

4. DESIGNING AN AGENTIC AI TRANSACTION MONITORING SYSTEM

4.1 Theoretical design

This section outlines a theoretical design for an agentic AI system used in transaction monitoring, taking into account the purposes as set out in the previous chapter. Such a system follows the typical stages of the transaction monitoring process – screening, analysis, and reporting – but allocates these tasks to dedicated AI agents rather than to human operators or (partial) 'regular' automation. Each agent is assigned a specific, clearly defined function. More granular agents may be introduced where necessary for complex tasks.

Agentic AI transaction monitoring system design

To concretely design a (partial) agentic AI system, financial institutions could observe the day-to-day activities of their AML/CTF function, mapping all tasks and queries in the process. On the basis of such inventory of human tasks, the actual AI system architecture could be designed and rolled-out step-by-step.

Suggestion: The created AI agents all have specific tasks that in a fully human system represent the different roles and tasks of the human operators. Such AI agents could be mapped in a comprehensive organisational chart, representing human and AI tasks in one overview.

Suggestion: The process-flow of the agentic system should be automated as far as possible, so that straight-through-processing is facilitated. This ensures the best possible efficiency result, whilst reducing human-AI interaction biases. Automation should be accompanied by appropriate and robust human control.

4.1.1 Orchestration layer

The agents operate within a coordinated orchestration framework that ensures communication, error logging and reruns, handover of results, and traceability across all steps. CoT outputs, metadata, and decision logs must be retained and structured such that each agent's input and reasoning can be traced across the workflow. Where relevant, agent interactions are governed by API protocols or rule-based process orchestration engines. In other words, the AI agents need an AI manager that coordinates their efforts and communication.

Each agent's role should be clearly documented in the financial institution's policies and procedures, with supporting governance structures in place (see paragraph 5) and associated risks mitigated (this will be covered in Part II of the White Paper Series on agentic AI). **Only when these conditions are met should a financial institution consider deploying an agentic AI transaction monitoring system.**

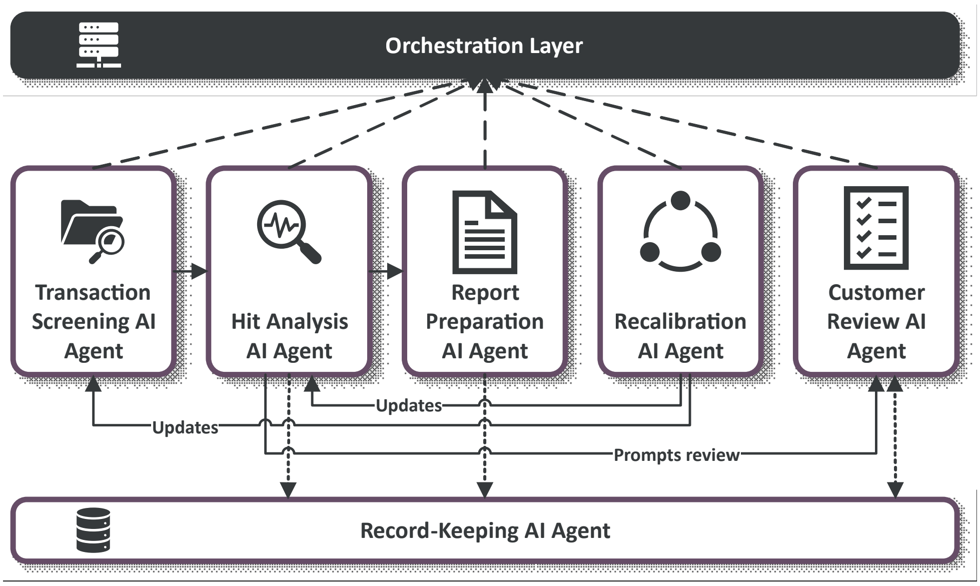

This system is a conceptual model. It is not intended as a *one-size-fits-all* solution or guaranteed to meet regulatory requirements in all cases. Agentic AI should always be implemented **in addition to** a team of human AML/CTF analysts and an appropriately functioning AI Governance Taskforce (see paragraph 5.1.1). Financial institutions, including their management bodies, should ensure that they have a clear understanding of the agentic AI system supported process at all times (know-your-structure).

4.2 Different roles within the system

4.2.1 Transaction screening AI agent

The transaction screening agent assesses transactions against business rules and transaction profiles defined by the financial institution. These rules and profiles may be based on typologies and indicators of money laundering or terrorism financing. This *analytical* agent may rely on traditional rule-based or classical machine learning models and does not necessarily need to be based on an LLM. Optionally, it may be split into pre-transaction and post-transaction screening agents, each with their own underlying models.

The AI agent should be capable of refining its own rules and customer profiles over time (e.g., through decision-tree models). These changes must be *transparent* to the institution, with a clear understanding of why a rule was adjusted and what new behaviour it is intended to detect. The agent may also identify unusual customer behaviour and outliers, potentially including network or relationship-based analysis. This latter capability could also be assigned to a dedicated network analysis agent.

4.2.2 Hit analysis AI agent

The hit analysis agent classifies alerts from the screening agent as either **true-positives** or **false-positives**. The classification process is based on human-defined rules and parameters, which may be adjusted dynamically by a separate recalibration agent. The analysis should be performed in an *explainable* and *interpretable* manner.

The agent uses data available to the financial institution, supplemented – where appropriate – with external data sources that the institution has previously vetted and approved. If a hit is classified as a true positive, it is passed to the report preparation agent. If a hit is classified as a false positive, the result is stored in the central data repository for future reference.

The agent should operate within pre-defined confidence thresholds. Where classification confidence falls below a minimum threshold, the alert should be escalated to a human AML/CTF analyst. In addition, a build-up of false-positives beyond a predefined limit should trigger re-evaluation of the alert by the hit analysis agent (or the recalibration agent).

In cases where the hit analysis agent cannot reach a decision within the set parameters – due to data gaps, conflicting indicators, or ambiguous transaction patterns – it should flag the alert for escalation. Escalated alerts may be routed to either a human review or, where defined, a coordination mechanism within the AI governance framework. All such escalation events should be logged and reviewed periodically to assess systemic blind spots or threshold misalignments.

4.2.3 Report preparation AI agent

Once a hit is confirmed as a true positive (either by the AI system or following human intervention), it is passed to the report preparation agent. This agent compiles a report for submission to the FIU, including the following:

I Identity of the customer, UBOs, and – where relevant – the party on whose behalf the transaction was made;

II Identity document type and number;

III Nature, timing, and location of the transaction;

IV Value, origin, and destination of the funds involved; and

V The reasons the transaction was flagged as unusual.

While the first four points are generally extractable from existing structured data, the fifth requires reasoning. The agent may incorporate the CoT analysis (or post-hoc analysis methods such as SHAP or LIME)[24] conducted by the hit analysis agent to provide a concise rationale. The report preparation agent may also produce internal management reports with further detail or suggested follow-up actions.

All reports must be subject to review and approval by a human compliance officer, who remains ultimately responsible for submissions to the FIU and for oversight of the transaction monitoring process within the internal governance structure.[25] The system must ensure that reports are *interpretable* and *auditable* by human reviewers, and that such reviews are appropriately documented.

[24] As expanded upon in Part II of the White Paper Series on agentic AI.

[25] Art. 2(3) Wwft.

4.2.4 Recalibration AI agent

The recalibration agent tests and recalibrates the screening and hit analysis agents on a regular basis. It ensures that the agents' output aligns with the institution's risk assessments (e.g., SIRA or GRA) and that performance remains consistent over time. In essence, this agent ensures that a functioning feedback loop is maintained that continuously improves the monitoring process.

The recalibration process may involve:

I Comparing current outputs to a baseline model to detect **drift**;

II Back testing against historical transaction data and prior unusual transaction reports; and

III Reassessing whether business rules and analysis thresholds remain valid.

Drift detection is especially critical to ensure the AI system does not evolve in unintended ways or deviate from institutional risk tolerance. Data for recalibration is sourced from the CDD data management agent. Depending on governance rules, adjustments may be implemented directly or only after human review. All changes must be logged with **justification** and **versioned accordingly**.

Box: Drift vs. Hallucinations

AI drift and **AI hallucinations** are distinct risks in the use of AI systems. *AI drift* refers to the gradual degradation of a system's performance over time, typically caused by changes in the underlying data, market conditions, or external environment that no longer match the data on which the original AI model was trained. In contrast, *AI hallucinations* occur when an AI system generates outputs that are factually incorrect, illogical, or fabricated, despite appearing plausible. While drift undermines accuracy by misalignment with reality, hallucinations undermine reliability by introducing false or misleading information.

Though hallucinations are more salient with high-impact cases, drift is potentially more harmful in the long-run, as model performance may degrade significantly and systemically without being noticed.

In addition to performance testing, the recalibration agent (or transaction screening agent) may be tasked with identifying emerging risk patterns or previously unclassified transaction typologies. This function may rely on *unsupervised learning* or *anomaly detection* models that surface deviations from baseline behaviour or cluster unusual transaction patterns. Alerts from this process should be routed for review by human AML/CTF analysts (or the AML AI Governance Taskforce – see paragraph 5.1.1) and may ultimately feed into refinement of the screening or analysis logic of the other agents.

4.2.5 Customer review AI agent

The customer review agent is triggered by specific events (e.g., a true positive hit) or periodic schedules to reassess the customer's CDD file. The agent accesses internal data to verify whether:

I CDD information remains complete and accurate; and

II The customer's risk profile has materially changed.

Where an elevated risk is identified, the agent may initiate additional information gathering or suggest contact with the customer. However, human intervention is required before EDD is applied or a customer is offboarded (i.e., before the 'exit protocol' is triggered). Any suggestions for exit must be supported by a clear and satisfactory rationale, and based on an *explainable* and *interpretable* reassessment of the CDD file.

4.2.6 Record-Keeping AI Agent

The record-keeping AI agent serves as the backbone of the agentic AI system.

I It maintains a logically structured, navigable database;

II It converts data between machine-readable and human-readable formats;

III It manages retention periods and alerts the oversight team when updates or deletions are due;

IV It stores process logs to ensure auditability and explainability; and

V It supplies relevant data to all other agents within the agentic AI system.

The database maintained by the AI agent does not necessarily have to be the entire CDD database of the financial institution, but it should be comprehensive for the activities of the agentic AI system. The curation of the data is of paramount importance, as the performance of the agentic AI transaction monitoring system will be largely based on the quality of the data used. The data pertaining to specific customers should be easily accessible to the compliance function and the first line analysts operate next to the agentic AI system.

The record-keeping AI agent also supports oversight by generating periodic reports with relevant performance statistics. Metrics included in these reports should be pre-defined by the institution and should meaningfully support management's understanding and control of the system.

5.

GOVERNANCE REGULATION – ORGANISATIONAL EMBEDDING OF AGENTIC AI

5. GOVERNANCE REGULATION – ORGANISATIONAL EMBEDDING OF AGENTIC AI

The deployment of agentic AI in transaction monitoring raises several regulatory issues, including how it is embedded and governed within the financial institution (governance). Various EU-level regulatory frameworks impose governance obligations on financial institutions (e.g., CRD IV,[26] MiFID II,[27] PSD2). We focus here on the Dutch implementation of these EU frameworks into the Dutch Financial Supervision Act (*Wet op het financieel toezicht*, "**FSA**"), as predominantly implemented in art. 3:17 and 4:14 FSA.

Generally speaking, financial institutions must have a sound internal governance structure under the FSA, including in respect of AI.[28] Such governance structure should ensure that the risks related to AI system use are efficiently detected and mitigated, but equally that its opportunities are recognised and used.

5.1 Organisational model

As part of setting up a proper internal governance, financial institutions should create clear reporting lines, allocate responsibilities and ensure that accountability is properly divided.[29] Responsibility and accountability for the oversight of the financial institution is divided in accordance with the three lines of defence model ("**3LoD Model**"). Under the 3LoD Model, the responsibilities for controlling the business is spread out over three layers:

First-line-of-defence – The business lines, i.e., the units that are taking risks and are responsible for their own operational management directly and on a permanent basis.

Second-line-of-defence – The risk management function and the compliance function. Financial institutions may set up additional specific control functions (such as ICT security, AML/CTF function and AI Governance function).

[26] Directive 2013/36/EU of the European Parliament and of the Council of 26 June 2013 on access to the activity of credit institutions and the prudential supervision of credit institutions (2013) (OJ L176/338).

[27] Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments (2014) (OJ L173/349).

[28] DNB and AFM (2024), p. 24.

[29] Art. 3:17(2) FSA and art. 17 DPR.

Third-line-of-defence – The independent internal audit function conducting risk-based and general audits and reviews of the internal governance arrangements, processes and mechanisms.[30]

Typically, AML/CTF compliance is allocated within the second line, specifically within the realm of the compliance officer (though a separate AML/CTF officer may be appointed).[31] However, transaction monitoring can be structured within the first line, which structure is considered a good practice by DNB.[32]

Agentic AI systems, by their nature, challenge traditional notions of delegation and require careful consideration of who is responsible for what – and under what conditions human override or escalation is possible.

[30] Art. 2d(4) Wwft.

[31] See art. 2d(2) Wwft.

[32] DNB (2024), p. 58.

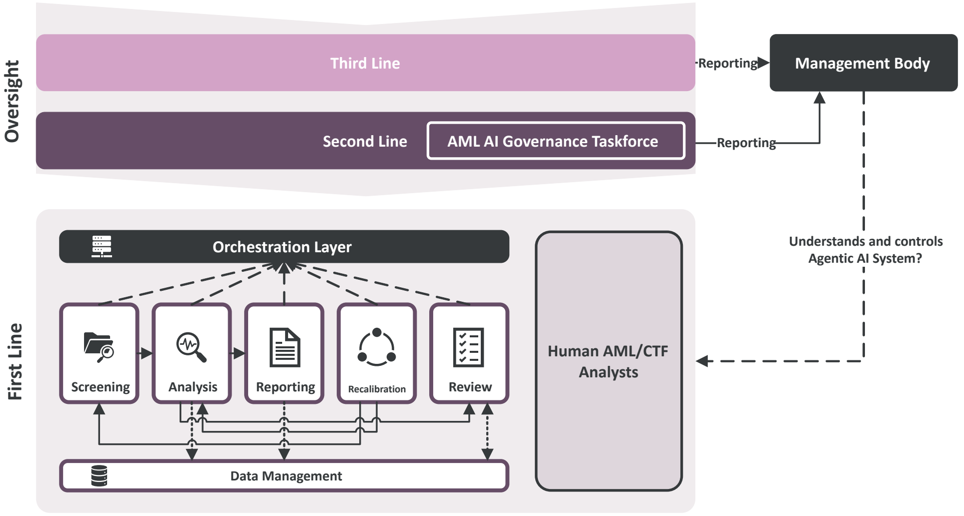

5.1.1 A dedicated AML AI Governance Taskforce

Generally, the use of an AI system should not mean that a financial institution has to completely rework its existing governance. Rather, where possible, the financial institution should focus on implementing additional or different governance arrangements where necessary to control the concrete use of AI systems. For instance, if a financial institution chooses to *only* use an (agentic) AI system for transaction monitoring, then *principally only* the governance surrounding that process should be altered.

The agentic AI transaction monitoring system itself should be overseen by the second and third line functions, as it produces risks that ought to be mitigated.[33] However, this oversight should be dynamic to accommodate the fast changing pace of the agentic AI system.

[33] In light of agentic AI, the embedding into the 3LoD Model becomes even more difficult, as the agentic AI system may perform parts of the first and second line function itself. For instance, the agentic AI system might comprise an AI agent that continuously analyses for material deficiencies in the applied process using certain key compliance indicators ("**KCIs**"). However, for the purposes of this paper, we assume that the agentic AI system is a separate 'entity' overseen by the second line.

5.1.1.1 AML AI Governance Taskforce

What could be particularly useful in this regard would be a combined second line taskforce of AI and AML/CTF specialists that can track and understand the developments of the agentic AI transaction monitoring system: the AML AI Governance Taskforce ("**Taskforce**"). The AI part of the expertise would monitor the governance of the model itself, whilst the AML/CTF expertise would monitor the continued alignment and appropriateness of the agentic AI system in the compliance function. The Taskforce should be equipped with sufficient seniority to ensure organisational buy-in. At the same time, the Taskforce should comprise plenty of 'first line' AI, ICT and AML/CTF expertise to facilitate a bottom-up style governance.[34]

In light of the substantial autonomy of an agentic AI transaction monitoring system, and its potential impact for customers and the integrity of the financial institution, it seems difficult to imagine a governance structure where there would not be a dedicated Taskforce (full-time) overseeing the AI system.

The AML/CTF policy, and/or the AI policy, of a financial institution should clearly set out the tasks, roles and responsibilities of the Taskforce and its interaction with the other first and second line functions involved in AML/CTF compliance. For instance, such policy should:

I Define escalation procedures, involving the Taskforce and other relevant staff, whereby we deem a (very) short escalation procedure to the management body appropriate;

II Allocate the accountability and ultimate responsibility for AI systems use to the management body; and

III Ensure that all involved personnel has a sufficient knowledge of how AI systems are used in the organisation and the related risks.

[34] Without infringing on the independence of the second line function.

5.1.2 Human oversight of the AI system

To ensure greater *accountability*, supervisory guidance on AI use emphasises the role of human intervention. The degree of human oversight should be calibrated to the autonomy level of the AI system, following a risk-based approach.[35] This is to minimise the risk that AI-based decisions result in harmful outcomes, especially if the AI outputs have significant potential impact on customers. Transaction monitoring certainly has the potential for a large impact on customers, wherefore the AMLR specifically requires the *meaningful human intervention* in relevant decisions.

There are varying ways in which human intervention can be achieved in respect of an AI system. Currently, the most prevalent approaches are termed:

- The '**human-in-the-loop**' – Human intervention in the decision cycle of the AI;
- The '**human-on-the-loop**' – Human intervention during the design cycle and subsequent reviews; and
- The '**human-in-control**' – Primacy of humans in making critical decisions.

The Taskforce could fulfil just such a role, combining all three of the above approaches. This Taskforce substitutes for the absence of real-time human involvement by exercising heightened, system-level oversight. It should hold the responsibility for review, approval, and intervention throughout the lifecycle of the AI system and each of its AI agents. Ideally, each AI agent within the agentic AI system would be assigned a specific human overseer within the Taskforce.

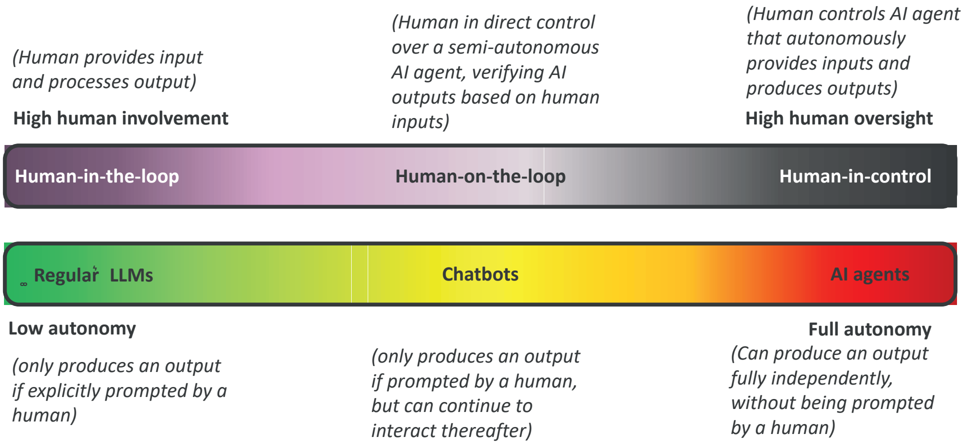
(Human provides input and processes output)
High human involvement
(Human in direct control over a semi-autonomous AI agent, verifying AI outputs based on human inputs)
(Human controls AI agent that autonomously provides inputs and produces outputs)
High human oversight
Human-in-the-loop
Human-on-the-loop
Human-in-control
Regular LLMs
Chatbots
AI agents
Low autonomy
(only produces an output if explicitly prompted by a human)
(only produces an output if prompted by a human, but can continue to interact thereafter)
Full autonomy
(Can produce an output fully independently, without being prompted by a human)

[35] Compare art. 14 EU AI Act.

First, there will be a human-in-the-loop despite an agentic AI model being largely autonomous. This will be in the form of, at least, the compliance officer that submits suspicious transaction report to the FIU. This ensures that at the end of the automated decision chain, there is always a human verifying the proposed decision (**output**) and – not unimportantly – taking responsibility for the submitted report. The same applies to offboarding or reviewing decisions (see paragraph 4.2.5).

Second, there will be a human-on-the-loop seeing as the Taskforce will continuously oversee the agentic model during the development (**input**) phase and the production (**throughput**) phase (covered in Part II of the White Paper Series on agentic AI). The Taskforce should ensure that at all moments in time, i.e., on a revolving basis, it is able to (swiftly) explain why the agentic AI system is doing something (e.g., by means of leveraging on CoT technology) and whether that fits into the governance of the financial institution involved.

Third, there will be humans-in-control as the Taskforce should be delegated sufficient *authority*, *agility* and *autonomy* to promptly take decisions as to the governance applied to, and the tenability of, the agentic AI system. It is of the utmost importance that the Taskforce has such 'executive' powers – within the guardrails set by the management body in the relevant policies – to match the comparative autonomy of the agentic AI system. Where the agentic AI system would exceed any pre-set guardrail, or would no longer be explainable, it should either be:

- Frozen and rolled-back to the last iteration of the AI system that did meet the governance requirements set by the financial institution; or
- Disabled altogether and (temporarily) replaced by a (human) fallback solution.

We discuss such operating limits and fallback solutions in Part II of the White Paper Series on agentic AI.

5.1.3 Reporting lines

Crucial in ensuring an adequately functioning governance structure around an agentic AI system, is clear reporting lines. Depending on the oversight model used to oversee the agentic AI system (see the box below), reporting lines may be longer or shorter. It will be of paramount importance that relevant information is passed through the reporting lines forthwith when it arises, ensuring that the entire organisation remains up-to-date with the calibration of the agentic AI system and the measures applied to control it.

It is crucial for the applied responsibility structure surrounding AI to be clear and comprehensible, particularly for the management body.[36] The management body – and all other staff – should at all times be able to understand and control the use of AI within its organisation, requiring constant **training** and frequent discourse with the Taskforce overseeing the AI system.[37] Ultimately, the management body remains *accountable* for the proper use of the AI system.

Box: AI governance model

Several theories exist as to how AI governance is best embedded in the 3LoD model. Broadly speaking, a financial institution can apply either a **horizontal** or a **vertical** approach, or a variant of both (e.g., the Taskforce).

[36] EBA Guidelines (2021), par. 78.

[37] Compare art. 14(4)(a) EU AI Act.

The **vertical approach** comprises an organisational structure where each business line has a (small) AI governance team that is dedicated to the use of AI within such business line.

✓ **Benefits**:

- specialised oversight of a particular AI application;
- a short distance between the actual AI system that is being overseen and the oversight, permitting for rapid decisions and control (i.e., more bottom-up oversight);
- short reporting lines between the first line and the AI governance team (second line).

× **Downsides**:

- resource intensive if the financial institution wishes to deploy AI systems in varying business lines; and
- disseminated oversight that requires a strong coordination effort to align the oversight on a more abstract level.

The **horizontal approach** comprises of a dedicated AI compliance team that oversees the use of AI across several business lines.

✓ **Benefits:**

- comprehensive oversight of AI usage – and thus risks – across the organisation (i.e., a more top-down style of oversight);
- concentrated expertise in a single team with a comparatively lower human staff requirement; and
- a harmonised and clearer reporting line of the business lines to the '**Taskforce**', and of the latter towards the management body.

× **Downsides:**

- relatively inflexible with more distance between the AI oversight and the business line, providing for a lesser familiarity and slower decision making; and
- might long-term lead to issues in timely adoption of AI system updates, facilitating the (incremental) build-up of 'small' risks.

5.2 Key principles for agentic AI system use

Any use of AI should be subject to certain key principles. These principles are used throughout this paper and are absolute edge-conditions to ensure the safe, prudent and ethical application of AI in finance.

- **Explainability** – An explainable AI system makes *transparent* how it arrived at a certain outcome (e.g., through CoT). Explainability is especially emphasised in a case such as transaction monitoring, where the output generated by the AI system may have a significant impact on the financial institution's customers and integrity. Arguably the biggest challenge to agentic AI use is managing the fundamental opacity of the underlying AI models of the AI agents that are part of the larger agentic AI system, which may or may not be exacerbated by their interaction with one another.
- **Auditability** – Transaction monitoring is one of the core processes of the control structure of a financial institution. As such, it is imperative that the agentic AI system itself and its performance is auditable for the third-line. Therefore, the each of the life-cycle phases of an agentic AI system requires proper documentation.
- **Interpretability** – Interpretability differs from explainability in the sense that it does not answer the question *how* the AI system came to a decision, but why. A financial institution may, for example, be able to explain *how* the AI system arrived at suspecting a suspicious transaction as a decision was reached using CoT technology, but explaining *why* that is a valid outcome requires a different analysis. This also comprises the requirement to have sufficient human skills available to interpret the AI system.

In essence, explainability, interpretability and auditability involve internal disclosure or transparency particularly to the board and senior management so they can better understand the risks and implications of AI use.

- **Transparency** – Transparency leads to greater accountability. Transparency is, for example, achieved through the documentation of how the AI system works, the governance surrounding it and the output generated by it, facilitating the internal and external assessment of the AI system.
- **Fairness** – The use of an agentic AI system for transaction monitoring must be accompanied with a strong focus on fairness of the generated output. As has become increasingly clear, automated

(AI) transaction monitoring has a high-risk for discrimination and *biased* outputs. Any agentic AI system should therefore safeguard the fairness of its generated output by a rigorous ethical analysis of the entire agentic AI system (extensively covered in Part II of the White Paper Series on agentic AI).

- **Reliability** – Expectations regarding reliability (or soundness) of AI systems resemble those for 'regular' (non-AI) models.[38] These involve the usual regular independent testing or monitoring to confirm that a model is performing as intended. They include monitoring metrics on validity, accuracy, robustness and reliability of both the data used and the methodologies applied.

[38] E.g., that the reliability of financial and non-financial information reported both internally and externally must be ensured; see, EBA Guidelines (2021), par. 145.

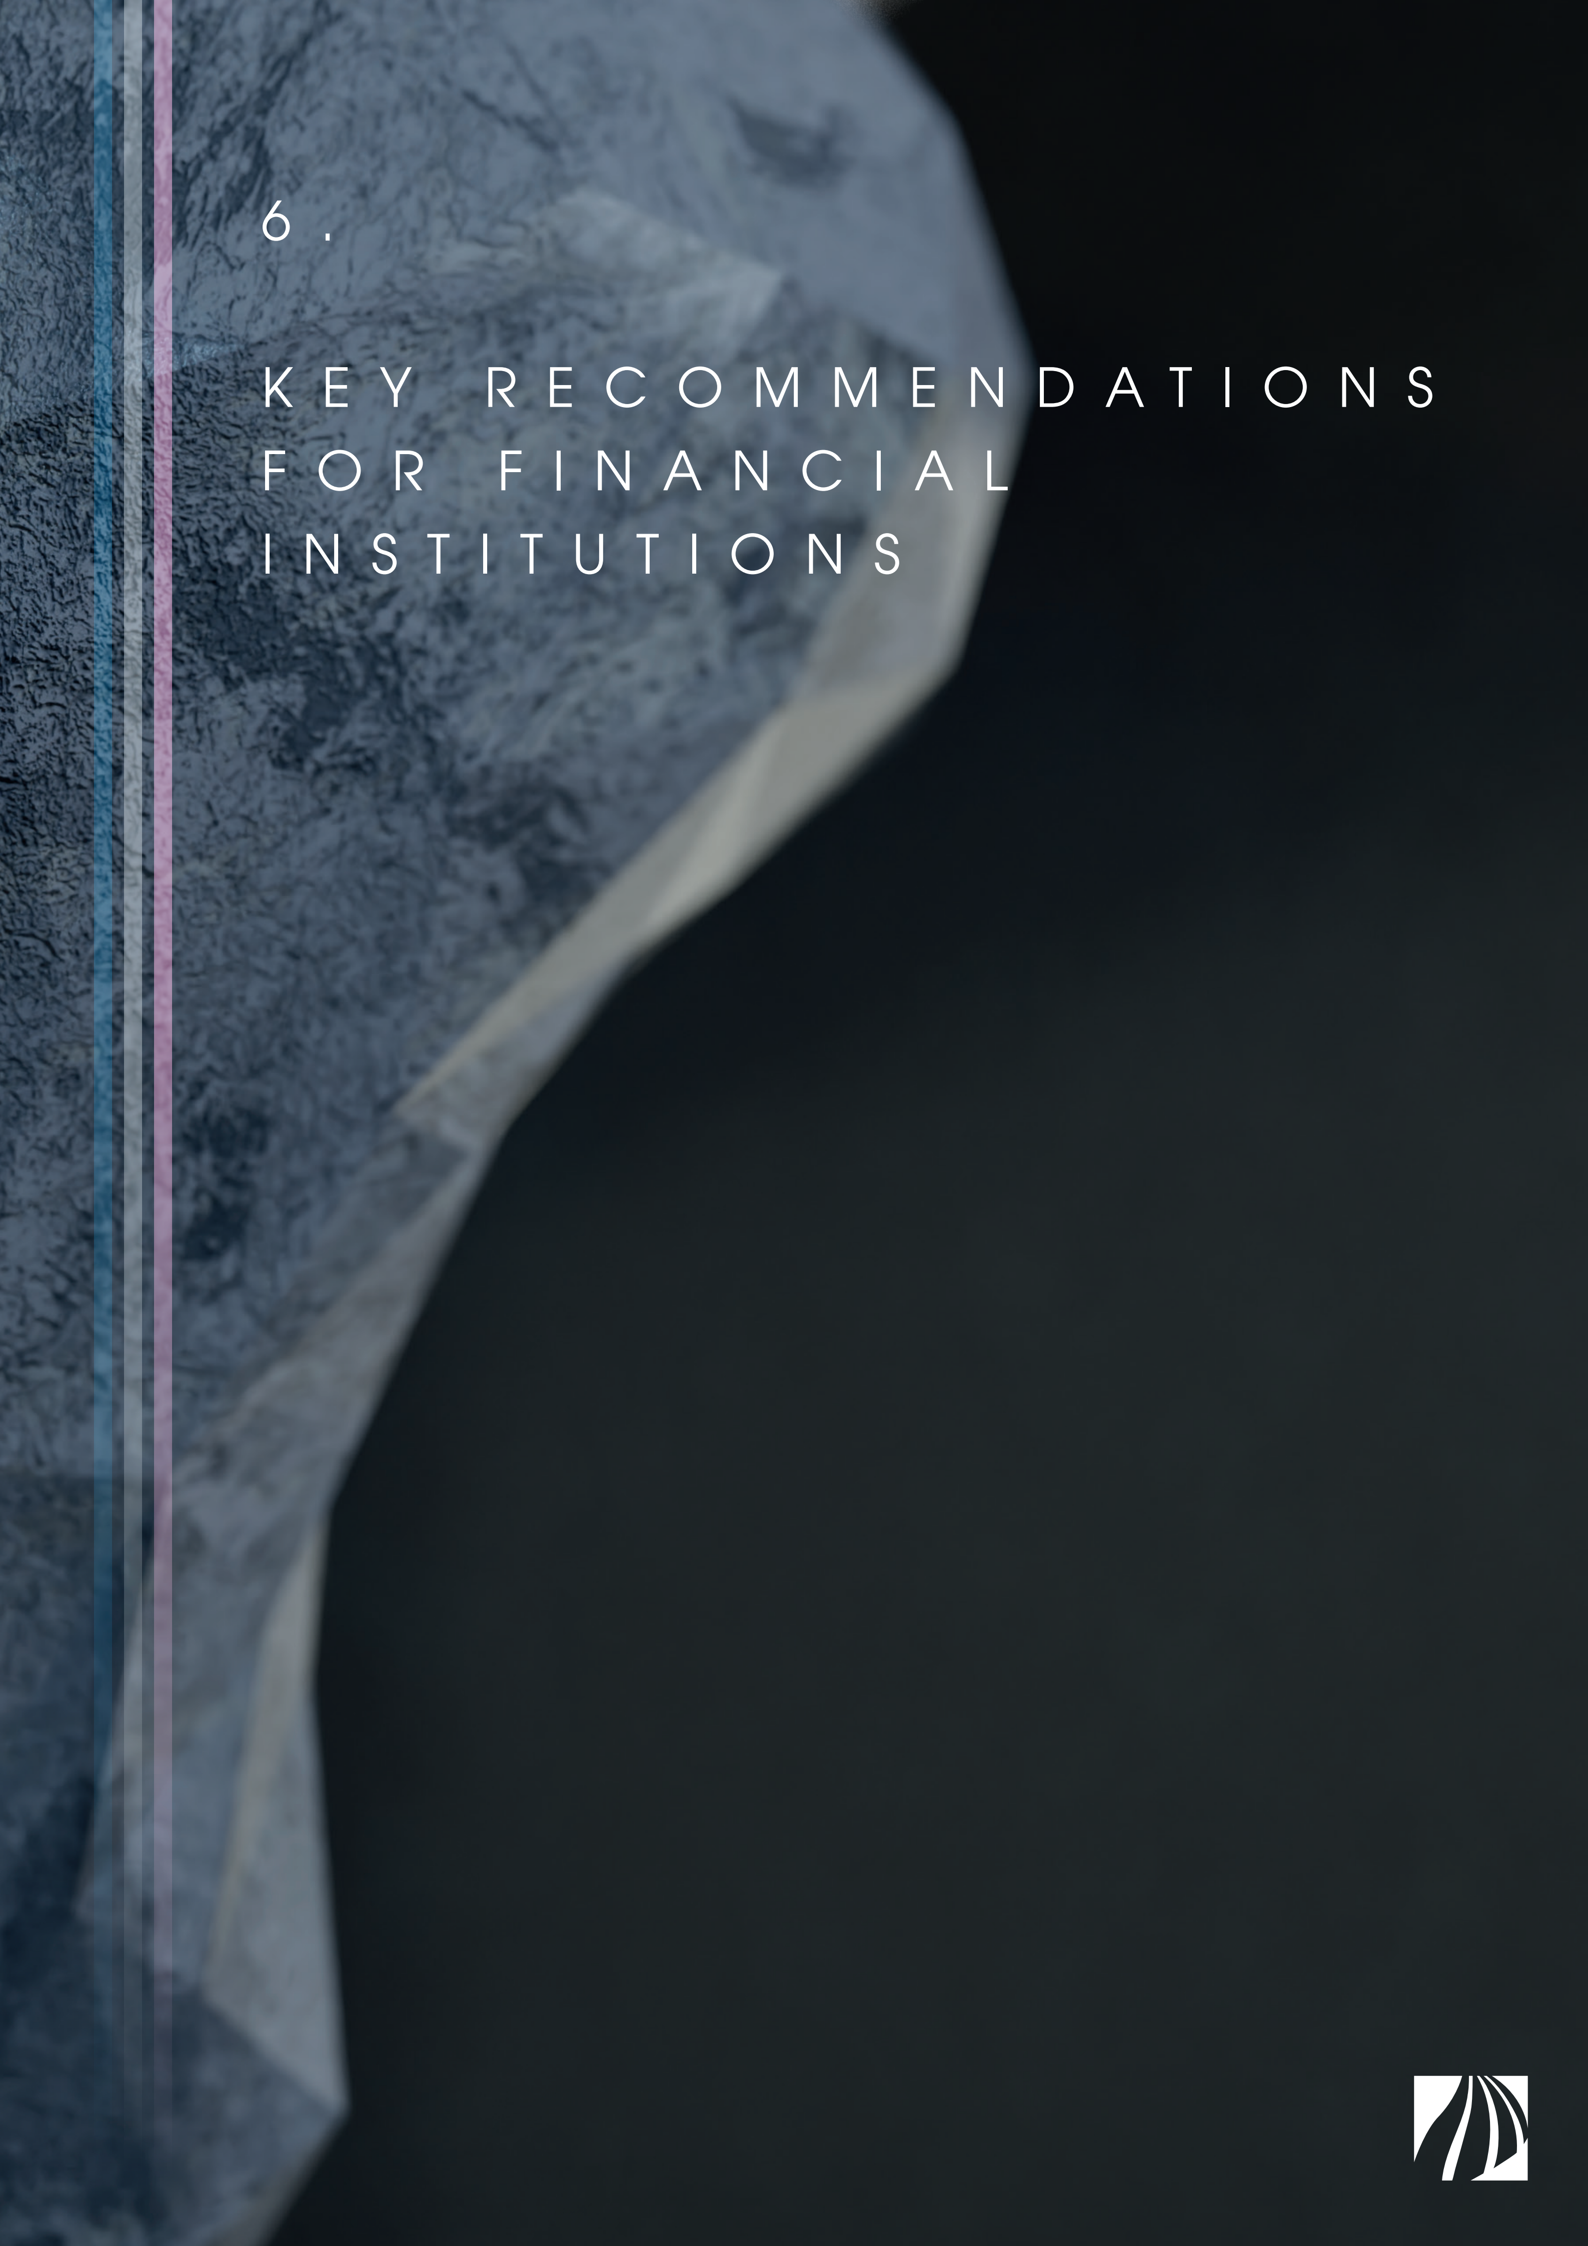

6.

KEY RECOMMENDATIONS FOR FINANCIAL INSTITUTIONS

6. KEY RECOMMENDATIONS FOR FINANCIAL INSTITUTIONS

To responsibly deploy agentic AI in transaction monitoring, financial institutions should take a structured and proactive approach. The following recommendations could be followed to implement such an approach:

I **Establish Strong Governance and Oversight**

a Create a Taskforce that combines expertise in AML/CTF, AI governance, compliance, and ICT. This taskforce should have sufficient seniority, independence, and authority to oversee the full lifecycle of agentic AI systems;

b Clearly allocate roles and responsibilities within the 3LoD (business, risk & compliance, and internal audit), and ensure that accountability for AI use ultimately rests with the management body;

c Implement escalation protocols to ensure rapid human intervention when AI exceeds set guardrails or produces unexplained results.

II **Ensure Explainability, Interpretability, and Auditability**

a Require all AI decisions to be documented in a way that can be explained both technically (explainability) and legally (e.g., CoT logs) (*interpretability*);

b Build audit trails into the system from the outset, enabling independent validation by internal audit and regulators;

c Develop internal training programs to equip staff with the skills to interpret and challenge AI-driven outcomes.

III **Maintain Human-in-the-Loop Controls**

a Ensure that critical decisions (such as customer offboarding, escalation to EDD, or reporting to the FIU) remain subject to meaningful human intervention;

b Calibrate oversight to the risk profile of the AI application; higher-risk decisions should require closer human review (and control);

c Define fallback procedures (e.g., freezing or rolling back the AI system) in case of malfunction, drift, or governance breaches.

IV Prioritise Data Quality and Ethical Use

a Invest in high-quality, well-curated datasets, recognising that data integrity directly affects detection accuracy and compliance outcomes;

b Conduct regular fairness and bias testing to prevent discriminatory outcomes in transaction monitoring;

c Align data use with regulatory requirements, particularly the AMLR's (and data privacy law) limitations to CDD data.

V Engage Proactively with Regulators and Supervisors

a Maintain transparent dialogue with supervisors (e.g., DNB), sharing lessons learned from pilot projects or (internal) sandboxes (if existent);

b Document compliance with AMLD4, AMLR, PSD2/PSR, and EU AI Act requirements, particularly regarding explainability, proportionality, and customer rights;

c Prepare for evolving supervisory expectations by adopting a flexible, principle-based compliance strategy.

VI Adopt a Phased and Controlled Rollout

a Begin with limited-scope pilots in lower-risk areas of transaction monitoring before scaling up;

b Apply "progressive automation": move from partial to greater autonomy only once human oversight mechanisms are proven robust;

c Continuously monitor performance, adjusting system design and governance in line with regulatory guidance and institutional learning.

REGULATORY COUNSEL
FINANCIAL SERVICES

Address: A.J. Ernststraat 595-F, 1082 LD Amsterdam, The Netherlands
Website: www.recofise.eu
E-Mail: info@recofise.eu
Tel: +31 (20) 520 68 54