SARTHAK GUPTA

Proposal and Comment Information

Title: Anti-Money Laundering and Countering the Financing of Terrorism Programs, R-1835

Comment ID: FR-0000-0133-02-C20

Submitter Information

Name: Sarthak Gupta

Submitted Date: 08/15/2026

Submitted by Sarthak Gupta in a personal capacity. The views are my own and not those of my employer, which has neither reviewed nor approved this comment. The full four-page letter is attached.

I work in model risk management and quantitative finance. My comment is narrow. The proposal makes an effective program turn on controls that are risk-based and reasonably designed and on independent testing of whether the program works. In a large bank most of that machinery is a detection system with thresholds, rules, and increasingly a statistical or learned component behind it. The Board rewrote its supervisory guidance for those systems twelve weeks before this proposal published, and the proposal does not mention it. Four points concern that seam.

1. The proposal encourages generative AI in AML/CFT while current model risk guidance disclaims it. Section IV.B states that innovative approaches could involve generative artificial intelligence and that banks incorporating innovative technologies will not on that basis incur additional supervisory risk. SR 26-2, issued April 17, 2026, superseded SR 11-7 and SR 21-8, and footnote 3 of the attached guidance states that generative AI and agentic AI models "are not within the scope of this guidance." SR 21-8 was the only issuance ever addressed to controlling AML detection systems. A bank is encouraged into the technology and left with no current statement of what controlling it requires. I recommend the final preamble state that the innovation provision goes to the choice of technology and not to the sufficiency of controls over it.

2. The proposal is silent on a classification that now determines whether monitoring is validated at all. SR 26-2 excludes deterministic rule-based processes from the model definition, while the FFIEC manual page cited at footnote 37 says surveillance monitoring systems include rule-based systems and that their programming methodology and effectiveness should be independently validated. When a monitoring engine leaves the model inventory it also leaves the change control, documentation standard and performance evidence that inventory status administers. I recommend the final preamble state that independent testing covers the design, calibration and detection performance of transaction monitoring regardless of classification.

3. Who may review the detection system is no longer stated anywhere. SR 11-7 said validation should be done by people not responsible for development or use. SR 26-2 refers to independence once and locates quality in the rigor of the review rather than organizational structure. Paired with point 2, a monitoring engine outside the model definition has no current statement of who is competent to challenge it.

4. The regulatory impact analysis books the benefit of fewer incorrectly flagged transactions and no paired cost of reduced coverage. When a change narrows coverage the routine metrics improve, because alert volume and productivity are computed on what the system still surfaces. The rescinded SR 21-8 stated that these systems place greater emphasis on coverage over efficiency. I recommend the final analysis record the paired cost and that banks retain evidence of the detection effect of coverage-narrowing changes, including sampling below revised thresholds.

Full letter attached, with every citation quoted and sourced.

Sarthak Gupta

Data Scientist II, Finance Models, Amazon

August 15, 2026

Benjamin W. McDonough

Secretary, Board of Governors of the Federal Reserve System

20th Street and Constitution Avenue NW

Washington, DC 20551

Re: Docket No. R-1835; RIN 7100-AG78, Anti-Money Laundering and Countering the Financing of Terrorism Programs

Dear Mr. McDonough:

I work in model risk management and quantitative finance. I have previously submitted comments to the Financial Stability Board on its consultation on artificial intelligence in finance and to FINRA on Regulatory Notice 26-14. I write in a personal capacity. The views here are my own and not those of my employer, which has neither reviewed nor approved this letter.

My comment is narrow. The proposal makes an effective program turn on controls that are risk-based and reasonably designed, and on independent testing of whether the program works. In a large bank, most of that machinery is a detection system with thresholds, rules, and increasingly a statistical or learned component behind it. The Board rewrote its supervisory guidance for those systems twelve weeks before this proposal published, and the proposal does not mention it. The four points below concern that seam.

1. The proposal encourages generative AI in AML/CFT while current model risk guidance disclaims it (Question 1)

Section IV.B states that "Innovative approaches could involve machine learning, generative artificial intelligence (GenAI), digital identity, blockchain monitoring and analytics, or application programming interfaces (APIs)," and that "Banks that responsibly incorporate innovative technologies into their AML/CFT programs will not incur on that basis any additional risk of being subject to a significant supervisory action or enforcement action solely based on the use of innovative technologies."

On April 17, 2026, the Board issued SR 26-2, Revised Guidance on Model Risk Management, which "supersedes and replaces SR letter 11-7" and "SR letter 21-8, Interagency Statement on Model Risk Management for Bank Systems Supporting Bank Secrecy Act/Anti-Money Laundering Compliance." Footnote 3 of the attached guidance states: "Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance." The same

footnote adds that "a banking organization's risk management and governance practices should guide the determination of appropriate governance and controls for any tools, processes, or systems not covered" by the guidance. That sentence points a bank back at its own judgment; it is not a control standard, and it is the only thing standing where a control standard used to be. SR 26-2 does not mention the Bank Secrecy Act anywhere. SR 21-8, the only issuance ever addressed to controlling AML detection systems, was withdrawn the same day.

A bank reading both documents is encouraged to put generative AI into suspicious activity detection, told that doing so carries no additional supervisory risk on that basis, and given no current statement of what controlling it requires. OCC Bulletin 2026-13 notes that "the agencies plan to issue in the near future a request for information that addresses model risk management generally and considers, in particular, banks' use of AI, including generative AI and agentic AI and AI-based models." An undated request for information is not a control framework, and the proposed rule would take effect twelve months after a final rule issues.

Recommendation: state in the final preamble that the innovation provision goes to the choice of technology and not to the sufficiency of the controls over it, and that a bank deploying generative AI in suspicious activity monitoring remains subject to the independent testing requirement and to the requirement that its controls be reasonably designed. If the Board intends to address the substance in the planned request for information, say so, so that the interval is a known gap rather than an implied permission.

2. The proposal is silent on a classification that now determines whether monitoring gets validated at all (Questions 8, 17, 19(a))

SR 26-2 defines a model as "a complex quantitative method, system, or approach that applies statistical, economic, or financial theories to process input data into quantitative estimates," and states that the term "excludes simple arithmetic calculations, such as those found within spreadsheets, as well as deterministic rule-based processes and software where there are no statistical, economic, or financial theories underpinning their design or use."

The FFIEC manual page this proposal cites at footnote 37 describes the affected systems directly: "Surveillance monitoring systems include rule-based and intelligent systems," and further down the same page, "the monitoring system's programming methodology and effectiveness should be independently validated to ensure that the models are detecting potentially suspicious activity."

The narrowing was deliberate and has been described publicly. Its effect on AML has not been addressed in any rulemaking document. FinCEN opened the question in its April 10, 2026 proposal, noting that it "shares certain concerns articulated in the comments to the RFI that these models, which are designed to assess different types of risks with different information input, processing, and reporting components may be overly burdensome and ill-fitted to address illicit finance risks." Its footnote cites the 2011 guidance at a Board web address that now returns a 404, for a document the Board superseded a week later.

The practical consequence is larger than one lost validation. When a monitoring engine leaves the model inventory, it also leaves the regime that administered its change control, its documentation standard, and its performance evidence, because in most banks those are enforced through inventory status rather than through the compliance function.

Recommendation: state in the final preamble that independent testing under the rule covers the design, calibration, and detection performance of transaction monitoring regardless of whether the bank classifies the system as a model under SR 26-2, and confirm in response to Question 19(a) that monitoring rules and parameters, together with the basis for their calibration, are among the components required to be in writing. Coordinate that statement with FinCEN so the two rules do not rest on inconsistent assumptions about which framework applies.

3. Who may review the detection system is no longer stated anywhere (Question 17)

The proposal sets a testable independence standard for program testing: "independent testing should be based on objective criteria designed to assess whether a bank has established and implemented an effective AML/CFT program and allocated resources consistent with its risk assessment processes."

Current model risk guidance no longer supplies a default for the system itself. SR 11-7 said, "Generally, validation should be done by people who are not responsible for development or use and do not have a stake in whether a model is determined to be valid." SR 26-2 refers to independence once, as "sufficient independence to maintain objectivity," and states that "The quality of validation process depends on the rigor and effectiveness of the review rather than on organizational structure of the banking organization's risk management function." That is defensible for models generally. Paired with the classification issue above, it leaves a monitoring engine outside the model definition with no current statement of who is competent to challenge it.

Recommendation: state that the independence expectation attaches to the function under review rather than to the classification of the system, so that review of monitoring design and calibration meets the same independence standard the rule sets for program testing.

4. The regulatory impact analysis books the benefit of recalibration and no paired cost (Question 20(b), and the Board's invitation on the economic analysis)

The analysis states that fewer low-risk customers would have "their non-illicit transactions being incorrectly flagged by an AML/CFT program," and adds, "The Board lacks the data to quantify the scale of this benefit." The costs section concludes that "the Board anticipates no increase in ongoing compliance costs resulting from the proposed rule."

The Board already names the offsetting failure in section IV.C.2, as an example of failing to implement a program in all material respects: "systems used to monitor for potentially suspicious activity failing to capture material volumes or types of transactions." Nothing in the description of independent testing asks for evidence that would surface it.

This is the part that is hard to see from outside a monitoring operation. When a change narrows coverage, the routine metrics improve. Alert volume falls and productivity ratios rise, because both are computed on what the system still surfaces. Activity that stopped being surfaced leaves no record to measure. The only routine evidence that speaks to it is sampling below the revised thresholds, and that evidence gets produced only when someone requires it. The rescinded SR 21-8 stated the priority plainly: "the objectives of most BSA/AML models place greater emphasis on coverage over efficiency," and a bank "may choose to accept a reduction in efficiency (such as by producing more alerts) in exchange for greater coverage in its automated transaction monitoring system." The impact analysis assumes the trade runs the other way.

Recommendation: record in the final analysis the paired cost of reduced coverage of higher-risk activity that has not yet been identified as higher risk. In the final preamble, state that where a change narrows monitoring coverage the bank should retain evidence of the detection effect of that change, including sampling below revised thresholds. And answer Question 20(b) yes for the narrow case: a change that materially reduces detection coverage warrants re-approval at the level that approved the program.

I would be glad to expand on any of the four points, particularly the evidence that makes a coverage change auditable after the fact.

Respectfully submitted,

Sarthak Gupta Data Scientist II, Finance Models, Amazon Submitted in a personal capacity