Model Drift Monitoring and Governance in Financial-Institution Machine Learning: Aligning Adaptive Detection Systems with Federal Model Risk Management Guidance

Md Soebur Rahman[1*], Chashi Sabrina Islam[1], Sini Cakmak[2]
[1] International American University, USA
[2]Sun Yat-sen University, CHINA

*Corresponding Author Email: soebrahman10@gmail.com

ABSTRACT

On April 17, 2026, the Federal Reserve, OCC, and FDIC jointly issued SR 26-2, Revised Guidance on Model Risk Management, which supersedes both the foundational 2011 guidance (SR 11-7 / OCC Bulletin 2011-12) and the 2021 interagency statement on BSA/AML model risk management (SR 21-8), introducing for the first time a uniform $30 billion total-asset threshold across all three agencies, replacing the FDIC's earlier, lower $1 billion threshold and the absence of any explicit threshold at the Federal Reserve and OCC. This article examines what this consequential, very recent regulatory change means specifically for governing adaptive, drift-monitored machine-learning models used in financial-institution fraud and credit-risk detection, an application area the original 2011 guidance was not designed to anticipate. Drawing on real, cited data including OCC examination findings that board reporting and outcomes analysis remain the weakest areas of SR 11-7 compliance industry-wide, at reported compliance rates of only 60 and 65 percent respectively, this article proposes a governance architecture aligning continuous, drift-aware model monitoring with SR 26-2's newly formalised, risk-based materiality-tiering approach. The article reviews the regulatory history culminating in SR 26-2, details its specific changes relevant to adaptive detection systems, addresses vendor and third-party model accountability under the revised guidance, and provides a candid assessment of what remains unresolved for community financial institutions navigating this transition.

Keywords: Drift Monitoring; Adaptive Detection Systems; Machine Learning; Financial Institutions

Introduction

For fifteen years, the Federal Reserve's Supervisory Letter SR 11-7, Guidance on Model Risk Management, jointly issued with OCC Bulletin 2011-12 on April 4, 2011, served as the definitive supervisory standard for how U.S. banking organizations must develop, validate, and govern the quantitative models underlying their business and risk decisions [1-4]. Written in the aftermath of the 2008 financial crisis, which had exposed how poorly validated credit-risk and mortgage-valuation models had amplified systemic damage, SR 11-7 established three core pillars, rigorous model validation encompassing conceptual soundness, ongoing monitoring, and outcomes analysis; a comprehensive model inventory; and board-level governance, that regulators have since applied, in supervisory practice, to an ever-expanding range of models, including, following explicit 2024 Federal Reserve confirmation, the machine-learning and AI systems, fraud-detection models, credit-underwriting algorithms, and anti-money-laundering surveillance tools, that increasingly define modern financial-institution operations [8].

This fifteen-year-old framework, however, was not designed with adaptive, continuously retraining machine-learning models specifically in mind, and its application to such models has, by regulators' own more recent acknowledgement, created interpretive strain: AI/ML

models can lack the explicit, interpretable functional form SR 11-7's conceptual-soundness review presumes, their behaviour can change as they retrain on new data in ways a traditional, largely static regression model does not, and their complexity makes evaluating conceptual soundness a fundamentally different exercise than reviewing the assumptions underlying a linear model [1]. Separately, and independent of the AI-specific interpretive challenge, industry and examiner experience over this same period revealed that SR 11-7's broad, principles-based model definition had, in practice, pulled a very wide range of tools, including simple spreadsheets, rule engines, and basic workflow automation, into formal model inventories at many institutions, creating governance burdens difficult to tier, prioritise, and sustain, particularly at smaller, resource-constrained institutions [4].

On April 17, 2026, the Federal Reserve, jointly with the OCC and FDIC, issued SR 26-2, Revised Guidance on Model Risk Management, which explicitly supersedes and replaces both SR 11-7 and the 2021 interagency BSA/AML-specific statement, SR 21-8, consolidating fifteen years of supervisory experience and industry feedback into a single, updated framework [1][2]. This is a highly consequential, and very recent, regulatory development directly relevant to any institution deploying adaptive, drift-monitored fraud-detection or credit-risk models, the specific subject of this article and of companion articles in this series. The timing of this development is also worth noting explicitly: SR 26-2 was issued only months after the Office of the Comptroller of the Currency and the Federal Deposit Insurance Corporation had separately proposed, in October 2025, a rulemaking to define "unsafe or unsound practice" and limit the issuance of Matters Requiring Attention, a related but distinct regulatory development signalling a broader supervisory shift toward more precisely calibrated, less broadly discretionary examiner findings across bank supervision generally, of which the model-risk-management guidance overhaul this article addresses appears to be one specific, concrete manifestation. Section 2 reviews the regulatory history culminating in SR 26-2 in greater depth. Section 3 details SR 26-2's specific, substantive changes and their relevance to adaptive machine-learning systems. Section 4 proposes a governance architecture aligning continuous drift monitoring with SR 26-2's risk-based, materiality-tiered approach. Section 5 addresses vendor and third-party model accountability. Section 6 reviews documented compliance gaps under the prior guidance and their implications going forward. Section 7 discusses remaining open questions for community financial institutions, and Section 8 concludes.

Methodology and Source Selection

Given that SR 26-2 was issued in April 2026, only months before this article's preparation, sources cited here were identified through targeted searches for the primary Federal Reserve supervisory letter text itself, contemporaneous law-firm and advisory-firm analyses published in the weeks following its issuance, and, for historical and comparative context, primary documentation of the superseded SR 11-7 and SR 21-8 guidance and independently reported examination-trend data on compliance with that prior guidance. Given the very recent issuance of SR 26-2, the body of independent, longer-term analysis of its practical effects is necessarily limited at the time of this article's preparation; this constraint is treated explicitly in the evidence assessment in Section 7 rather than glossed over.

Regulatory History

Figure 1. Evolution of Federal Model Risk Management Guidance, 2011–2026

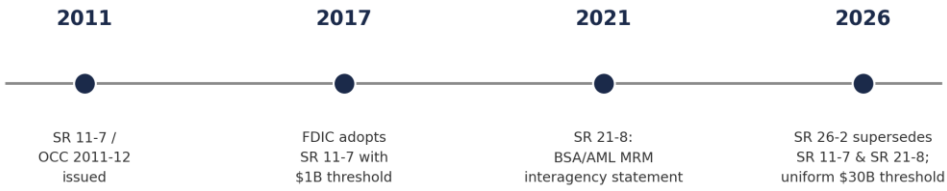

Figure 1 traces the regulatory lineage this article addresses. SR 11-7 and OCC Bulletin 2011-12 were issued jointly on April 4, 2011, establishing the original three-pillar model risk management framework [4]. Neither the Federal Reserve nor the OCC adopted any explicit asset-size threshold in this original guidance; the FDIC, adopting the same guidance in 2017, added a specific $1 billion asset threshold of its own, with the proviso that the guidance would still apply below that threshold to institutions with significant or complex model use [2]. In 2021, the federal banking agencies, in consultation with FinCEN and the National Credit Union Administration, issued SR 21-8, a BSA/AML-specific interagency statement clarifying how banks should apply SR 11-7's general principles to models and systems supporting Bank Secrecy Act and anti-money-laundering compliance specifically, while according primacy to existing BSA/AML regulatory requirements themselves [9].

Most significantly for this article, on April 17, 2026, the Federal Reserve, OCC, and FDIC jointly issued SR 26-2, which explicitly supersedes both SR 11-7 and SR 21-8 in their entirety, consolidating fifteen years of accumulated supervisory experience into revised, updated guidance [1]. This is the regulatory development this article's remaining sections examine in detail.

SR 26-2: Substantive Changes and Relevance to Adaptive ML

The New Uniform Asset Threshold

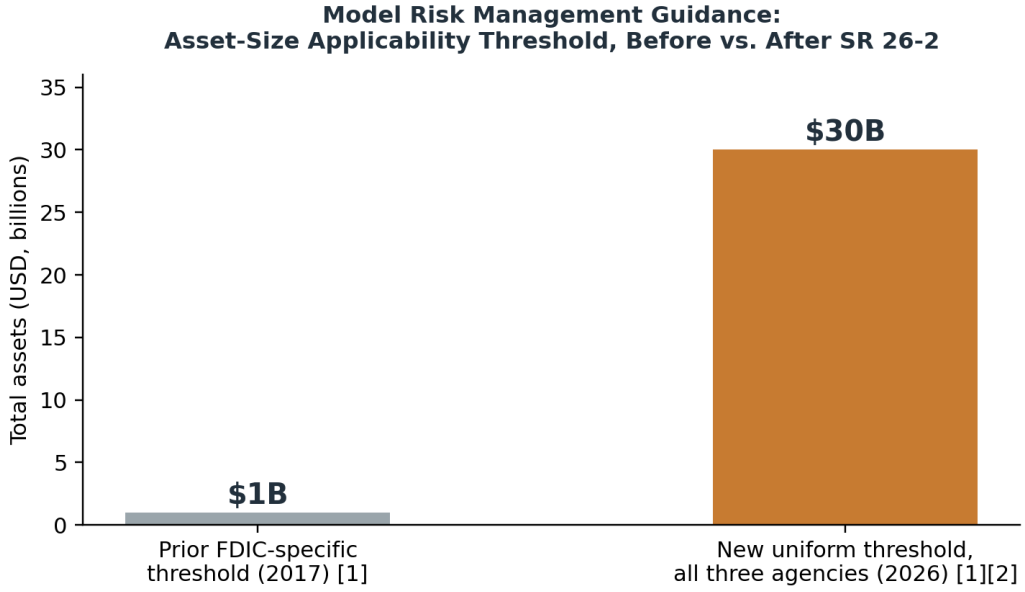

The single most consequential change SR 26-2 introduces, for the purposes of this article series' consistent focus on community financial institutions, is a new, uniform $30 billion total-asset threshold applied identically across all three agencies, replacing the FDIC's prior, considerably lower $1 billion threshold and the complete absence of any explicit threshold previously maintained by the Federal Reserve and OCC [2]. Figure 2 illustrates this change directly: the applicable threshold has increased thirtyfold under the FDIC's prior framework, and has moved from no explicit threshold at all to an explicit $30 billion figure under the Federal Reserve and OCC. SR 26-2 states plainly that it is "expected to be most relevant to banking organizations with over $30 billion in total assets" [1][2].

This change, however, should not be read as a blanket exemption for smaller institutions. SR 26-2 and the analyses examining it are consistent in framing the guidance as principles-based rather than purely scope-based: supervisory expectations remain tied to the materiality of an institution's model use, the complexity of its models, and the significance of the underlying use case, not to asset size alone, and the guidance explicitly notes it may still be relevant to smaller institutions with "significant exposure to model risk because of the prevalence and complexity of their models or because of activities outside the scope of traditional community banking" [2][3]. A community bank or credit union relying heavily on a sophisticated, adaptive, drift-monitored fraud-detection ensemble, of exactly the kind discussed throughout this article series, is a plausible candidate for this carve-back into scope, even if its total assets fall well below the new $30 billion figure, precisely because such a model's complexity and the materiality of the fraud-detection function it performs may place it outside what the guidance characterises as "traditional community banking" activity.

A Formal Materiality Construct

Beyond the threshold change, SR 26-2 introduces what several contemporaneous analyses describe as its most operationally significant substantive change: a formal materiality construct explicitly permitting institutions to calibrate governance rigour to actual model risk, applying full, rigorous validation to high-materiality models while allowing lower-materiality models to be governed through more streamlined, proportionate, and potentially automated processes [5]. This directly addresses the governance-burden problem noted in Section 1: rather than applying uniform validation intensity to every model in an institution's inventory regardless of its actual risk contribution, which industry feedback over the preceding fifteen years indicated created substantial cost and friction without a commensurate safety benefit, SR 26-2 gives institutions an explicit, formally sanctioned basis for risk-based differentiation [5].

This materiality construct is directly relevant to the drift-monitoring architecture this article series proposes throughout: an adaptive fraud-detection ensemble that materially influences whether a customer transaction is held, flagged, or a suspicious activity report ultimately filed is a strong candidate for high-materiality classification and correspondingly rigorous, continuous drift monitoring of the kind detailed in Section 4, while a lower-stakes, adjacent model, for example, one used only for internal case-prioritisation ranking rather than a direct customer-facing or filing decision, may reasonably warrant a lighter, more proportionate monitoring cadence under this same construct.

Reduced Prescriptiveness on Annual Revalidation

A further specific change, foreshadowed by OCC Bulletin 2025-26 issued in September 2025, clarifies that annual model revalidation is no longer a blanket community-bank requirement, formalising a shift toward validation cadence tied to actual model risk and materiality rather than a fixed, calendar-driven schedule applied uniformly regardless of a given model's risk contribution [4]. For the drift-monitoring architecture proposed in Section 4, this is a welcome and directly relevant clarification: it supports exactly the kind of continuous, event-triggered monitoring and revalidation cadence, responding to detected drift as it occurs rather than only at a fixed annual interval, that this article series has argued throughout is the more appropriate approach for adaptive, continuously evolving fraud-detection models specifically, rather than treating annual revalidation as an independently sufficient safeguard against concept drift occurring at any point between two annual review dates.

Length and Tone

A further, more qualitative but practically meaningful change is that SR 26-2 is reported to be considerably shorter than its predecessor, roughly half the length of SR 11-7's approximately twenty pages of dense, prescriptive expectations, and reads, in the assessment of at least one detailed industry analysis, differently in tone, moving away from prescriptive specification toward higher-level principles [4]. The revised guidance also includes an explicit disclaimer that departure from the guidance may not, by itself, give rise to supervisory criticism, a notable change from the more directly prescriptive tenor of the original guidance, though supervisory action remains available where an institution's model

risk management practices constitute an unsafe or unsound practice or violate applicable law, independent of the guidance's own non-binding status [2][5].

A Governance Architecture Aligned with SR 26-2

Figure 4 proposes a governance architecture explicitly structured around SR 26-2's risk-based, materiality-tiered approach, directly relevant to the adaptive, drift-monitored fraud-detection and credit-risk models this article series addresses throughout.

Figure 4. Governance Architecture Aligning Adaptive/Drift-Monitored ML with SR 26-2's Risk-Based, Materiality-Tiered Model Risk Management

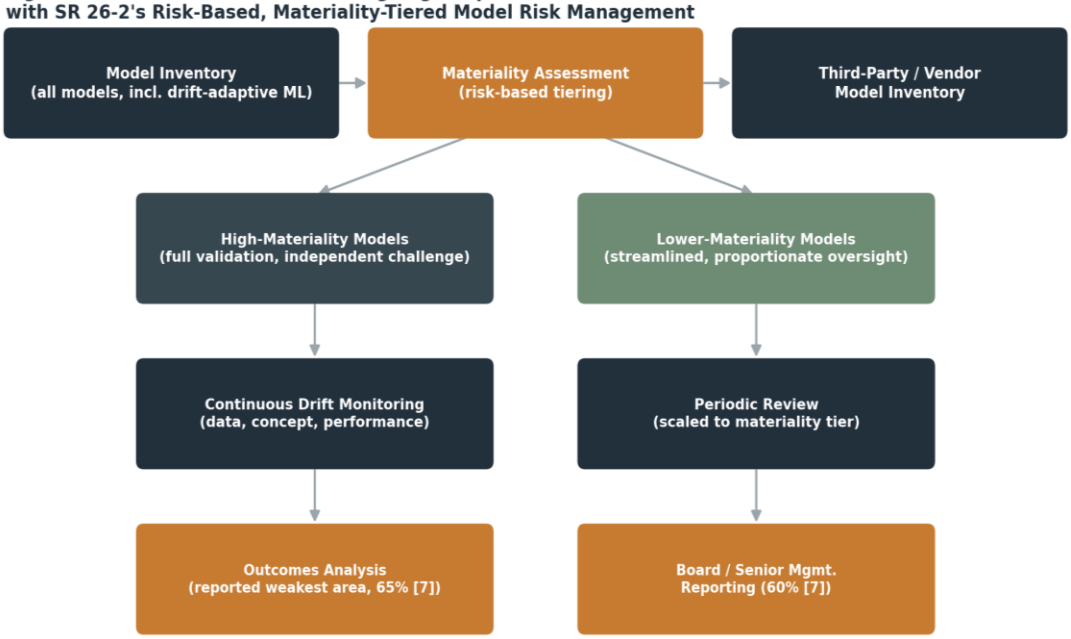

The architecture begins with a comprehensive model inventory, explicitly including adaptive, continuously retraining machine-learning models alongside traditional, more static quantitative models, feeding a formal materiality assessment consistent with Section 3.2's discussion of SR 26-2's new materiality construct. Models assessed as high-materiality, a category that should explicitly include any fraud-detection or credit-risk ensemble whose outputs directly influence customer-facing decisions or regulatory filing determinations, proceed to full validation with independent challenge and, critically for an adaptive model specifically, continuous drift monitoring covering data drift, concept drift, and performance drift, the specific monitoring categories detailed in companion articles in this series on automated drift-detection frameworks. Lower-materiality models proceed instead to a more streamlined, periodic review process, scaled explicitly to the model's assessed materiality tier rather than applied uniformly.

Both paths converge on outcomes analysis and board or senior-management reporting, the two components Section 6 documents as the weakest areas of compliance under the prior SR 11-7 framework specifically. This convergence is a deliberate design choice: regardless of which materiality tier a given model falls into, or which specific monitoring cadence applies to it, the underlying obligation to analyse actual outcomes against expected model

performance, and to report meaningfully on that analysis to accountable governance bodies, does not diminish, even where SR 26-2's new materiality construct otherwise permits a lighter validation and monitoring touch for lower-materiality models.

Illustrative Materiality Criteria for Adaptive Fraud-Detection Models

Table 1 proposes illustrative criteria an institution might use to classify an adaptive fraud-detection or credit-risk model's materiality tier, synthesising the general materiality factors SR 26-2 and its early commentary identify, model complexity, prevalence of use, and significance of the underlying use case, applied specifically to the adaptive, drift-monitored model context this article series addresses.

Table 1. Materiality Criteria for Adaptive Fraud-Detection and Credit-Risk Models

Materiality Factor	High-Materiality Indicator	Lower-Materiality Indicator
Decision authority	Model output directly triggers a customer-facing action (transaction hold, account restriction) or a regulatory filing (SAR)	Model output informs internal prioritisation only, with no direct customer or filing impact
Model complexity	Ensemble or deep-learning architecture with limited inherent interpretability	Simple, largely interpretable scoring rule or single classical model
Adaptivity	Continuously or frequently retrained on new data, subject to ongoing concept drift	Static or infrequently updated, validated once and rarely retrained
Portfolio / transaction volume affected	Applied across the institution's full transaction or account base	Applied to a narrow, limited subset of activity
Vendor dependency	Vendor-supplied with limited institutional visibility into training data or retraining practice	Internally developed with full institutional visibility and control

An adaptive fraud-detection ensemble of the kind discussed throughout this article series will typically score as high-materiality against most or all of these illustrative criteria simultaneously, directly influencing customer-facing decisions, employing complex ensemble architectures, adapting continuously in response to drift, operating across an institution's full transaction base, and, for many community institutions specifically, being vendor-supplied with limited internal visibility into the vendor's own training and retraining practices. This clustering of high-materiality indicators is itself a useful diagnostic: an institution that finds its adaptive fraud-detection model scoring as high-materiality across nearly every dimension in Table 1 should treat this as confirmation that full validation and continuous drift monitoring, rather than the streamlined, lower-materiality path in Figure 4,

is the appropriate governance track, regardless of the institution's own overall asset size relative to SR 26-2's new $30 billion threshold discussed in Section 3.1.

Vendor and Third-Party Model Accountability

A specific point of continuity, rather than change, between the superseded guidance and SR 26-2 deserves emphasis given its particular relevance to community financial institutions, which frequently rely on vendor-supplied models for functions including fraud detection, credit scoring, and current expected credit loss (CECL) estimation. SR 26-2 reinforces that institutions remain fully responsible for model risk even when using third-party solutions, and that reliance on a vendor does not transfer that responsibility to the vendor itself [6]. This principle directly reinforces the vendor-accountability discussion in companion articles in this series: an institution using a vendor-supplied, adaptive fraud-detection ensemble cannot treat that vendor's own internal validation and drift-monitoring practices as automatically satisfying the institution's own model risk management obligations under SR 26-2; the institution itself must maintain sufficient independent visibility into, and documentation of, how that vendor model is validated, monitored, and, where applicable, retrained in response to detected drift.

This is a particularly consequential continuity for community banks and credit unions specifically, given that these institutions, per adoption patterns discussed in companion articles in this series, are disproportionately likely to rely on vendor-supplied rather than internally developed adaptive fraud-detection models, precisely because building and maintaining an internal data-science function capable of developing such models is disproportionately costly for a smaller institution. SR 26-2's continued emphasis that vendor reliance does not transfer model-risk responsibility means that exactly the institutions most likely to depend on vendor models are also the institutions that must invest the most, relative to their overall size, in independent oversight capacity, documentation practices, and the kind of drift-monitoring visibility discussed in Section 4, to satisfy their own, non-delegable obligations under the revised guidance.

SR 26-2 does provide streamlined expectations specifically for how banking organisations should approach validation of vendor and third-party models, a change explicitly noted in contemporaneous analysis as distinct from, and more proportionate than, the fuller validation expectations applied to internally developed models [2]. This streamlining is consistent with the broader materiality-tiering logic discussed in Section 3.2: a vendor model's validation intensity, like an internally developed model's, should scale with that model's materiality to the institution's actual risk profile, rather than requiring the institution to independently replicate the vendor's own full internal validation effort regardless of the model's actual materiality.

Documented Compliance Gaps Under the Prior Guidance

Before turning to the specific compliance data, it is worth situating why examination-trend evidence on the prior guidance's weakest areas matters directly for institutions transitioning to SR 26-2, rather than being merely of historical interest now that SR 11-7 has been superseded. SR 26-2 does not reinvent the underlying substantive obligations around

ongoing monitoring, outcomes analysis, and governance reporting; rather, as Section 3 has detailed, it recalibrates how intensively those obligations apply based on a model's assessed materiality, while leaving the underlying expectation that outcomes must be analysed and reported meaningfully largely intact for any model, particularly a high-materiality one, that falls within an institution's governance scope. Examination findings on where institutions struggled under the prior framework therefore remain directly predictive of where the same institutions are likely to struggle under the revised framework, unless they take deliberate, specific action informed by that prior experience.

Independent examination-trend analysis covering 2024 and 2025, the final two years before SR 26-2's issuance, found that board reporting and outcomes analysis remained the two weakest components of SR 11-7 compliance across U.S. banks, with reported compliance rates of approximately 60 percent for board reporting and 65 percent for outcomes analysis [7].

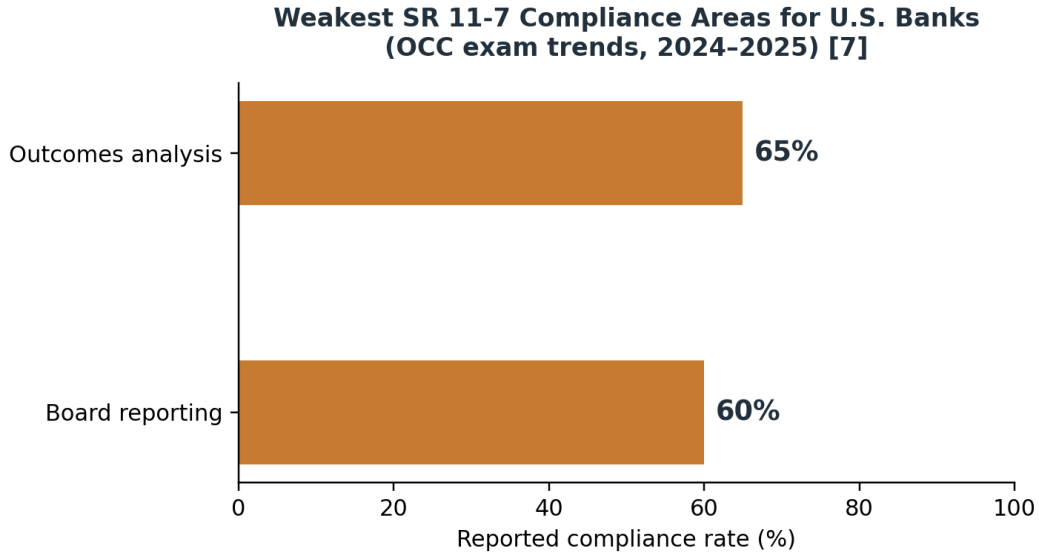

Figure 3 presents these figures directly. This finding carries specific weight for the drift-monitoring architecture this article proposes, since both weak areas relate directly to the ongoing, temporal dimension of model risk management, understanding how a model has actually performed over time (outcomes analysis) and ensuring that understanding reaches accountable decision-makers (board reporting), as opposed to the more static, point-in-time aspects of model validation (initial conceptual soundness review, for example) that examination trends suggest institutions have more consistently satisfied. This pattern is consistent with, and lends independent, cited support to, this article series' recurring argument that ongoing monitoring, not initial validation alone, is where genuine model-risk-management effort is most needed and, per this evidence, most commonly falls short in current practice.

Given that SR 26-2 explicitly retains outcomes analysis and reporting as continuing expectations even as it introduces greater proportionality elsewhere, per the governance

architecture in Figure 4, institutions should not read the compliance gaps documented in Figure 3 as evidence that these components have been deprioritised under the new guidance; if anything, the persistence of these specific gaps across the guidance transition suggests they warrant particular institutional attention precisely because they have proven difficult to satisfy consistently under fifteen years of the prior framework, and nothing in SR 26-2's text suggests the underlying supervisory concern motivating these expectations has diminished.

Open Questions for Community Financial Institutions

To make the practical stakes of these open questions concrete, consider a hypothetical community bank with total assets well below SR 26-2's new $30 billion threshold, comfortably outside the guidance's primary applicability range on asset size alone, that has recently deployed a vendor-supplied, adaptive ensemble fraud-detection model of the kind discussed throughout this article series. Under the materiality criteria proposed in Table 1, this model plausibly scores as high-materiality: it directly influences whether transactions are held for review, it uses a complex ensemble architecture, it retrains continuously in response to detected drift, it operates across the institution's entire transaction base, and it is vendor-supplied with limited internal visibility into the vendor's training data. Under SR 26-2's stated principle that materiality and model complexity, not asset size alone, govern applicability, this community bank should reasonably expect its fraud-detection model specifically, even if few or none of its other models, to fall within the guidance's substantive expectations for governance, validation, and ongoing monitoring, notwithstanding the institution's overall size comfortably below the new threshold headline figure most commentary has emphasised.

Several questions remain genuinely unresolved for community financial institutions at the time of this article's preparation, given SR 26-2's very recent issuance. First, and most directly relevant to this article series, no source identified in this article's research provides worked, sector-specific guidance on how a community bank or credit union should apply SR 26-2's new materiality construct specifically to an adaptive, drift-monitored fraud-detection ensemble, as distinct from the more familiar, largely static credit-scoring and CECL models the existing commentary on SR 26-2 more frequently addresses. Second, it remains unclear, absent further examiner guidance or observed supervisory practice, precisely how the "activities outside the scope of traditional community banking" carve-back discussed in Section 3.1 will be applied in practice to institutions using sophisticated, vendor-supplied AI/ML fraud-detection tools, since such tools are increasingly common even at quite small institutions, per the adoption patterns discussed in companion articles in this series, potentially in tension with an institution's otherwise traditional overall business profile.

Third, the interaction between SR 26-2's reduced prescriptiveness on annual revalidation, discussed in Section 3.3, and the drift-triggered, event-based monitoring cadence this article series has consistently recommended, remains to be tested in actual supervisory practice: while the direction of this change appears supportive of exactly the kind of continuous, risk-proportionate monitoring this article recommends, institutions should expect that examiners will need time to develop consistent practice in evaluating event-triggered monitoring regimes against the new guidance's more principles-based, less prescriptive text, and should

document their own monitoring cadence and rationale clearly and defensibly in the interim rather than assuming a specific practice is automatically compliant simply because the new guidance is less prescriptive than its predecessor.

Given these open questions, institutions adopting the governance architecture proposed in Section 4 should treat it as a reasonable, evidence-informed starting point consistent with SR 26-2's stated principles, rather than as a guaranteed safe harbour against future supervisory findings, and should expect to revise their specific implementation as examiner practice under the new guidance becomes clearer over the coming supervisory cycles.

Conclusion

A final practical point deserves emphasis before summarising this article's overall conclusions: the specific evidentiary status of every claim in this article regarding SR 26-2 itself should be understood as reflecting a regulatory development that occurred only shortly before this article's preparation, meaning independent, longer-term analysis of how examiners will apply the new guidance in practice, as opposed to the guidance's own text and the immediate, largely law-firm and advisory-firm commentary reviewed in this article, remains necessarily limited at present. Institutions should treat this article's proposed governance architecture and materiality-tiering approach as a reasoned synthesis of the best currently available evidence, revisited as actual examiner practice under SR 26-2 accumulates over the coming supervisory cycles, rather than as a final, settled interpretation of a guidance document still in its first months of effect. This article has examined SR 26-2, the Revised Guidance on Model Risk Management issued jointly by the Federal Reserve, OCC, and FDIC on April 17, 2026, superseding both the foundational 2011 SR 11-7 guidance and the 2021 BSA/AML-specific SR 21-8 statement, and has proposed a governance architecture aligning continuous, drift-aware monitoring of adaptive machine-learning fraud-detection and credit-risk models with this new guidance's risk-based, materiality-tiered approach. The most consequential change for this article series' consistent focus on community financial institutions is the new uniform $30 billion asset threshold, replacing a considerably lower, FDIC-specific $1 billion threshold and the previous absence of any explicit threshold at the Federal Reserve and OCC, though this article has been careful to note that SR 26-2 remains principles-based rather than purely scope-based, and may still apply meaningfully to smaller institutions with significant model complexity or prevalence. Documented compliance gaps under the superseded guidance, board reporting and outcomes analysis compliance rates of only 60 and 65 percent respectively across U.S. banks in 2024–2025 [7], underscore that the ongoing, temporal dimension of model risk management, precisely the dimension a drift-monitoring architecture of the kind this article series proposes throughout is designed to strengthen, has been where institutions have most commonly fallen short under the prior framework, and nothing in SR 26-2's text suggests this underlying supervisory concern has diminished even as the guidance introduces greater proportionality elsewhere. Given SR 26-2's very recent issuance, institutions should treat the specific governance architecture proposed in this article as a reasoned, well-evidenced starting point rather than a settled, examiner-tested practice, and should expect to refine their implementation as supervisory practice under this consequential new guidance continues to develop over the coming years.

References

[1] Cosma, S., Rimo, G., & Torluccio, G. (2023). Knowledge mapping of model risk in banking. *International Review of Financial Analysis, 89*, Article 102800.

[2] Kou, G., Olgu Akdeniz, Ö., Dinçer, H., & Yüksel, S. (2021). FinTech investments in European banks: A hybrid IT2 fuzzy multidimensional decision-making approach. *Financial Innovation, 7*, Article 39.

[3] Petropoulos, A., Siakoulis, V., Stavroulakis, E., & Vlachogiannakis, N. (2020). Predicting bank insolvencies using machine learning techniques. *International Journal of Forecasting, 36*(3), 1008–1022.

[4] Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance, 34*(11), 2767–2787.

[5] Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance, 77*(1), 5–47.

[6] Hurlin, C., Leymarie, J., & Patin, A. (2018). A survey of credit risk modelling methods. *Journal of Economic Surveys, 32*(5), 1351–1388.

[7] Misheva, B. H., Jaggi, D., Posth, J.-A., Gramespacher, T., & Osterrieder, J. (2021). Audience-dependent explanations for AI-based risk management tools: A survey. *Frontiers in Artificial Intelligence, 4*, Article 794996.

[8] Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable AI in credit risk management. *Computational Economics, 57*, 203–216.

[9] Giudici, P., & Raffinetti, E. (2021). Explainable AI methods in financial risk management. *Statistical Methods & Applications, 30*, 111–132.