July 31, 2026

Office of the Secretary
Public Company Accounting Oversight Board
1666 K Street NW
Washington, DC 20006-2803

Re: PCAOB Release No. 2026-005 — Request for Public Comment on PCAOB Standard Setting

Members of the Board:

I write as a public-finance and audit practitioner with over a decade of cross-sector experience. This comment addresses the Board's request for input on its general approach to standard setting and on the firm and engagement performance metrics item. It does not address NOCLAR.

The distinction I ask the Board to draw

A requirement to disclose produces a disclosure. It does not produce a true disclosure. Where verification is expensive, a false disclosure is cheap to make and durable once made, and the cost of discovering it falls on whoever is willing to spend months looking.

I urge the Board to treat verifiability as a design property separate from disclosure, and to prefer it wherever a mechanical check is possible. A disclosure requirement asks a firm to state something. A verifiability requirement asks that the statement be checkable by someone who does not work for the firm, does not trust the firm, and holds only what the firm has published. As the Commission weighs permitting semiannual in place of quarterly reporting — an impact this release itself asks about — verification rises in importance as disclosure frequency falls: fewer statements, more weight on each.

Why the distinction matters

The theoretical case is long-established and settled. Grossman (1981) and Milgrom (1981) showed that where disclosure is costless and verifiable, every type discloses in equilibrium: silence is priced at the worst case, so withholding is worse than revealing. In that model the mandate is not doing the work; the verifiability is.

The later literature explains why real firms do not unravel. Verrecchia (1983) supplies proprietary cost: a firm withholds because the disclosure would help its competitors. Dye (1985) and Jung and Kwon (1988) supply uncertainty about whether the firm holds the information at all, so silence stops being informative.

There is a third friction, and it is the one a standard-setter can remove. Where a claim cannot be checked without substantial independent effort, silence and false statement look alike to a reader,

and the equilibrium that would otherwise compel disclosure never forms. A firm does not need a proprietary-cost story to withhold or to shade; it only needs the reader to be unable to tell.

Proprietary cost is a real constraint the Board cannot legislate away. Uncertain endowment is a fact about the world. Verification cost is a design choice, and it is inside the Board's remit.

The Board's own evidence: thresholds become levels, in both directions

The empirical record on the Board's inspection regime supports a sharper claim than "disclosure regimes get gamed." It supports convergence: a published threshold becomes the operating level from both sides.

Aobdia (2018) found that auditor effort increases on engagements receiving a Part I finding, with spillover to non-inspected engagements of the same partner or office, and that effort and financial reporting quality subsequently decline on inspected engagements that received no finding. Firms and clients converge on the level the Part I bar sets. Abbott and Buslepp (2021) reach the companion result by a different method: firms learn the regime and adjust opportunistically toward its threshold. Two studies, two methods, one finding: the measure becomes the ceiling as well as the floor.

The improvement evidence is real and belongs beside it. Blann, Kleppe, and Shipman (2023) show audit quality improved significantly where PCAOB office expansion changed real oversight exposure, with triennially inspected firms most affected. Blankley, Hong, Kerr, and Wiggins (2014) and Alam, Cheng, Rickett, and Skomra (2024) document responses in remediation and reporting behavior. Carcello, Hollingsworth, and Mastrolia (2011) find early inspection-regime effects, scoped to Big 4 firms; that scoping matters when generalizing.

Read together: oversight that changes what is checkable changes behavior durably; oversight that publishes a threshold invites convergence on the threshold. That is the case for preferring verifiability mechanisms over disclosure-plus-bar mechanisms wherever the Board can specify one.

Recommendation: specify the computation and the input commitment

The Board has recent, direct experience with this item. Firm and engagement metrics rules adopted on November 21, 2024 were withdrawn from Commission consideration on February 11, 2025, after comment concentrated on reporting burden and its weight on smaller firms. I take the withdrawal as an opening: a re-approach can build verifiability in at design time, and in one specific sense it asks less of firms than the withdrawn design did. An input commitment binds figures a firm has already computed. It adds no new metric, no new threshold, and no new reporting category; it adds only the property that what is later exhibited can be shown to be what was held.

For the metrics item, I urge the Board, wherever it requires a metric, to specify in the same requirement: (1) the computation, exactly, so that two parties holding the same inputs derive the

same number; and (2) an input commitment — a cryptographic hash of the metric's inputs, fixed at a stated time and published or retained under audit trail — so that the inputs a firm later exhibits can be proven to be the inputs it committed to when the metric was reported. Where a metric cannot be specified to that standard, saying so plainly is better than a requirement that produces confident, unfalsifiable reporting.

The machinery is ordinary. Adjacent regimes already require the underlying property of records whose integrity can be demonstrated after the fact: activity recordkeeping under MiFID II, logging obligations under Article 12 of the EU AI Act, integrity controls under ISO 27001, SOC 2, and DORA. None of these mandates a published hash commitment; each mandates tamper-evident record integrity, which the same construction is routinely used to implement. The increment asked for here is small and specific: fix the metric's inputs at a stated time, so that the record a firm later exhibits can be proven to be the record it held. The marginal cost is a hash construction. The marginal benefit is that restatement becomes detectable rather than deniable.

Input commitments are also threshold-neutral. They publish no bar for behavior to converge on and bind only what the firm itself claimed, so they add no new ceiling of the kind the inspection literature documents.

A working demonstration at small scale

I can answer the practicality question from operation rather than argument, because I run a working instance, deliberately small so that its costs are legible.

I maintain a public forecasting record under a pair of self-issued standards written to be citable by paragraph and testable by third parties. Under the current construction, every scored input — the probability asserted and the deadline it is asserted over — sits inside a salted SHA-256 commitment sealed before any outcome exists. The log of commitments is append-only in public version-control history, and the published state is timestamped by an RFC 3161 authority whose token I cannot amend or revoke. Two verifiers, written against the standard library alone and with no knowledge of my desk, accept a URL from anyone and return a conformance verdict with paragraph citations: whether every seal recomputes, whether any pre-registered figure changed between two snapshots of the record, and whether the published counts derive from the records rather than from my say-so.

Two properties bear on the Board's question. The first is cost. The entire verifiability layer — the hash construction, the anchoring, and both verifiers — is a few hundred lines of code and required no license, no firm, no budget, and no trust in the operator. The marginal cost of a checkable record over a merely stated one was a design decision, not an infrastructure program. If one desk can make its numbers verifiable by an untrusting third party at that price, firm-level metrics reporting can.

The second is what the discipline does to its own author. This month, an adversarial review of my own construction found that its commitment originally excluded the two inputs every score is

computed from: a sealed prediction proved I had said a thing, and proved nothing about the confidence or horizon being scored, so a probability could be restated after sealing while every hash still verified. The repair was not to quietly fix the record. The construction was version-bumped forward; the legacy entries were disclosed on the record's face as supported by anchored history rather than by commitment; the finding was printed with paragraph numbers and an external citation; and the verifier gained the check that makes a restated figure fail as loudly as an altered hash. A record that names its own defects, dated, where the reader cannot miss them, is the only kind of self-maintained record worth reading, and a reporting standard can require exactly that behavior.

The record, its standards, its findings, and both verifiers are public at https://retroprescientaudit.com/ and in the repository the record's anchor declaration names. I invite the Board's staff to run the verifiers rather than take my description of them.

Disclosure without binding: a current example at standards scale

The first comprehensive independent security analysis of C2PA, the content-provenance specification an industrial coalition is deploying against synthetic media, including the first formal-methods analysis of its core protocols, found that the specification fails its own stated security goals (Golaszewski, Krawetz, Sherman, Zieglar, et al., Cryptology ePrint Archive 2026/804, analyzing Version 2.2; the coalition reports partial fixes in Version 2.3). The most consequential finding: nothing in the signed data references the trusted timestamp, so the date a validator displays can be removed or replaced without detection — a field on the face of the record, outside the material the seal protects. The analysis also found conforming validators returning inconsistent results, and a conformance program that certified products without technical review or defined requirements. These are not exotic attacks. They are the ordinary consequence of specifying what must be stated without specifying what binds it. A standard-setter that requires a metric can, in the same sentence, require the input commitment that makes the metric checkable.

Closing

Standards codify practice; they do not precede it. A standard that requires a claim nobody can check codifies the claim rather than the practice. Where the Board can specify a computation and an input commitment so that a party holding the inputs derives the same number, it should. And whatever threshold accompanies it, the Board's own evidence suggests the threshold will become the level, in both directions.

I appreciate the Board's consideration.

Respectfully submitted,

NebelKrähe

Operator, the Retro-Prescient Audit Desk
(the pseudonym under which the record, its standards, and its copyright are published)
audit@retroprescientaudit.com
https://retroprescientaudit.com/