TRAICEBACK INSIGHTS

AI Governance at the Speed of Adoption

Why the emerging record-keeping risks in AI belong on the board agenda — and why a well-built governance framework accelerates adoption rather than restraining it.

JULY 2026

Two articles published within days of each other in July 2026, by authors approaching the problem from different directions, arrive at substantially the same conclusion. Neither is about model performance, algorithmic bias, or any of the topics that have dominated AI governance discussion for the past three years. Both are about paperwork.

Writing in the CLS Blue Sky Blog, Patrick Meson identifies what he calls the documentation imperative: the observation that a company's protection against legal liability rests on the quality of the record it can produce. Directors defend oversight claims with minutes and charters. Issuers sustain the securities due-diligence defense with the documented trail of their investigation. The contemporaneous record, in Meson's formulation, is not merely evidence of diligence — it has become the condition of legal protection. For regulated firms the same holds on the regulatory side, where the examination record and the compliance file perform an identical function.

The paradox Meson then identifies is that the AI tools now spreading through ordinary workflows generate records of their own, automatically and at scale. A single meeting attended by an AI notetaker produces a recording, a transcript, a summary, an action list, a prompt history, and several drafts where the older practice produced one set of approved minutes. These machine records are discoverable, frequently inconsistent with the company's formal account, and capable of waiving privilege. The law presses companies to document more at precisely the moment their tools have begun producing records that no one decided to create.

Jeff Kelly and Scott Sherman, writing in Law.com's Corporate Counsel, reach the same terrain from the adoption side. Most companies, they observe, never made a single board-approved decision to deploy AI. It arrived one feature, one vendor, and one employee at a time — a meeting platform that began offering summaries, a director who used a consumer chatbot to condense a board packet, a finance team that switched on an assistant during the close. By the time AI reaches the board agenda, it is already embedded in how the company creates and retains information. Their practical guidance runs to five steps, and the standard they set at the close is exacting: a workable program lets the company reconstruct a consequential use — who chose the tool, what went in, what came out, who relied on it, and what record remains.

Read together, the two pieces describe one problem stated twice. The exposure is no longer a failure to document. It is a failure to decide, in advance, what the documentation is.

The Direction of Travel

What makes these articles worth the attention of chief executives and directors, rather than only of counsel, is what they imply in combination. The legal community is converging — through commentary, through early rulings on privilege and discovery, and through statutes that increasingly reward documented governance with affirmative defenses, as the Texas Responsible AI Governance Act now does for substantial compliance with a recognized risk management framework — on a standard that no firm can satisfy retroactively.

That last point deserves emphasis. An evidentiary record cannot be reconstructed after the fact. A firm that has not decided which of its AI-generated artifacts constitutes the official record will not be able to produce one under examination or in discovery; it will produce all of them, including the drafts it never adopted, the transcripts it never reviewed, and the prompt logs it never knew it was keeping. The remedy is not available at the moment it is needed. It has to have been built ex ante.

This is the sense in which a defensible AI governance framework has stopped being a matter of good practice and become a practical requirement. Not because a regulator has mandated one in so many words, but because the protections firms already rely on — the oversight defense, the diligence defense, the reasonableness of a compliance program — have all come to depend on an evidentiary record that ungoverned AI adoption quietly degrades.

Governance as Accelerator

The natural reading of both articles is cautionary, and it invites a conclusion that deserves to be resisted: that governance is a brake on AI adoption, a set of constraints that trades speed for safety. The opposite is closer to the truth, and there is a useful precedent.

The Wright brothers flew in 1903. Commercial aviation remained a marginal business for more than two decades afterward — not because the technology was inadequate, but because no passenger could distinguish a competent operator from a reckless one, and no investor could underwrite the difference. The era’s barnstorming was thrilling, and it was destroying the industry’s commercial credibility. What changed was the Air Commerce Act of 1926, which directed the Secretary of Commerce to license pilots, certify aircraft airworthiness, and establish federal airways.

The detail usually omitted from this history is the one that matters most. The Act was passed at the industry’s own urging. Aviation operators had concluded that they could not attract passengers or capital without a credible safety regime, and that federal certification was the precondition for their business becoming investable rather than a constraint upon it. They were right. Accident rates fell, the airmail routes passed to private carriers, and substantial capital moved into airports and navigation infrastructure. Regulation did not slow the industry down. It was the thing that allowed the industry to exist at commercial scale.

The parallel is not that firms should wait for a mandate. It is that the operators recognized, before any regulator compelled them to, that ungoverned operation was capping their own growth.

Kelly and Sherman's title observes that AI moves faster than governance. It does — but it does not have to, and the reason it usually does is instructive. Ungoverned adoption is fast only until the first subpoena, the first examination request, the first privilege dispute. At that point it becomes extraordinarily slow: the firm must reconstruct what it did, discover it cannot, and then negotiate from a position of evidentiary weakness. The apparent speed was borrowed against a liability that comes due later, with interest.

A firm that has resolved its governance questions in advance is genuinely faster. It knows which deployment postures are approved and what evidence each one produces. It knows what its record is and what is disposable. It knows who may deploy which class of tool against which class of data. When a new capability arrives, the question is not whether to allow it but which existing posture it falls into — a question answerable in days rather than quarters. An artificial intelligence governance framework does not slow adoption down; rather, it is what makes rapid adoption of AI tools possible.

This depends on a framework built to be iterative and dynamic rather than static. A policy set adopted once and filed is quickly overtaken because the technology it tries to address will not hold still. AI governance is not a one-and-done exercise, and a framework that assumes otherwise will lag adoption no matter how carefully it was drafted. Built and operated properly, an AI governance framework can absorb any pace of adoption, because it governs classes of use rather than individual tools.

Certification did not make aircraft fly. It made them financeable, investable, and accessible to the public. The same logic applies to the firm that can demonstrate, rather than assert, that its AI use is governed.

What the Guidance Leaves Open

Both articles establish what must be provable. Neither fully addresses how a firm makes it provable, and the honest answer is that this cannot be resolved by a checklist.

A framework that holds up runs in one direction. The business defines the risks and opportunities: what the firm actually does, for whom, with what data, subject to which regulatory obligations and under what competitive conditions. The risks and opportunities determine the controls — the policies and procedures that address those specific exposures. And the controls produce the evidence, as a byproduct of operating rather than as a separate documentation exercise. Each link is derived from the one before it.

Run that chain backward and the failure mode becomes obvious. Evidence that no control produced is an artifact, not proof of anything; it may show that something happened, but not that anyone intended it or reviewed it. A control that no identified risk called for is a policy without a purpose, and it will be observed inconsistently because no one can say what it is for. And a risk assessment that did not begin with the business it purports to describe is describing someone else's firm. This is why the policy copied from another organization (or generated

by an AI chatbot) fails in a way that is invisible until it is tested: it was derived from a different set of risks, arising from a different business, or from a large language model pre-trained on data pre-dating current conditions, and the evidence it generates proves the wrong things.

This chain is also what makes AI model deployment posture a governance question rather than a technical one. Kelly and Sherman are right that an enterprise label on a vendor agreement is not sufficient assurance, and that training rights, retention, subprocessors, and exportability all require examination. But this framing assumes the AI in question is a third-party service being configured. Increasingly it is not. A firm running open-weight models on its own hardware faces a materially different profile: there is no vendor to negotiate with, no subprocessor to diligence, and no egress at all — but it inherits questions about model validation and data provenance and integrity that simply do not arise in the vendor case. A firm augmenting a commercial model with its own document corpus in a retrieval-augmented generation (or "RAG") pipeline sits somewhere between the two, with exposures peculiar to that arrangement.

These are different risks, so they call for different controls and different decisions, so they yield different evidence. Guidance written for one case does not reach the others. A governance framework that is going to hold has to distinguish among them systematically, so that the evidentiary obligations attaching to a given use follow from how it is deployed rather than from an ad hoc judgment made at the point of adoption.

The last link deserves particular attention, because it is where most programs can quietly break. Kelly and Sherman's closing standard — reconstruct who chose the tool, what went in, what came out, who relied on it, what record remains — is a demand for provenance, and provenance is a technical discipline before it is a legal one. It requires that artifacts be hashed at creation, that audit logs be tamper-evident, that the chain from source authority to generated output be traceable and verifiable months later by someone who was not present. Most firms reading that standard will recognize its force but have no idea how to satisfy it. The gap between agreeing with the principle and being able to manifest compliance with evidence is where the actual work sits.

This difficulty is about to sharpen. The record-keeping problem both articles describe arises when AI assists a person — a notetaker transcribes a meeting someone convened, an assistant summarizes a document someone supplied. The move toward agentic systems, which take sequences of actions across tools without a human reviewing each step, removes the person from the moment of creation altogether. When the answer to who chose the tool and who relied on the output is another system rather than an employee, the reconstruction Kelly and Sherman require becomes materially harder — and it becomes harder precisely for the firm that adopted the capability without first deciding how its actions would be recorded. A framework built to be extended can absorb this. A static one cannot, and the distance between the two is widening.

Why This Belongs to the Board

Kelly and Sherman allocate AI oversight to the board or a committee, with an accountable management owner. That is correct as far as it goes, but oversight is a review function — the board satisfies itself that management has done the work. The argument advanced here is stronger.

The board's duty of oversight is not a new idea, and it is not a soft one. Since *In re Caremark*, the law has held that directors must make a good-faith effort to put in place a reasonable system for monitoring and reporting the risks that matter to the business — and a line of more recent Delaware decisions has shown these claims can survive dismissal where the risk is central to the enterprise. The lesson those cases press is directly on point here: a board discharges the duty by ensuring a monitoring system genuinely exists and functions, not by pointing to minimal regulatory compliance after the fact. Oversight of how the firm creates and keeps its records, in an environment where its tools now generate those records autonomously, is precisely the kind of system a board is expected to assure.

If a governance framework determines the pace at which a firm can safely adopt AI, then it is a strategic instrument rather than a compliance artifact, and it belongs to the chief executive and the board as a matter of ownership rather than review. Treating it as something that lives in the compliance or technology function guarantees that it will lag adoption, for a structural reason: neither function controls the adoption decisions. Business units adopt. Compliance and technology discover, respond, and remediate. A framework owned by the responders will always be reconstructing a state of affairs it did not authorize.

There is a corollary that senior management should take seriously. Chief compliance officers and chief information security officers are the first line of defense against precisely the risks these articles describe, and they are frequently the only people in an organization who can see the exposure clearly. But they are rarely positioned to compel the resource allocation that addressing it requires, and a request that originates from compliance or security is easily received as a departmental ask rather than an enterprise priority.

Boards should not wait to be told. The emerging risks described in these two articles — the discoverability of machine-generated records, the fragility of privilege when tools are introduced without counsel's involvement, the preservation obligations that attach to artifacts a firm did not know it was creating — are not technical details to be escalated upward when they become acute. They are governance questions of the kind boards exist to address, and a board that understands them well enough to lead the effort will get a materially better outcome than one that waits to evaluate a budget request.

The Question Worth Asking

Directors and chief executives need not approve every AI use or master every model. Kelly and Sherman are right about that. But there is a short set of questions that a board should be able to answer without consulting anyone, and the answers are diagnostic.

Every firm has its own barnstorming problem: the free accounts, the embedded features nobody switched off, the consumer chatbot used to condense a board packet. It is worth asking what the firm actually knows about it.

1. Do we know where AI is being used across the firm — including the free accounts, the embedded features, and the tools our outside advisers use on our matters?
2. Do we know what controls govern that use, and can we demonstrate that those controls operate rather than merely exist on paper?
3. Could we reconstruct a material AI-assisted decision six months from now — to a regulator, to an adversary in litigation, or to ourselves?

A firm that cannot answer the third question is not, in any meaningful sense, governing its AI use. It is documenting it, at volume, without curation — which is the precise condition Meson identifies as the source of the exposure rather than the protection against it.

The evidentiary record has always been the measure of diligence and the proof that business judgment was actually exercised by the board. The tools are now producing that record whether or not anyone has decided what it should be. Closing that gap is not a matter of documenting less, or more. It is a matter of deciding, in advance, what the firm's record is. The best defense against a claim of negligence has always been the manifestation of due care — and a defensible AI governance framework is how a firm manifests it.

Traiceback Insights examines the convergence of AI governance, cybersecurity, fiduciary oversight, and evidence-based accountability.

Gordon Eng is the founder and Managing Director of Traiceback Solutions LLC. He is a former general counsel and chief compliance officer, and a member of the New York and Connecticut bars.