New State AI Laws Create Dual Misrepresentation Risk

By **William Pao, Tijana Brien and Sean Quinn** (June 11, 2026)

William Pao

Tijana Brien

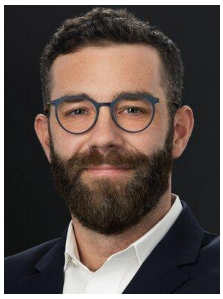

Sean Quinn

Artificial intelligence transparency laws are sweeping the nation. And every compliance record they require — every transparency report, training data summary and incident disclosure — may now be measured against what a company tells its investors, its customers and the public. When those records diverge, the result is litigation and regulatory risk.

The pace of AI disclosure regulation has accelerated sharply. In the past few months alone, several states have enacted or amended AI disclosure laws. The result: Companies must now produce an expanding body of AI-specific disclosures — frameworks, transparency reports, training data summaries, consumer disclosures, regulator-facing catastrophic-risk submissions and incident reports — directed at different audiences, on different timelines, in different formats.

Those records do not exist in a vacuum. They sit alongside U.S. Securities and Exchange Commission filings, earnings calls, investor decks, website claims, and marketing materials. When the two sets of statements diverge, the company faces a dual misrepresentation risk — exposure under AI-specific statutes on one side, and securities, fraud or consumer protection law on the other.

The underlying liability theory is not new; anti-fraud and consumer protection law already punish false or misleading AI-related statements. What is new is the volume and specificity of the compliance record that AI statutes will generate — and the number of points at which that record can be compared against a company's public narrative.

This is not a future problem. Cornerstone Research identified 16 AI-related securities class action filings in 2025 alone, a 7% increase from 2024 and a 129% increase from 2023.[1]

This article does four things: (1) surveys the AI-specific disclosure duties taking effect; (2) examines enforcement actions and securities litigation — including the General Motors Co./Cruise LLC case — that show how cross-record inconsistency has already drawn scrutiny; (3) analyzes the ways those records can diverge from other public statements; and (4) sets out practical steps to manage the risk.

The Emerging AI Disclosure Landscape

Several recently enacted or pending state laws illustrate the emerging regime's two-sided nature: Each generates records — some public-facing, some regulator-facing — that can later be measured against a company's other statements.

New York

New York's amended Responsible AI Safety and Education, or RAISE, Act — effective Jan. 1, 2027 — illustrates the two-sided disclosure regime.

On the public-facing side, the statute requires frontier model developers to adopt and publish an AI safety framework, review it annually, and publish any material modification and its justification within 30 days. It also requires a transparency report before, or concurrently with, deployment, which may be satisfied by a system card.

On the regulator-facing side, the RAISE Act requires confidential summaries of catastrophic risk assessments every three months, or on another reasonable schedule; reporting of critical-safety incidents within 72 hours; and a 24-hour escalation where there is imminent risk of death or serious physical injury — which are explicitly prohibited from being materially false or misleading.

By comparison, California's S.B. 53 allows 15 days for reporting of critical-safety incidents — underscoring that differing timelines across states compound compliance burdens.

California

California's AI Training Data Transparency Law, or A.B. 2013, creates a purely public-facing disclosure record focused on training-data provenance. It requires covered developers of generative AI systems to post documentation online about training data — including dataset sources, whether the data includes copyrighted or personal information, and collection dates.

On March 4, a U.S. District Court for the Central District of California judge denied xAI's motion to preliminarily enjoin the statute, reasoning that consumers could use such information to assess model quality and reliability. The ruling does not resolve the merits, and xAI is appealing.

The resulting disclosure creates a reference point that plaintiffs or regulators could later compare against a company's public claims about data sourcing, licensing or model suitability — making cross-channel vocabulary alignment essential.

Connecticut

Connecticut's recently enacted S.B. 5 regulates AI across multiple domains — including consumer disclosures, employment decision notices and synthetic-content labeling.

The provision most relevant here requires large frontier developers to prepare and submit quarterly reports to their officers and directors summarizing all anonymized internal reports of catastrophic AI risk.

That requirement creates an indirect dual misrepresentation exposure: If a company's public statements — in SEC filings or at investor conferences — characterize AI-related risks in terms inconsistent with what those internal reports reflect, and a catastrophic AI event later causes a stock drop, a plaintiff could use the reports to allege the company knew its public risk characterizations were false when made.

Other State Laws

Other new state laws create additional litigation vectors — by state attorneys general or private plaintiffs invoking existing consumer protection and anti-fraud statutes — any of which may also trigger securities litigation exposure.

For example, both California's AI Transparency Act, S.B. 942, and Washington's H.B. 1170 impose provenance and transparency obligations on providers of widely used generative AI systems.

H.B. 1170 requires certain providers of publicly accessible generative AI systems with more than one million monthly users in Washington to embed provenance data in AI-created or materially altered image, audio and video content. It also requires government agencies offering AI to consumers to disclose that the consumer is interacting with AI. The law is enforceable by the attorney general under the state Consumer Protection Act and is effective Feb. 1, 2027.

Utah's mental health chatbot law, H.B. 452, enacted in June 2025, requires suppliers to disclose that the chatbot is AI — not human — at multiple interaction points. A company that markets its product, say, as a wellness tool while omitting those disclosures violates state consumer fraud and deceptive practices statutes.

Enforcement Precedents

Recent cases and federal enforcement actions show how mismatches across a company's records can fuel both securities litigation and regulatory scrutiny — and the new incident reporting requirements discussed above are likely to sharpen that exposure.

In December 2023, investors sued GM, Cruise and several executives after a Cruise autonomous vehicle struck a pedestrian and then dragged her roughly 20 feet. The suit, Shamoon v. GM, filed in the U.S. District Court for the Eastern District of Michigan, alleged that Cruise and GM had misled investors about both the technology and the accident. In March 2025, the court denied the motion to dismiss in part and allowed key claims to proceed.

Two aspects of the opinion matter most here.

First, the court drew a sharp line between two kinds of autonomy claims. Cruise executives had repeatedly described the company's vehicles as both "fully driverless" and having "Level 4 autonomy." Cruise argued that the terms were interchangeable. The court found that they were not.

"Level 4" is a term defined by the SAE Taxonomy of Driving Automation — and according to the court, it carried a specific legal meaning. Whereas the "fully driverless" statements were not actionable — reasonable investors would have understood them to mean only that no safety driver sat in the car — the "Level 4" label had a technical definition, which gave the court a benchmark against which to measure the company's statements.

The lesson is clear: Once a company adopts terminology tied to a recognized standard, it cannot then ignore that standard. What might otherwise pass as puffery becomes a testable factual claim.

Second, the accident exposed a cross-record problem. After the October 2023 collision,

Cruise publicly suggested the vehicle stopped on impact. It did not disclose that the vehicle dragged the pedestrian.

But federal and state reporting obligations required Cruise to describe what actually happened — and the full account reached the public only after the California DMV suspended Cruise's permit. The court used that regulatory record as a benchmark, compared it against Cruise's public statements and found the gaps sufficient to plead both falsity and scienter.

For AI companies, the lesson is critical: Once a critical incident generates mandatory regulator-facing reports, those reports may become a benchmark against which courts compare the company's public narrative. Mismatches across that record can support both falsity and scienter allegations in a securities claim.

Federal regulators have independently signaled that AI-related consumer representations are subject to active scrutiny under existing law.

In September 2025, the Federal Trade Commission launched a compulsory inquiry under Section 6(b) of the FTC Act — which lets the agency demand detailed business practice information without alleging a specific violation — targeting companies that offer AI-powered companion chatbots. The inquiry asked what those companies have told users and parents about features, intended audiences, negative impacts and data practices.

In addition, California's attorney general has said that existing state law applies to entities developing and using AI. The Consumer Financial Protection Bureau has emphasized that creditors using AI must still provide specific and accurate adverse-action reasons.

These are just a few examples. And any enforcement action by these regulators is likely to draw scrutiny from the securities plaintiffs bar as well.

The Dual Misrepresentation Risk: How AI Law Records Conflict With Other Disclosures

These AI disclosure laws create exposure on multiple fronts — securities litigation, regulatory enforcement and consumer protection claims.

Scope Mismatch

AI law disclosures are likely to be more specific than public statements because they are drafted to satisfy specific regulatory requirements, covering defined model types, identified risk categories and specified deployment contexts. Public statements, by contrast, tend to be broader, more aspirational and often untethered from the precise scope of the underlying compliance document.

The gap between those two levels creates a scope mismatch risk.

A company's AI transparency report may describe an incident reporting structure that applies to only one model family. If the company's CEO later tells investors the company has robust incident reporting structures across all AI systems built by customers on its model family for any significant incident, that statement may be challenged as false or misleading because the incident reporting structure in the AI disclosure is more narrowly limited to the model level and not all AI systems built on those models.

Companies should review AI-related investor statements for scope words like "all," "every," "enterprisewide," "global" and "fully." When paired with specific legal frameworks and disclosure regimes, they risk converting a general governance message into a testable statement.

Terminology Mismatch

The AI industry has not converged on standard definitions for governance and transparency concepts. The new statutes use differing terminology. In securities litigation, courts treat statutory or industry-standard definitions as benchmarks against which challenged statements are measured. Companies must therefore ensure that their use of a technical term aligns with its meaning under applicable law and recognized industry frameworks.

The Cruise case illustrates the point. As discussed above, the court distinguished between "fully driverless" and "Level 4" autonomy, and held the Level 4 statements actionable.

The lesson: Until the industry settles on uniform definitions, companies must track the terminology in each applicable AI statute and maintain strict alignment across regulatory filings and public disclosures. In public statements, the safer course is to use plain, nontechnical language that is less likely to be measured against a specific statutory or industry definition.

Practical Steps to Manage Dual Misrepresentation Risk

The state law boom means AI companies will be required to speak more often, more precisely and to more audiences about the same systems. To mitigate dual misrepresentation risk, companies should:

- Establish cross-functional coordination among AI compliance, securities disclosure and communications teams, routing all AI-related statements through a shared review process that keeps the legal mandates for AI law reporting, securities materiality and cybersecurity escalation distinct;

- Maintain a unified vocabulary map and incident taxonomy across all disclosure channels — public, investor-facing and regulator-facing — so that the same term carries the same meaning wherever it appears;

- Audit public statements for scope words that may convert aspirational messaging into testable representations; and

- Document the factual basis for each AI-related claim before it is made, so that consistency across channels can be demonstrated if challenged.

The objective is not to eliminate high-level messaging. Generalized statements about safety, compliance or responsible governance are often treated as puffery.

But every transparency report filed, every incident submission made and every training data disclosure posted adds a data point against which those broad statements can be tested.

The new AI-disclosure laws do not just regulate AI — they build a parallel evidentiary record that sits alongside everything a company says to investors, consumers and the public. Companies that recognize this dynamic early — and treat their AI-compliance records and

public narratives as a single, integrated body of statements — will be best positioned to navigate the dual misrepresentation risks ahead.

William Pao, Tijana Brien and Sean Quinn are partners at Cooley LLP.

Cooley associates Rebecca Kahn, Adam Silow and Julian Piroli contributed to this article.

[1] https://www.cornerstone.com/wp-content/uploads/2026/01/Securities-Class-Action-Filings-2025-Year-in-Review.pdf.