The Liability Map:
The Three Channels Through Which AI Creates Legal Exposure

Basil C. Puglisi, MPA
A Human-AI Collaboration

For the executives who deploy AI, the lawyers who defend them, and the insurers who price the risk.

The Map

Most organizations track AI risk by watching for new laws. That approach fails for a simple reason: the laws keep moving, and the exposure does not wait for them. Colorado offers the proof. The state enacted the first comprehensive AI statute in the United States in 2024, delayed it, and then repealed and replaced it on May 14, 2026, six weeks before it would have taken effect. An organization that spent two years building compliance around that statute's impact assessments now owns a program built for a law that never arrived. An organization that spent the same two years building governed, documented AI practice owns evidence that serves it in every jurisdiction, under every version of every statute.

That is the orientation. AI legal exposure does not travel through one channel that a single compliance project can close. It travels through three, and each channel asks a different question, demands a different artifact, and punishes a different failure.

The first channel is regulatory enforcement: governments fining, restricting, or barring the use of AI systems. The second is civil and product liability: courts assigning responsibility when AI output harms someone. The third is contract and insurance: the private ordering of AI risk through warranties, indemnities, and coverage, where exposure is reassigned quietly at every renewal and every signature. Mapping all three is the work of this piece, which walks the channels and names the one demand they share.

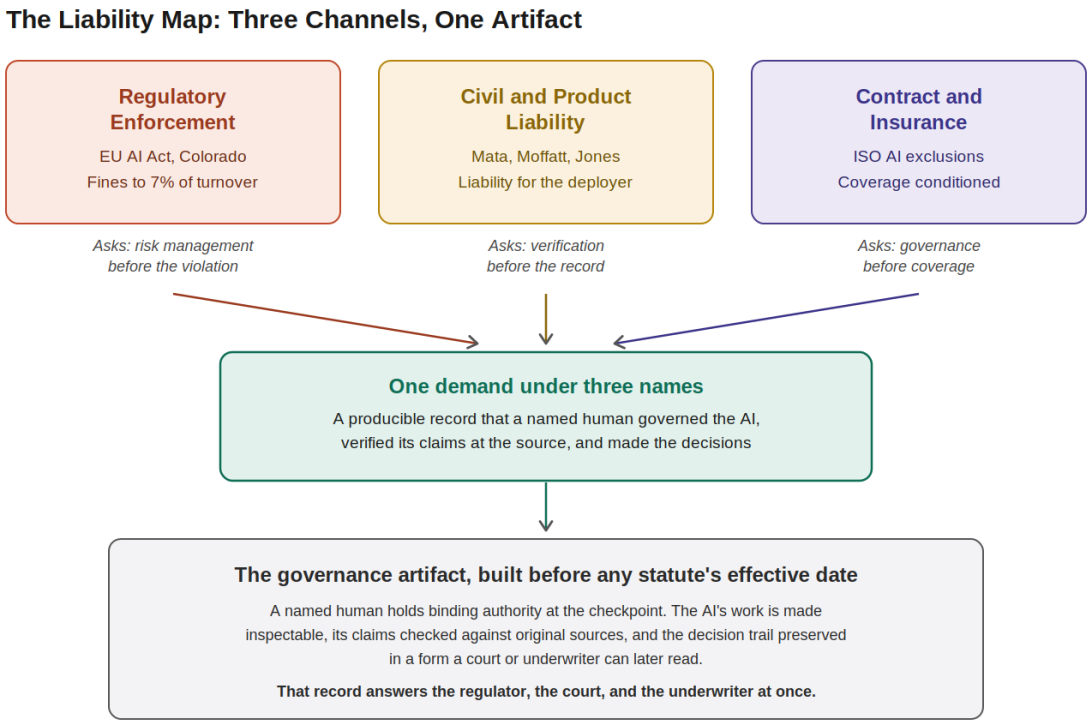

Figure 1. Three channels of AI legal exposure converge on a single demand: a producible record that a named human governed the AI.

Channel One: Regulatory Enforcement

The European Union operates the most developed enforcement regime. The EU AI Act, Regulation 2024/1689, classifies AI systems by risk and attaches a three-tier penalty structure under Article 99: fines up to 35 million euros or 7 percent of worldwide annual turnover for prohibited practices, up to 15 million euros or 3 percent for violations of high-risk system obligations, and up to 7.5 million euros or 1 percent for supplying misleading information to authorities. These ceilings exceed the GDPR's maximum of 4 percent. The timeline is staged and already partly live. Prohibited practices have been enforceable since February 2, 2025. Obligations for general-purpose AI models took effect on August 2, 2025. The high-risk timeline then moved. On May 7, 2026, the Council and the European Parliament reached a provisional agreement on the Digital Omnibus on AI, which defers obligations for stand-alone high-risk systems under Annex III to December 2, 2027, and for AI embedded in regulated products under Annex I to August 2, 2028. Those deferrals take legal effect only after formal adoption and publication in the Official Journal. Publication is expected before August 2, 2026, and that date remains the live one for transparency obligations under Article 50, with national market surveillance authorities empowered to enforce.

The United States offers no single statute and a far less stable map. Colorado's repeal and replacement is the defining example: the original Colorado AI Act's duty of care, risk management programs, impact assessments, and attorney general reporting were eliminated in favor of a narrower disclosure regime, the Colorado Automated Decision-Making Technology Act, effective January 1, 2027. Layered over the state activity is a December 2025 executive order seeking to limit state AI regulation, with preemption litigation expected to shape how much of the state layer survives.

The lesson of this channel is not any single statute. It is that regulatory exposure attaches to the use of AI in consequential decisions, whatever the statute of the season calls it, and that enforcement regimes reward one thing consistently: documented, demonstrable risk management that exists before the regulator asks.

Channel Two: Civil and Product Liability

Courts have not waited for AI statutes. They are assigning responsibility under doctrines that already exist, and three decisions mark the perimeter.

In *Mata v. Avianca*, the Southern District of New York sanctioned attorneys who filed AI-fabricated case citations. The court was explicit that there is nothing inherently improper about using a reliable AI tool for assistance, and that the duty the lawyers violated was the gatekeeping duty to verify what the tool produced before filing it. In *Moffatt v. Air Canada*, a British Columbia tribunal held the airline responsible for misinformation its customer-service chatbot provided, rejecting the argument that the chatbot was somehow a separate entity from the company that deployed it. The output crossed the point of no return, reaching a customer with no human checkpoint between the system and the harm, and the deployer owned the result. And in *Jones v Family Court at Whangārei*, which concerned a self-represented litigant who filed hallucinated citations, the Supreme Court of New Zealand cautioned in 2026 that reliance on the unverified outputs of AI applications in court filings may, in serious cases, amount to obstruction of justice or contempt. The caution was guidance from the bench rather than the decision's holding, and the phrase that matters carried across every channel: unverified outputs.

Three jurisdictions, three doctrines, one finding: liability attaches at the missing verification step, and it lands on whoever let unverified output enter the record. The product liability layer extends the same logic to AI embedded in products. The European Union's revised Product Liability Directive, Directive 2024/2853, brings software, including AI systems, within the strict liability regime, with transposition by Member States due by December 9, 2026.

Channel Three: Contract and Insurance

The quietest channel is repricing AI risk fastest. In January 2026, the Insurance Services Office introduced endorsements CG 40 47 and CG 40 48, giving commercial general liability carriers standard forms to exclude losses arising out of generative AI. Major carriers moved in parallel. Trade-press reporting indicates that Berkshire Hathaway, Chubb, and Travelers sought state regulatory approval to exclude AI-related damages from general liability policies, that a large majority of those requests were approved, and that exclusions began taking effect as early as January 2026. At the aggressive end, absolute AI exclusions for directors and officers, errors and omissions, and fiduciary policies eliminate coverage for claims arising out of AI use or development, with at least one carrier's endorsement reported to list inadequate AI governance and chatbot communications among the excluded grounds. The carrier-level figures here trace to industry reporting rather than primary regulatory filings, and they are presented as such.

The era insurers call silent AI coverage, where AI losses were paid by default because no policy language addressed them, is closing. What replaces it is instructive. The affirmative AI coverage now emerging, from Lloyd's syndicates and specialty entrants, underwrites hallucinations, model drift, and harmful outputs, and it conditions that coverage on detailed governance documentation. The same demand flows through the contract channel itself: vendor warranties, customer indemnities, and procurement terms increasingly turn on whether AI use was governed and whether the governance can be shown.

Read this channel plainly. The market that prices risk for a living has concluded that ungoverned AI use is becoming uninsurable, and that documented governance is the artifact that restores insurability. That conclusion was not reached by an ethics board. It was reached by underwriters.

The Incentive: One Artifact Serves All Three Channels

Each channel punishes a different failure, and all three reward the same preparation. The regulator asks for demonstrable risk management before the violation. The court asks for evidence that a human verified the output before it entered the record. The underwriter asks for governance documentation before binding coverage. Strip the vocabulary and the demand is identical: a producible record showing that a named human governed the AI's work, verified its load-bearing claims against original sources, and made the decisions that mattered.

That artifact does not have to wait for any statute's effective date, which is precisely its value. Colorado's repeal stranded compliance programs built to a statute's text. It stranded nothing for an organization whose AI use was already governed and documented, because the evidence such practice produces answers the regulator, the court, and the underwriter

regardless of which version of which law survives. The standard-of-care frameworks that anchor that practice, the NIST AI Risk Management Framework and ISO/IEC 42001 among them, are voluntary rather than binding, but legal analysts increasingly treat the NIST framework as a likely reference point for the reasonable-care standard in AI negligence, and insurers have begun to weigh governance maturity, human oversight, documentation, and clear accountability, as a factor that may shape future underwriting.

Major brokers and researchers have since converged on this exposure logic. The Willis Research Network's May 2026 *Risk and Resilience Review* frames AI as a risk amplifier across existing lines and concludes that the central question is no longer whether AI is used but where reliance occurs and whether governance, controls, and documentation are aligned to the failure mode. The argument set out here was published earlier in the author's working paper, *The AI Risk Economy: Why Insurance Cannot Price What Governance Cannot Prove* (Puglisi, May 24, 2026), which first mapped the five-tier governance-to-insurability relationship that this convergence now reflects. Both lines of work trace in part to the same cross-disciplinary convening, *The Intention Advantage*, a summit series organized by the Institute for Advertising Ethics with the IEEE Standards Association as knowledge partner, whose founding session in April 2026 surfaced the connection between human oversight and the insurance pricing of governance for a cross-disciplinary audience.

The Economic Case for Governing Now

This analysis closes against four questions, the same four that decide exposure in every channel above.

What duty or liability does this create, and who is exposed? All three channels converge on the deployer. The organization that lets AI output reach a customer, a court, or a counterparty without verification owns the result, whether the bill arrives as a regulatory fine, a judgment, or a denied claim.

What would an insurer need to see to underwrite the exposure? The emerging affirmative market answers directly: documented AI governance, named human oversight, and records of verification. Absent those, the standard market's answer is the exclusion endorsement.

What evidentiary standard would a lawyer or an examiner demand? Mata, Moffatt, and Jones supply it: proof that a human verified the output before it acted on the world. The defensible position is a contemporaneous record of who checked what, against which source, and who approved the result.

Where would an organization produce that evidence before a dispute begins? Through a method, not a tool: a named human holds binding authority at each consequential decision point, the AI's work is made inspectable rather than asserted, its load-bearing claims are

checked against original sources before they enter the record, and the decision trail of who approved what on which evidence is preserved in a form a court or underwriter can later read. That method is what governance means in practice once the principles and the engineering controls are stripped away. The frameworks that carry it, including Checkpoint-Based Governance, RECCLIN, CAIPR, CARCS, and SCOPE, are published as open-source proposals and named here only as one worked example of the method. The risk is real, the insurance industry is reacting faster than the regulators, and the methods to solve for this are already published publicly and provide the roadmaps to proactive economic rewards.

Sources

Brownstein Hyatt Farber Schreck. (2026, March 4). Colorado's landmark AI law coming online: What developers and deployers should know.

European Commission, AI Act Service Desk. Article 99: Penalties. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-99 (penalty tiers, including the 1 percent misleading-information tier).

Gibson Dunn. (2026, May 27). EU AI Act Omnibus agreement: Postponed high-risk deadlines and other key changes (Annex III deferral to December 2, 2027; Annex I to August 2, 2028; August 2, 2026 transparency obligations remain live pending Official Journal publication).

Goodwin. (2026, June). Colorado enacts law repealing and replacing landmark AI Act (first comprehensive state AI law; repeal and replace signed May 14, 2026; effective January 1, 2027).

Hogan Lovells. (2026, May 7). EU legislators agree to delay for high-risk AI rules (provisional agreement reached May 7, 2026).

Hunton Andrews Kurth. (2026, May). Colorado AI Act amended and effective date delayed.

Insurance Services Office endorsements CG 40 47 and CG 40 48 (effective January 1, 2026), as reported in Lathrop GPM, The AI coverage gap (2026, May 4), and Traverse Legal, AI insurance requirements (2026, April 24). Carrier-level approval figures and specific endorsement language reflect trade-press reporting, not primary regulatory filings.

Jones v Family Court at Whangārei [2026] NZSC 1 (self-represented litigant; AI-misuse caution from the bench).

Mata v. Avianca, Inc., No. 1:22-cv-01461 (S.D.N.Y. 2023).

Moffatt v. Air Canada, 2024 BCCRT 149.

Morrison Foerster. (2026, May 15). Colorado hits reset on AI regulation with a new AI Act.

PYMNTS. (2026, May 1). Big insurance backs away from AI risk and startups rush in (carrier exclusion reporting).

Puglisi, B. C. (2026, May 24). *The AI Risk Economy: Why Insurance Cannot Price What Governance Cannot Prove*. basilpuglisi.com. https://basilpuglisi.com/ai-risk-economy-insurance-governance/

Regulation (EU) 2024/1689 (EU Artificial Intelligence Act), Article 99 penalty tiers and staged application dates.

Revised Product Liability Directive, Directive (EU) 2024/2853 (software and AI within strict liability; transposition due December 9, 2026).

Willis Research Network. (2026, May). *AI in Action: The Road to Responsible Adoption. Risk and Resilience Review*. WTW.

NIST AI Risk Management Framework (AI RMF 1.0), voluntary framework, and ISO/IEC 42001:2023 AI management system standard, as discussed in relation to the tort duty-of-care standard in Catastrophic Liability: Managing Systemic Risks in Frontier AI Development (arXiv:2505.00616) and in industry analysis of AI governance and insurance underwriting (Johnson Lambert LLP, 2026; StackAware, 2025, citing The Geneva Association, 2024).

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

#AIassisted using the HAIA Ecosystem