

Selling Control Without a Backstop: AI Governance Frameworks, the Uninsurability of Artificial Intelligence, and the Vendor in the Civil Liability Chain

Travis Gilly*

Working draft, June 2026 (v2)

Abstract

A market has formed around the sale of artificial intelligence governance. Vendors offer frameworks, audits, maturity models, and certifications that promise an enterprise it can deploy AI responsibly and demonstrate that it did. This article argues that the governance market is selling a form of control over a risk that the insurance market has just declared it will not underwrite, and that the gap between those two facts has a legal shape. As of late 2025 and into 2026, major carriers filed to exclude artificial intelligence liabilities from standard corporate policies, and the standard form publisher for United States commercial general liability released endorsements that carve generative AI out of the baseline coverage that most businesses carry. Regulatory enforcement, where it reaches at all, terminates at the deployer. The result is a civil loss-allocation vacuum: when a deployed system causes a third party harm, the loss lands on the deployer's balance sheet with no insurer behind it and no statute reallocating it. Into that vacuum the governance vendor has inserted itself, not by accident but through its own marketing, holding out its framework as the thing that makes deployment safe and defensible. This article develops three claims. First, control over a contemporary AI system cannot honestly be sold, because the system is stochastic in output, open in action space, and indifferent as to substrate, so no framework can promise the determinacy the sale implies. Second, when a framework is marketed as the assurance on which a deployer reasonably relied, the vendor has volunteered for a place in the civil liability chain that regulation does not give it and that an indemnity clause cannot fully cure. Third, the governance genre asks four familiar questions and never asks the fifth, which is whether the buyer and the people the buyer's system will affect have been told, in plain terms, that the risk is uninsured. The article closes on that fifth question.

Keywords: artificial intelligence governance; insurability; civil liability; loss allocation; vendor liability; negligent misrepresentation; EU AI Act; ISO CG 40 47.

*Travis Gilly is Founder and Executive Director of the Real Safety AI Foundation. ORCID: 0009-0007-2954-6313. Correspondence: t.gilly@ai-literacy-labs.org. The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All legal analysis, case synthesis, theoretical framing, doctrinal arguments, and conclusions are solely the author's own.

Contents

I	Introduction	3
II	The Insurability Collapse	4
	A The Carrier Retreat	4
	B From Silent Coverage to Explicit Exclusion.	5
III	Two Chains, One Vacuum.	6
	A The Regulatory Chain Stops at the Deployer	6
	B The Civil Loss-Allocation Vacuum	7
IV	The Framework as Evidence.	8
V	The Impossibility of Sold Control.	8
	A Stochastic Output	9
	B Open Action Space	9
	C Unspecified Substrate	10
VI	The Vendor Enters the Civil Chain.	10
	A Marketing as the Last Solvent Link.	10
	B Indemnity Cannot Cure It.	11
VII	The Fifth Question	12
VIII	Conclusion	13

I. Introduction

There is a genre of professional product that sells an enterprise the feeling of control over its artificial intelligence. The product takes many forms. It is a governance framework, a maturity model, a set of seven fundamental questions, a trademark-bearing methodology, a readiness assessment, a third party audit that issues a certificate. The common promise is that an organization that adopts the product can deploy AI responsibly and, just as importantly, can show that it did. The market for that promise is large and growing, and the people selling it speak with confidence about questions to ask, controls to install, and boxes to check.

This article is about a fact that the genre rarely mentions to the buyer. At roughly the same moment the governance market matured, the insurance market reached the opposite conclusion about the same technology. In November 2025, the Financial Times reported that major carriers, including AIG, Great American, and W.R. Berkley, had filed with United States regulators to exclude artificial intelligence liabilities from standard corporate policies.¹ One proposed W.R. Berkley exclusion would bar claims involving “any actual or alleged use” of artificial intelligence, reaching products that merely incorporate the technology.² The carriers were not raising premiums. They were declining to write the risk at all, which is what insurers do when a loss is one they cannot model.³

The exclusion is not confined to a few carriers’ bespoke endorsements. The standard form publisher made it standard. Effective January 1, 2026, Verisk’s Insurance Services Office released endorsement CG 40 47, which excludes bodily injury, property damage, and personal and advertising injury arising out of generative artificial intelligence under both Coverage A and Coverage B of the commercial general liability form, along with CG 40 48 for Coverage

1. Lee Harris & Cristina Criddle, *Insurers Retreat from AI Cover as Risk of Multibillion-Dollar Claims Mounts*, Fin. Times (Nov. 23, 2025), <https://www.ft.com/content/abfe9741-f438-4ed6-a673-075ec177dc62>.

2. *Id.*; see also *Major Insurers Move to Avoid Liability for AI Lawsuits*, Tom’s Hardware (Nov. 25, 2025) (quoting the “any actual or alleged use” language and describing the systemic, correlated, aggregated risk that drove it).

3. An underwriter described generative AI output to the Financial Times as “too much of a black box.” See Harris & Criddle, *supra* note 1. Aon’s head of cyber explained that the industry could absorb a single large loss but not a “systemic, correlated, aggregated risk” in which one upstream model failure produces a thousand losses at once. See Tom’s Hardware, *supra* note 2.

B alone and CG 35 08 for products and completed operations.⁴ Because ISO forms underlie roughly four-fifths of United States property and casualty policies, the practical reach of that standardization is the commercial baseline itself.⁵ What had been “silent” coverage, present because nothing excluded it, became an explicit and standard exclusion.

So the buyer of a governance framework occupies a position the seller seldom names. The deployer is being sold control over a risk that its own insurer has just refused to stand behind. This article asks what follows, in law, from that arrangement. Part II documents the insurability collapse. Part III separates the two chains that allocate AI loss, the regulatory chain and the civil chain, and locates the vacuum between them. Part IV examines the move by which a governance framework becomes evidence. Part V argues that control over a contemporary AI system cannot honestly be sold. Part VI shows how the vendor, through its own marketing, enters the civil chain that regulation leaves open, and why an indemnity clause does not fully cure the exposure. Part VII states the fifth question the genre never asks.

II. The Insurability Collapse

A. The Carrier Retreat

Insurance functions by pooling risks that are independent enough to be modeled across many insureds.⁶ Generative AI breaks the premise. A single model, a single provider, a single update,

4. Ins. Servs. Office, *Exclusion — Generative Artificial Intelligence*, Form CG 40 47 01 26; *Exclusion — Generative Artificial Intelligence (Coverage B Only)*, Form CG 40 48 01 26 (eff. Jan. 1, 2026). The forms define “generative artificial intelligence” as a machine-based learning system or model trained on data with the ability to create content or responses, including text, images, audio, video, or code. The trigger phrase is “arising out of,” which under settled insurance law requires only a causal connection, not proximate cause.

5. *See Cyber Re/insurance Market Hits New High as AI Risks Reshape Coverage*, Reinsurance Bus. (May 4, 2026) (reporting that ISO forms stand behind roughly 82% of United States property and casualty policies and quoting Munich Re’s Stefan Golling). By April 2026, carriers including W.R. Berkley, Chubb, Travelers, Cincinnati Financial, and Berkshire Hathaway had filed to adopt the endorsements or proprietary equivalents, with state regulators approving more than four-fifths of submitted filings.

6. The condition of insurability is a pool of exposures that are individually uncertain but collectively predictable through the law of large numbers; correlation across insureds is the property that defeats that predictability. *See Harris & Criddle, supra* note 1 (quoting Aon’s Kevin Kalinich, who explained that the industry could absorb a single \$400 to \$500 million loss but not a “systemic, correlated, aggregated risk” in which one upstream model failure produces a thousand simultaneous losses).

or a single vulnerability can produce thousands of simultaneous losses, which makes the risk correlated rather than independent and therefore resistant to ordinary pricing.⁷ The losses that prompted the retreat were not hypothetical. A search engine’s AI summary feature falsely stated that a solar company faced suit by a state attorney general, drawing a defamation action seeking roughly \$110 million.⁸ An airline was ordered to honor a discount its customer service chatbot had invented.⁹ An engineering firm lost the equivalent of roughly \$25 million when employees were deceived by a deepfake video impersonation of a senior executive.¹⁰

By early 2026, the move had spread from the small group of named carriers to the broader market. Industry reporting described carriers across errors and omissions, directors and officers, cyber, and general liability lines declining to write coverage for claims arising from AI-generated outputs, and treating AI vendors more harshly still, frequently declining to cover them at all.¹¹

B. From Silent Coverage to Explicit Exclusion

The shift the standard forms accomplished is a shift from silence to text. Before the endorsements, a commercial general liability policy did not mention AI, so a policyholder could argue that an AI-related loss fell inside the grant of coverage because nothing carved it out. That “silent AI” posture mirrors the “silent cyber” problem the market spent years unwinding

7. See Geneva Ass’n, *Gen AI Risks for Businesses: Exploring the Role for Insurance* 1–3 (Oct. 2025) (surveying 600 corporate insurance decision-makers across China, France, Germany, Japan, the United Kingdom, and the United States, and assessing the insurability of emerging generative AI risks). The report found that roughly seven in ten surveyed businesses had implemented generative AI in at least one function and that demand for risk transfer was strongest among firms that had already experienced severe failures.

8. See Harris & Criddle, *supra* note 1 (collecting incidents). The defamation suit over the AI summary is among the incidents the carriers cited in support of their exclusions.

9. *Moffatt v. Air Canada*, 2024 BCCRT 149 (Can. B.C. Civ. Resol. Trib.) (holding the airline liable for its chatbot’s misstatement and rejecting the argument that the chatbot was a separate legal entity responsible for its own representations).

10. See Harris & Criddle, *supra* note 1. The deepfake impersonation losses illustrate the social-engineering vector that carriers cite when describing AI as an aggregating exposure rather than a discrete one.

11. See *Insurance Carriers Quietly Back Away from Covering AI Outputs*, CIO (Apr. 2026) (describing carriers carving out coverage across cybersecurity and errors and omissions lines and noting that a company identifying itself as an AI company “has a good chance” of receiving a declination).

after 2018.¹² The CG 40 47 line converts that silence into an explicit exclusion at renewal.

Because the change attaches at renewal, often without a separate negotiation, many insureds will not learn that their coverage changed until a claim is denied.¹³

The carriers' withdrawal has a mirror image: a thin and conditional affirmative market. Munich Re writes performance coverage for an insured's own AI under its aiSure line, translating risks such as discrimination into a measurable error rate against a fairness metric.¹⁴ The conditions on that affirmative market matter to the argument that follows. Carriers increasingly tie coverage, where they will write it at all, to the insured's governance posture, and at least one carrier has introduced an endorsement that references compliance with the EU AI Act as a coverage criterion.¹⁵ That linkage is the hinge of this article: it is precisely what gives a governance framework its claimed value, and precisely what pulls the framework's seller toward the loss when the framework fails.

III. Two Chains, One Vacuum

When a deployed AI system injures a third party, two distinct systems decide who answers for it. They are easy to conflate and important to separate.

A. The Regulatory Chain Stops at the Deployer

The first chain is regulatory enforcement. In the European framework, the deployer of a high-risk AI system carries defined obligations of human oversight, monitoring, and use in accordance

12. See *ISO CG 40 47 AI Exclusion Forces a GL Rate Adequacy Rethink*, actuary.info (May 22, 2026) (drawing the parallel to the silent cyber experience and describing the "arising out of" trigger as one of the broadest causal phrases available in policy drafting).

13. See *ISO Just Made AI Exclusions Standard*, Insurance Intel (May 18, 2026) (observing that most commercial insureds are unaware their policies changed at renewal and warning that an undisclosed CG 40 47 attachment can shift errors-and-omissions liability onto the broker who bound it without notice).

14. Munich Re, *Insuring Generative AI: Risks and Mitigation Strategies* (white paper) (describing the aiSure product line and Munich Re's practice, since 2018, of treating model underperformance as a quantifiable error rate). Standalone offerings from other entrants, including performance warranties from specialty insurers and managing general agents backed by the London market, remain limited in capacity and conditioned on the insured's controls.

15. See *The Insurance Industry Just Became AI's Most Powerful Regulator*, Modulos (Mar. 24, 2026) (describing QBE's endorsement referencing EU AI Act fines and predicting that governance frameworks will become prerequisites for coverage rather than optional additions).

with instructions.¹⁶ A regulator that finds a violation proceeds against the deployer, and in some cases the provider, under the statute's penalty structure. That chain is real, but it is enforcement of public obligations, not a mechanism that compensates the injured third party for a private loss. It also operates on a legislative clock that has slipped. The European Commission's Digital Omnibus on AI, proposed in November 2025, would defer the high-risk obligations from August 2, 2026 to December 2, 2027 for standalone systems, with later dates for AI embedded in regulated products.¹⁷ The point is not that enforcement will never arrive. The point is that even when it does, it terminates at the operator of the system and answers a public-law question, not the private-law question of who pays the person the system harmed.

B. The Civil Loss-Allocation Vacuum

The second chain is civil loss allocation: tort and contract law deciding who bears a private loss. Ordinarily a deployer transfers residual exposure to an insurer through errors and omissions, directors and officers, and cyber policies, and the insurer absorbs the loss. Part II showed that transfer closing. The provider, meanwhile, has disclaimed the loss by contract; the major model terms place the risk of the output on the user.¹⁸ The European Commission's proposed AI Liability Directive, which would have eased a claimant's path against actors along the value chain, was withdrawn, leaving the civil allocation question to existing national law.¹⁹

Stack those facts and a vacuum appears. The insurer has stepped out. The provider has contracted out. The regulator, when present, enforces public duties against the deployer rather

16. Regulation (EU) 2024/1689, art. 26, 2024 O.J. (L) (laying down obligations of deployers of high-risk AI systems, including the duty to take appropriate technical and organizational measures to ensure use in accordance with the instructions of use and to assign human oversight to competent natural persons).

17. Proposal for a Regulation Amending Regulations (EU) 2024/1689 and (EU) 2018/1139 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI), COM (2025) 836 final (Nov. 19, 2025). A provisional political agreement reached in May 2026 replaced the proposal's conditional trigger with fixed deferral dates of December 2, 2027 for standalone high-risk systems and August 2, 2028 for high-risk systems embedded in regulated products.

18. See Arthur Charpentier, *Insurers and AI, a Systemic Risk*, *Freakonometrics* (Nov. 25, 2025) (observing that AI providers' contracts include drastic limitations of liability, exclusions of performance guarantees, and clauses transferring almost all risk to the user, enforced through the near-inelastic demand for the major models).

19. See Inst. of Int'l & Eur. Affairs, *The Transition to a New Digital Policy Agenda* (Dec. 15, 2025) (noting the withdrawal of the proposed AI Liability Directive and the ongoing debate over whether it will be reintroduced).

than compensating the third party. The loss therefore comes to rest on the deployer, uninsured, with the deployer's executives and board exposed to the residue. That is the space the governance vendor's product is sold into, and it is the space the next two Parts examine.

IV. The Framework as Evidence

A governance framework is not only an operational tool. Once deployed, it becomes evidence, and the parties know it. Carriers ask to see an applicant's AI usage policy, tool inventory, and governance documentation, and reward demonstrated human-in-the-loop controls with better terms.²⁰ A deployer facing a claim will introduce its adherence to the framework as proof of due care. A regulator weighing a penalty will read the same documentation.²¹ The framework, in other words, is offered and consumed as an evidentiary artifact whose value is its capacity to show that someone was reasonable.

That evidentiary function is where the sale becomes legally interesting. A framework introduced to prove due care is a representation about the adequacy of the deployer's control. If the vendor authored and certified the representation, and held it out as sufficient, the vendor has made a statement about a matter within its professed expertise, on which the buyer was meant to rely. The remaining Parts ask whether that statement can be true, and what attaches to the vendor when it is not.

V. The Impossibility of Sold Control

The governance sale implies a determinacy the underlying system does not have. Three properties of a contemporary AI system defeat the promise of sold control, and they compound.

20. *See Major Insurers Are Pulling Back from AI Liability*, Metropolitan Risk Advisory (Nov. 24, 2025) (advising insureds to submit AI usage policies, tool inventories, and governance frameworks, and to demonstrate human oversight, in order to secure better terms in a hardening market).

21. The evidentiary use runs in both directions: the same documentation that supports a due-care defense also supplies the regulator and the insurer with a record against which to measure the deployer's conduct. *See* Modulos, *supra* note 15 (describing governance frameworks moving from optional additions toward prerequisites for coverage, so that the framework becomes the criterion on which both insurability and a liability defense turn).

A. Stochastic Output

A generative model is a probabilistic system. The same input does not yield the same output, and the distribution has tails that include the harmful response. A reinsurer that writes AI coverage describes its own task as quantifying the probability and severity of model underperformance, not eliminating it.²²²³ A framework can lower the rate of bad outputs and document the effort. It cannot drive the rate to zero, because the system that produces the outputs is stochastic by construction. A promise of control that means a promise of determinate output is a promise about a different kind of machine.

B. Open Action Space

The second property is the open action space of agentic deployment. A system that only returns text occupies a bounded space of harms. A system empowered to send messages, move funds, file documents, or take steps in the world occupies an action space the designer did not fully enumerate, because enumerating it is the unsolved part. Governance that lists prohibited actions cannot foreclose the action nobody listed. The relevant loss is frequently the unenumerated one, which is why a framework keyed to a checklist of known failure modes does not bound the exposure it claims to bound.²⁴

22. Munich Re, *supra* note 14 (describing the pricing of AI performance coverage in terms of the predictive robustness of the model and the fluctuation of its loss ratio, and treating hallucination as a known error type rather than an eliminable defect).

23. This is why the carriers framed the exposure as one they could not model rather than one they could price. *See* Harris & Criddle, *supra* note 1 (underwriter describing model output as “too much of a black box”); *see also* Geneva Ass’n, *supra* note 7, at 2 (identifying opacity and unpredictability as the features that move generative AI outside ordinary insurance boundaries).

24. Carriers have signaled the same intuition by treating agentic deployments and AI vendors more severely than text-only use, in many cases declining vendor coverage outright. *See* CIO, *supra* note 11 (reporting that a company identifying itself as an AI company “has a good chance” of receiving a declination, and that carriers treat AI vendors differently from AI users). The reason tracks the doctrine: an open action space is an open class of foreseeable but unspecified harms, and underwriters price what they cannot enumerate as systemic rather than discrete. *See* Harris & Criddle, *supra* note 1.

C. Unspecified Substrate

The third property is substrate indifference. The deployer typically does not control, and often cannot inspect, the model on which its application runs. The model is external, opaque, and subject to update by the provider on the provider's schedule. A framework that certifies a deployment certifies a configuration that the provider can change underneath it without notice, and the provider has disclaimed responsibility for the consequences by contract.²⁵ Control sold over a substrate the seller does not govern is control in name.

Taken together, stochastic output, open action space, and unspecified substrate mean that no framework can deliver the determinate control the sale implies. A vendor may honestly sell documentation, process discipline, and a lower error rate. A vendor that sells the assurance that deployment is safe and defensible is selling something the technology cannot supply, and is doing so into the uninsured vacuum of Part III.

VI. The Vendor Enters the Civil Chain

A. Marketing as the Last Solvent Link

Regulation does not place the governance vendor in the civil liability chain. The vendor places itself there, through the content of its own marketing. When a vendor holds out its framework as the assurance that makes deployment responsible and defensible, and a buyer adopts the framework in reliance on that representation, the vendor has supplied information for the guidance of another in the other's business dealings. That is the predicate for liability for negligent misrepresentation, which reaches one who, in the course of a business in which it has a pecuniary interest, supplies false information for the guidance of others who justifiably rely on it, where the supplier failed to exercise reasonable care in obtaining or communicating

25. See Charpentier, *supra* note 18 (describing the contractual asymmetry by which providers impose limitation-of-liability and risk-transfer terms on deployers who lack viable alternatives, even in regulated sectors that require the deployer to control its models).

the information.²⁶ A vendor that markets certainty about a system that cannot be made certain, and whose representation is then introduced as the deployer's proof of care, has volunteered for the role of the last solvent link in a chain whose other links have stepped out: the insurer by exclusion, the provider by disclaimer.

The vendor's exposure is not limited to the buyer. Where the framework's failure contributes to a third party's injury, and the vendor held itself out as the professional assurance that the deployment was safe, the structure resembles the negligent-undertaking principle under which one who undertakes to render services it should recognize as necessary for the protection of a third person may be liable for a failure to exercise reasonable care.²⁷ The doctrinal labels will vary by jurisdiction and by the precise representations made. The structural point holds across them: the vendor that markets itself as the backstop has nominated itself for the position the backstop used to occupy.

B. Indemnity Cannot Cure It

The reflexive answer is the indemnity clause. The vendor will say its contract shifts any liability back to the buyer, so the marketing exposure is closed. The answer is incomplete for two reasons. First, an indemnity running from the buyer to the vendor does nothing about the vendor's exposure to an injured third party, who is not a party to the contract and whose claim sounds in tort. Second, a contractual term that purports to exempt a party from liability for harm caused by its own reckless or intentional conduct is unenforceable as against public policy, and a

26. Restatement (Second) of Torts §552 (Am. L. Inst. 1977) (information negligently supplied for the guidance of others). The elements that matter here are the supplier's pecuniary interest in the transaction, the supply of information for the guidance of the recipient in its business, and the recipient's justifiable reliance. The section has been widely adopted to govern the liability of professionals who furnish information for the guidance of others in their business affairs. *See, e.g., Nycal Corp. v. KPMG Peat Marwick LLP*, 688 N.E.2d 1368 (Mass. 1998) (adopting §552 and confining a professional's negligent-misrepresentation exposure to those persons and transactions the supplier knows its work product is intended to influence).

27. Restatement (Second) of Torts §324A (Am. L. Inst. 1977) (liability to third persons for negligent performance of an undertaking). The principle is long settled: having undertaken to act, a party must use due care, and is answerable where another reasonably relies on the undertaking. *See Indian Towing Co. v. United States*, 350 U.S. 61 (1955) (Coast Guard that undertook to operate a lighthouse and let the light fail, inducing reliance, could be held to the good-Samaritan duty to use due care in the undertaking). Whether the doctrine reaches a particular vendor depends on the specificity of the undertaking and the foreseeability of reliance, but the marketing of a framework as the protection against AI harm is the kind of undertaking the section contemplates.

term exempting a party from liability for its own negligence is subject to close scrutiny and to invalidation where it offends public policy.²⁸ A vendor that sold certainty it could not deliver, into a market it knew the insurers had exited, may find that the very conduct it most needs to indemnify away is the conduct the indemnity cannot reach.

VII. The Fifth Question

The governance genre is built on questions. The frameworks ask what the system does, who is accountable, what controls are in place, and how incidents are handled. Those are good questions. They are also the four questions a vendor can ask while never disturbing the sale, because each can be answered with a document the vendor provides.

There is a fifth question the genre does not ask, and its absence is not random. The fifth question is whether the buyer, and the people the buyer's system will affect, have been told in plain terms that the risk is uninsured. It is the question that, asked honestly, makes the product harder to sell, because the honest answer changes what the buyer is buying.²⁹ The buyer is not buying control over a model that produces correlated, unmodelable, now-excluded losses. The buyer is buying documentation of effort, which has value, and a representation of safety, which does not have the value implied, deployed into a space where the loss has nowhere to go but the buyer's own balance sheet.

28. Restatement (Second) of Contracts §195 (Am. L. Inst. 1981) (a term exempting a party from liability for harm caused intentionally or recklessly is unenforceable on grounds of public policy, and a term exempting a party from liability for negligence is unenforceable in defined circumstances, including where the party charged is a professional whose service is of importance to the public). The leading articulation of the public-interest limit is *Tunkl v. Regents of University of California*, 383 P.2d 441 (Cal. 1963), which sets out the factors that render an exculpatory clause unenforceable, among them that the party seeking exculpation performs a service of practical necessity, holds itself out as suitable for the task, and exercises superior bargaining strength over a party who must accept the term to obtain the service. A governance vendor that markets professional assurance to deployers who cannot otherwise transfer the risk fits each factor uncomfortably well.

29. The same non-disclosure dynamic has already produced liability one layer down, at the broker who binds an AI exclusion without telling the insured. See Insurance Intel, *supra* note 13 (warning that where a CG 40 47 endorsement is attached without policyholder communication and a claim is later denied, the errors-and-omissions liability flows to the broker). If non-disclosure of the coverage change exposes the broker, non-disclosure of the underlying uninsurability is not a smaller problem for the vendor selling assurance into the gap.

To ask the fifth question is to convert the governance conversation from a compliance exercise into an honest allocation of an uninsured risk. The deployer learns that the framework is a record of diligence, not a transfer of loss. The third party who will be affected gains nothing from the framework at all, because a certificate does not compensate an injury. And the vendor learns that the marketing which closes the sale is the same marketing that may place it in the chain when the loss arrives. The genre's silence on the fifth question is therefore not an oversight. It is the condition on which the sale depends.

VIII. Conclusion

Two markets reached opposite conclusions about the same technology at the same time. The insurance market decided that artificial intelligence liability is, for now, a risk it will not stand behind, and wrote that decision into the standard forms that define the commercial baseline. The governance market decided that the same technology can be controlled, certified, and made defensible, and sold that conclusion to the firms the insurers had just declined. Between them sits a deployer holding an uninsured loss, and behind the deployer sits a vendor whose marketing has nominated it as the backstop the insurers stopped being.

Control over a contemporary AI system cannot honestly be sold, because the system is stochastic in output, open in action space, and indifferent as to substrate. A framework can document care and lower a rate. It cannot deliver the determinacy the sale implies, and a vendor that promises that determinacy, into a market it knows the carriers have exited, has volunteered for a place in the civil chain that regulation does not assign it and that an indemnity clause cannot fully reclaim. The genre asks four questions and sells a feeling. The fifth question, the one about who actually pays when the model does what stochastic systems do, is the one the buyer and the affected public are owed. This article's modest proposal is that it be asked out loud.