PERSPECTIVES

THE EMERGING ROLE OF EXPERTS IN ARTIFICIAL INTELLIGENCE DISPUTES:
TECHNICAL CAUSATION, ECONOMIC HARM, AND GOVERNANCE FAILURE

JS|HELD

Our perspectives feature the viewpoints of our subject matter experts on current topics and emerging trends.

thebrief™ | JS|HELD

Intellectual property litigators, artificial Intelligence (AI) and technology litigation attorneys, chief intellectual property officers, and others should read this article to learn more about:

- Why AI disputes now extend beyond copyright into privacy, employment, securities, safety, antitrust, and governance claims.
- Why AI cases often require technical and damages experts to work from the same factual record.
- Why the toughest issues frequently arise in non-IP cases involving bias, privacy, safety, fraud, and pricing.
- How AI's opacity and non-deterministic behavior complicate causation, damages, and admissibility.
- How integrated expert analysis can reduce inconsistencies and strengthen defensibility.
- Why future disputes may focus increasingly on AI governance, oversight, controls, and risk management.

EXPERT VOICES

James E. Malackowski

James draws on decades of experience in intellectual property valuation and expert testimony to explain why AI disputes increasingly require a unified analysis of technical causation, legal liability, and economic harm, particularly where the same evidence addresses all three issues. He frames the paper's central insight that traditional sequential expert models are being reshaped by the unique evidentiary and valuation challenges of AI litigation.

J. Scott Womack

Scott brings a governance-focused perspective, examining how AI disputes are evolving beyond model performance to questions of board oversight, risk management, and organizational accountability. He highlights the emerging role of governance maturity and control frameworks in assessing liability, underwriting risk, and establishing standards of care in the next generation of AI litigation.

Executive Summary

As AI systems become more deeply embedded in business operations and consumer products, the disputes surrounding them have expanded well beyond the intellectual-property and copyright claims that first defined the field. The docket now reaches privacy, employment, consumer protection, securities, product safety, antitrust, and questions of corporate governance. Drawing on a study of more than 425 publicly filed AI-related disputes, this article examines how the opacity of these systems, their non-deterministic behavior, and the practical difficulty of obtaining model data complicate the assessment of causation, liability, and economic harm. It identifies a structural feature of AI litigation that the existing literature has largely left unnamed: in a substantial share of matters, the question of what an AI system did and the question of what that conduct is worth are not readily separable, because both turn on the same underlying technical facts. That interdependence blurs the conventional line between technical and damages issues, and in doing so, it strains the sequential division of expert labor on which litigation ordinarily relies. Looking ahead, the article argues that the most consequential AI disputes of the coming decade may arise less from what these systems create than from how the organizations deploying them choose to govern their use.

Introduction: A New Litigation Category

AI disputes are no longer confined to the copyright claims that first drew public attention. Across the matters examined in this study, the rate of new filings climbed from a low base before 2019 to 34 in 2022, 66 in 2023, 62 in 2024, and roughly 85 in 2025—a single-year increase of roughly 35%.[1] The disputes have diversified across legal theory, technology, and forums. Generative AI accounts for the largest share of the studied matters, at roughly 59%, but it now sits alongside AI-enabled decision systems (about 15%), computer-vision applications (about 11%), predictive and machine-learning models (about 9%), and autonomous systems and robotics (about 5%), each generating its own distinctive litigation. Increasingly, the disputes concern operational deployment rather than model development.

AI systems can occupy three roles at once: a technical artifact whose behavior must be explained, an economic asset whose value must be quantified, and an evidentiary instrument whose reliability must be defended in court.[2] The proposition this article tests against the data is that those roles are often not separable for purposes of expert analysis. In a measurable share of the studied matters, what data trained a model, whether a safety guardrail existed, and how a pricing algorithm behaved are, at the same time, the facts that establish liability and the facts that fix the measure of damages. Where traditional litigation retains a technical expert to establish what happened and then an economist to price it, many AI disputes reject that handoff.

This study analyzes more than 425 AI-related disputes to identify a structural feature of AI litigation that existing trend analyses largely overlook, and from which we draw three things: an account of what the docket actually looks like; a structural argument, illustrated by a worked example, about why expert work in these matters is frequently integrated rather than sequenced; and a candid statement of the limits of that argument. This study establishes the empirical landscape and the structural claim and then develops a taxonomy of paired expert functions, the admissibility problem, the emerging frontier of governance-failure litigation, and practical guidance for counsel.

The Empirical Landscape

The observations in this article rest on a study of more than 425 publicly filed AI-related disputes, current through the second quarter of 2026. Each matter was coded along five dimensions: the underlying AI technology, the economic-harm theory pleaded, the technical-failure mode alleged, the jurisdiction, and the affected market segment. The sources are independent and publicly accessible filings, enriched by expert review. Two features of the data deserve emphasis before any inference is drawn from it.

First, the figures describe filings, not adjudications. A complaint reflects an allegation and a litigation theory, not a finding of liability; the distribution of claims therefore measures what plaintiffs and their counsel are willing to assert, which is informative but distinct from what courts will ultimately accept. Second, the study captures public filings and is necessarily a snapshot of a rapidly expanding series, so recent high-salience matters are over-weighted relative to the steady-state docket. The percentages that follow should be read as the shape of a moving

[1]This article draws on a J.S. Held study of more than 425 publicly filed AI-related disputes, coded by underlying technology, economic-harm theory, technical-failure mode, jurisdiction, and market segment, current through June 30, 2026. Filing-year counts and category and segment distributions cited herein derive from that study and the accompanying J.S. Held AI Disputes Monitor unless otherwise noted. Because a single matter frequently pleads more than one theory, the issue-level shares reported below overlap and do not sum to one hundred percent.

[2]On the analytical distinction among AI systems as technical artifacts, economic assets, and evidentiary tools, see J.S. Held, *Expertise for AI Litigation Risk* (Apr. 2026).

target, not as fixed proportions. Because many complaints plead several theories at once, the issue-level shares overlap and are reported as approximate.

Review of the Docket

The disputes cluster on a small set of model developers. A handful of firms representing the developers of the most widely deployed large language and generative models appear as defendants far more often than any others, with the largest single concentration among general-purpose model providers, and secondary clusters around computer vision, biometric systems, and autonomous vehicle technology. This concentration is not incidental to the expert-services question. Because a comparatively small number of system architectures generate a disproportionate share of disputes, forensic methodology developed against those architectures (e.g., log reconstruction, prompt replay, version comparison) can be reused across matters rather than rebuilt case by case.

By market segment, legal and intellectual property matters predominate at roughly 47% of the studied disputes, followed by civil rights and privacy (about 15%), financial services (about 9%), criminal justice and law enforcement (about 8%), and employment (about 7%). The jurisdictional distribution follows a more conventional pattern, as one might expect. Federal district courts exercising original jurisdiction handle the clear majority of AI-related matters, well outpacing the combined volume of state trial courts and foreign courts; the federal forum remains the center of gravity for this litigation. Even so, the meaningful minority of matters proceeding in state and foreign venues means experts cannot assume a single governing standard for admissibility and method will apply across every case.

Intellectual Property in the Foreground

Intellectual property claims are the largest issue cluster in the data: on an approximate coding, close to half of the studied matters implicate the unauthorized use of copyrighted content—text, images, audio, video, code, and proprietary databases—to train AI models, and a substantial share of those plead lost licensing revenue, market displacement, or unjust enrichment as the theory of economic harm.

One pattern, present in roughly 1 in 6 of the studied matters, is large-scale acquisition through scraping and the circumvention of platform protections, including allegations that developers bypassed technological protection measures in violation of the anti-circumvention provisions of the Digital Millennium Copyright Act. A second, narrower pattern alleged in a smaller subset is the use of known pirated corpora such as shadow libraries and torrent-sourced datasets as primary training material, an allegation that converts ordinary infringement into willful infringement and strains a fair-use defense. The complaint by Encyclopedia Britannica and Merriam-Webster against OpenAI, alleging the unauthorized copying of roughly 100,000 reference articles together with trademark harms from misattributed outputs, exemplifies the publisher-plaintiff variant.[3] The suit by BMG Rights Management against Anthropic, alleging infringement of 493 musical compositions and acquisition of training data through pirate libraries, exemplifies the rights-holder variant and squarely presents the question of model memorization (i.e., whether a model reproduces protected content at output).[4]

[3] *Encyclopedia Britannica, Inc. v. OpenAI, Inc.*, No. 1:26-cv-02097 (S.D.N.Y. filed Mar. 13, 2026) (alleging unauthorized copying of approximately 100,000 articles and dictionary entries, together with Lanham Act trademark claims premised on misattributed AI outputs).
[4] *BMG Rights Mgmt. (US) LLC v. Anthropic PBC* (N.D. Cal. filed Mar. 17, 2026) (alleging infringement of 493 musical compositions and acquisition of training data through pirate libraries and torrent sites).

The Non-IP Tail, Where the Novel Questions Cluster

The analytically demanding questions concentrate not in the IP foreground but in the smaller, more heterogeneous set of non-IP matters. Algorithmic-discrimination claims arising from automated decision systems in hiring, lending, housing, and education, and from the institutional misuse of unreliable AI-detection tools appear in roughly 1 in 7 of the studied disputes. The suit against Yale arising from a disputed AI-detection score, which alleged both unreliability and disparate impact on non-native English speakers, illustrates how the technical question of a classifier's false-positive rate can drive both the liability theory and the measure of harm.[5]

Biometric and privacy matters form a comparably sized cluster, exemplified by litigation over AI-integrated consumer hardware. The class action against Meta and Luxottica over smart glasses marketed as privacy-protective but alleged to have transmitted captured video to overseas human reviewers signals an emerging concern with hardware-level transparency and genuine consent for always-on sensors.[6] Product-safety and wrongful-death matters, fewer in number at roughly 1 in 10, raise the most consequential standard-of-care questions. The class action against xAI alleging that its generative-image model produced child sexual abuse material depicting identifiable minors frames inadequate safety architecture as a product defect, with corresponding exposure to punitive damages.[7]

Fraud and misrepresentation matters round out the tail. They include securities claims premised on the overstatement of AI capabilities, as in the class action against monday.com over the durability of its AI-driven revenue projections[8] and claims that AI systems were deployed to systematically underpay obligations, as in the suit against Liberty Mutual and its AI valuation vendors over a tornado damage claim.[9] A smaller set of antitrust matters alleges algorithmic price coordination through shared pricing software.

Across the corpus, filing frequency and analytical novelty appear inversely correlated. Intellectual property claims are the most numerous, yet they rest on the most mature damages frameworks (i.e., reasonable royalty, lost profits, and unjust enrichment) and on causation theories that, though contested in application, are familiar in form. The non-IP tail is thinner in terms of count but is where causation and damages remain genuinely unsettled: where the standard of care is still being defined, where the measure of privacy or discrimination harm is contested, and where the technical and economic questions are most tightly bound together. It is therefore in the tail, more than in the foreground, that integrated expert work tends to matter most.

Why AI Disputes Strain the Sequential Expert Model

In conventional commercial and IP litigation, expert work proceeds in sequence. A technical expert establishes what happened: how a product failed, what a system did, whether code was copied. An economic expert then takes that finding as a premise and quantifies the resulting loss. The division works because

[5] *Rignol v. Yale Univ.* (D. Conn. filed Feb. 2025) (alleging that disciplinary action based on a GPTZero AI-detection score was unreliable and discriminatory against non-native English speakers; preliminary injunction denied May 5, 2025).
[6] *Bartone v. Meta Platforms, Inc.*, No. 3:26-cv-01897 (N.D. Cal. filed Mar. 4, 2026) (alleging that AI-integrated smart glasses marketed as privacy-protective covertly transmitted captured video to overseas human reviewers).
[7] *Doe 1 v. X.AI Corp.*, No. 5:26-cv-02246 (N.D. Cal. filed Mar. 16, 2026) (product-liability and negligence class action alleging defective design of a generative-image model that produced child sexual abuse material depicting identifiable minors).
[8] *Potter v. monday.com Ltd.*, No. 1:26-cv-01956 (S.D.N.Y. filed 2026) (securities-fraud class action premised on alleged overstatement of the durability of AI-driven revenue growth).
[9] *Ryan v. Liberty Mut. Ins. Co.* (N.D. W. Va. filed Jan. 7, 2026) (alleging a coordinated claims-underpayment scheme executed through third-party AI valuation vendors, including CCC Intelligent Solutions, Snapsheet, and Claim Assist).

the two questions are usually separable. The cause can be established without reference to the price, and the price can be computed once the cause is fixed.

A skeptical reader will object that this is hardly unique to AI. Patent litigation has always involved facts such as the scope of the claims and the structure of the accused product that bear on both infringement and damages. Competent counsel have long coordinated their technical and economic experts accordingly. The objection is fair, and the difference this article presses is one of degree that, in the harder matters, becomes a difference in kind. In the ordinary patent case, the overlapping facts are stable and, through discovery, knowable; the accused device does not behave differently each time it is examined, and its design can be inspected. AI evidence has three features that strain that model past its usual tolerances: the behavior under examination is non-deterministic; the decisive facts are frequently unrecoverable without the developer's cooperation; and a single fact often serves three functions at once—liability, the damages base, and the input to a further technical analysis. The five features below develop the point.

Stochastic, Non-Deterministic Causation

AI models produce probabilistic outputs. The same input can yield different results across runs, so the traditional "but-for" inquiry must be re-engineered for a system whose behavior is described by a distribution rather than a fixed response. The consequence for expert work is that the probability structure of the harm tends to be the damages model. One cannot quantify expected loss without first characterizing the variance of the system's outputs, which is a technical exercise, and one cannot frame that technical exercise without knowing which outputs matter economically. In these matters, causation and quantification are computed together.

A Worked Example: Provenance, Liability, and the Royalty Base

Consider the rights-holder pattern exemplified by the BMG complaint against Anthropic.[10] The threshold technical exercise is a training-data provenance audit: determining which protected works were ingested into the model and through what channels. That audit is not a discrete preliminary to the legal and economic analysis; it produces, in one step, the facts on which all three depend.

- **The technical finding.** Forensic provenance analysis seeks to establish whether the 493 musical compositions at issue were present in the training corpus, and whether they were acquired through licensed channels or through pirate libraries and torrent sources. A related technical inquiry of memorization testing asks whether the model reproduces those works, in whole or in part, at output.
- **The legal implication.** Each technical answer maps directly onto a liability question. Ingestion of the works supports the claim of unauthorized copying; acquisition through pirate sources converts the allegation into one of willful infringement and undercuts a fair-use defense; and reproduction at output supports a distinct theory of output-level infringement separate from the act of ingestion.
- **The economic implication.** The same findings set the measure of damages. The set of 493 works is not merely the liability predicate; it is the royalty base on which any reasonable-royalty or unjust-enrichment computation must be built. The willfulness finding bears on the availability and magnitude of enhanced damages. And the memorization result fixes the damages denominator: if the model reproduces protected works, the scope of infringing output bounds the measure of harm, whereas if it does not, the theory of harm shifts entirely to ingestion and its associated licensing value.

[10] *BMG Rights Management (US) LLC v. Anthropic PBC*, No. 5:26-cv-02334, Compl. (N.D. Cal. Mar. 17, 2026), ECF No. 1.

Three professional questions are answered by a single body of evidence: what the forensic evidence shows, what it establishes in law, and what it is worth. The technical expert who scopes the provenance audit and the economic expert who builds the royalty model are not working in sequence; they are interpreting the same facts. This is the sense in which the sequential model is strained: not because coordination is merely convenient, as in the ordinary patent case, but because the dividing line between the technical and economic record has largely dissolved.

Multi-Layer Intellectual-Property Entanglement

A single AI system can simultaneously implicate patents covering its architecture, copyrights in its training data and code, trade secrets in its model weights, and contractual rights under the licenses governing its programming interfaces. Each layer carries a different valuation logic, and the layers interact. A valuation that treats one layer in isolation will tend to misstate the value of the asset as a whole, because the economic significance of, say, the trade-secret layer depends on facts established in the patent and contract layers.

Opacity, Explainability, and Provability

AI systems are frequently opaque, and their developers often treat models and training data as trade secrets, resisting disclosure even where the systems affect significant public interests. This opacity is not only a discovery problem; it is an economic variable. What can be proven about a system bounds what can be valued, because harm that cannot be demonstrated cannot be quantified. An expert who understands which technical facts are recoverable, such as through log reconstruction, version control, or telemetry, and which are not, is therefore also the best expert to advise which damages theories are viable. The discovery posture and the damages posture are aspects of one strategic problem.

Where the Model Still Holds

The interdependence described above is a tendency, not a universal law, and the argument is stronger for acknowledging where it does not apply. A meaningful number of AI-related disputes remain conventional in structure and are well served by the traditional sequence. A copyright matter with clean licensing facts and an admitted dataset may present an ordinary infringement-and-damages question in which the technical record can be settled before the economic analysis begins. A single-issue privacy claim under a statute with fixed statutory damages may require little economic modeling at all. A straightforward contract dispute over an AI procurement may turn on the agreement rather than the model's behavior. Not that every AI dispute demands an integrated expert engagement; it is that a distinctive and growing subset does, and that counsel are well served by recognizing which matters fall on which side of that line.

A Taxonomy of Paired Expert Roles

For the subset of matters in which the technical and economic questions are bound together, the practical task is to specify, for each recurring kind of dispute, the technical function and the economic function that must be performed in tandem. The following table maps the recurring dispute clusters observed in the study to their paired functions. It should be read not as a menu of separable specialties but as two vantage points on a single body of facts.

DISPUTE CLUSTER	TECHNICAL EXPERT FUNCTION	ECONOMIC / FINANCIAL EXPERT FUNCTION
Intellectual Property	Training-data provenance audit; scraping and circumvention	Lost licensing revenue; unjust enrichment; reasonable royalty; market displacement
Algorithmic discrimination (hiring, lending, housing, education)	Bias and disparate-impact testing; validation of detection tools, including false-positive-rate analysis	Lost wages; denied-opportunity and wealth-building loss; class-wide damages modeling
Biometric and data	Hardware and system data-flow forensics; review of consent and disclosure mechanisms	Data-as-property valuation; class-wide damages exposure
Safety, product liability, and wrongful death	Adequacy of guardrails, red-teaming, and crisis-intervention protocols against the standard of care	Personal-harm quantification; punitive-exposure analysis
Fraud misrepresentation (investor, insurance capability claims)	Verification of represented versus actual capability; telemetry of system performance	Investor-loss measurement; claims-underpayment analysis; restitution
Algorithmic collusion and pricing	Analysis of pricing-algorithm behavior and cross-firm synchronization	Antitrust overcharge and but-for price modeling

In the disputes where the two columns describe the same facts, the value of analyzing them together lies in coherence: a technical finding and an economic conclusion that rest on a shared record are harder to impeach than two reports assembled separately and reconciled after the fact.

Admissibility in a Stochastic World

Whatever the type of dispute, an expert opinion in an AI matter must withstand the gatekeeping standards governing the admissibility of expert testimony. The methodology described throughout this article is designed for that purpose: causation opinions are derived through reproducible forensic techniques (e.g., system-log reconstruction, prompt-replay analysis, model-version comparison, provenance auditing, and telemetry review) and cross-validated against known failure modes such as hallucination, drift, integration error, and bias.[11] Applied with discipline, these methods speak to the conventional questions a court asks: whether the technique can be and has been tested, whether it carries a known error rate, and whether it rests on accepted principles.[12]

The harder problem, and the one that sets AI matters apart from ordinary technical litigation, is that the system under examination itself is non-deterministic. A reliability standard built on the assumption that a properly conducted test reproduces a fixed result sits uneasily with a model that returns different outputs to the same prompt. Rather than reproduce a single output, the expert characterizes the distribution of outputs, reports the rate at which harmful behavior occurs, and pins the analysis to

[11]All technical-causation opinions described here are derived using reproducible forensic methodology—including system-log reconstruction, prompt-replay analysis, model-version comparison, training-data provenance auditing, and integration-telemetry review—aligned with the OWASP Top 10 for Large Language Model Applications and designed to satisfy the governing standards for the admissibility of expert testimony. *OWASP Foundation, OWASP Top 10 for LLM Applications 2025* (2024); Fed. R. Evid. 702; Daubert v. Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579, 589–95 (1993); Kumho Tire Co. v. Carmichael, 526 U.S. 137, 141–42, 147–49 (1999).
[12]*Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 509 U.S. 579 (1993); Fed. R. Evid. 702.

a specific model version under controlled conditions, so that the experiment is repeatable even though any single generation is not. Error rate, in this approach, is a property of the distribution rather than of a point estimate, and the stochastic character of the system becomes what the methodology measures rather than an obstacle to measuring it.

A second consideration is strategic. Separately retained experts who reason from different premises tend to produce reports that diverge at the edges, and those divergences are a natural target for a motion to exclude. An engagement that develops technical and economic analyses from a shared factual record narrows that exposure, not by suppressing disagreement, but by ensuring that the technical finding on which the economic expert relies is the same finding the technical expert will defend on cross-examination.

The New Frontier: Governance-Failure Litigation

The disputes examined in this study are, for the most part, first-generation AI cases: they concern what a model did, what it copied, what it generated, and how it decided. The litigation that follows may concern something one step removed: not what the model did, but whether the organization deploying it exercised reasonable oversight over a system it knew, or should have known, carried material risk. This is the part of the field where the analysis is necessarily more forward-looking, and the claims that follow are offered as a reading of the trajectory rather than as settled doctrine.

The analogy is to the evolution of cybersecurity litigation. Early data-breach cases asked whether an intrusion had occurred and who was responsible. Over time, the focus of liability migrated upwards, toward the board and officers, and the question became whether the organization had implemented and monitored a reasonable system of controls. There is reason to expect AI litigation to follow a similar arc, because the doctrinal vehicle already exists. Under the Delaware oversight doctrine first articulated in Caremark and reaffirmed in Stone v. Ritter, directors may face liability where they fail to implement a reporting and monitoring system, or, having implemented one, consciously fail to monitor it or to heed red flags.[13] L ater decisions sharpened the doctrine for risks "mission critical" to the business, holding that oversight of such risks must be exercised more rigorously, and extended the duty to officers within their domains.[14]

What makes this more than an analogy is that the first-generation d isputes a lready i n t he study contain recognizable governance-failure claims in embryonic form. Two examples from the corpus illustrate the point. The securities-fraud action against monday.com, pleaded as a misrepresentation of AI capability, is in substance a claim that the company's public statements outran its internal knowledge of the systems it was selling—an information-and-reporting failure of exactly the kind the oversight doctrine addresses. The product-safety action against xAI, pleaded as defective design, is in substance a claim that a known and foreseeable risk was left unmonitored and unmitigated despite the means to address it. Recast at the level of the enterprise rather than the product, each of these is a claim about the adequacy of oversight, and each points toward the derivative suits, fiduciary-duty claims, and professional-negligence theories that a maturing field tends to generate.

[13] *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959 (Del. Ch. 1996); *Stone v. Ritter*, 911 A.2d 362 (Del. 2006).
[14] *Marchand v. Barnhill, 212 A.3d 805* (Del. 2019) (oversight of "mission critical" risks must be "more rigorously exercised"); *In re Boeing Co. Derivative Litig.*, 2021 WL 4059934 (Del. Ch. Sept. 7, 2021); *In re McDonald's Corp. S'holder Derivative Litig.*, 289 A.3d 343 (Del. Ch. 2023) (extending oversight duties to officers).

That progression carries direct consequences for insurance. Governance failure claims are, characteristically, claims against directors, officers, and professional advisers, and they fail to be tested against directors-and-officers, professional-liability, and cyber policies that were not written with AI oversight in mind. The questions now surfacing are whether a misrepresentation-of-AI-capability claim triggers the D&O tower, whether bodily injury exclusions reach AI-enabled product harms, and how cyber and technology error policies interact where an AI system both processes data and makes decisions, each mirrors the coverage litigation that accompanied the maturation of cybersecurity risk a decade earlier. For carriers, the implication is a shift toward pricing AI governance maturity at the point of underwriting; for insureds, it is a growing incentive to document that a reasonable system of controls was in place before a claim arose.

This frontier is where regulatory standards begin to harden into a standard of care. The emerging frameworks—the EU AI Act, the NIST AI Risk Management Framework, and ISO/IEC 42001—provide reference points against which the reasonableness of an organization's oversight may be measured, and non-compliance is increasingly framed not merely as regulatory exposure but as evidence bearing on a breach of duty. In that environment, the expert's work begins to shift from post-hoc reconstruction toward pre-dispute risk engineering: assessing an organization's AI governance maturity before a claim arises and establishing the baseline evidence that reasonable controls existed. Standardized governance-scoring approaches are one emerging tool for that assessment, offering a benchmarkable measure of governance maturity for internal risk management and, in time, for litigation; the AIQ Score™ is one example of such an approach.

Conclusion

A study of more than 425 AI-related disputes suggests that a growing subset of these matters cannot be analyzed through the traditional sequence in which technical experts establish what occurred and economic experts subsequently value the result. In these disputes, the same body of evidence often determines technical causation, legal liability, and economic harm. That interdependence places a premium on expert analysis developed from a shared factual record and capable of withstanding scrutiny in environments characterized by probabilistic systems, incomplete transparency, and evolving standards of proof. As AI-related claims continue to expand beyond intellectual-property disputes and into questions of safety, discrimination, privacy, and organizational oversight, the expert's role is likely to extend beyond reconstruction and valuation to include the assessment of governance systems and controls. The central challenge in many AI disputes, therefore, may be less understanding what a model produced than evaluating the technical, economic, and governance conditions under which it was deployed.

Acknowledgments

J.S. Held would like to thank James E. Malackowski and J. Scott Womack for providing insights and expertise that greatly assisted this research.

More About J.S. Held's Contributors

James E. Malackowski, CPA, CLP, is the Chief Intellectual Property Officer of J.S. Held and the firm's **Intellectual Property** Practice Leader, with shared responsibility for the firm's artificial-intelligence strategy. He is a co-founder and

Senior Managing Director of Ocean Tomo, a part of J.S. Held, where his work centers on intellectual property valuation, strategic advisory, and economic expert testimony. Mr. Malackowski has served as an expert witness on more than one hundred occasions, addressing intellectual-property economics, including valuation, reasonable royalty, lost profits, price erosion, licensing terms, copyright fair use, and injunction equities. His testimony spans U.S. federal, state, bankruptcy, and Chancery courts as well as international venues, including the Ontario Superior Court of Justice and the Federal Court of Australia. A past president of the Licensing Executives Society International, he received the LES International Gold Medal in 2025 and was inducted into the IP Hall of Fame in 2022, where he also received the Q. Todd Dickinson Award.

James can be reached at
james.malackowski@jsheld.com or
+1 312 327 4410.

J. Scott Womack, PMP, is a Senior Director in the Office of the Chief Intellectual Property Officer at J.S. Held, working with the Ocean Tomo team. His practice combines intellectual property valuation with the practical implementation of AI governance, advising clients on operationalizing AI governance metrics and board-oversight structures. His work bridges the technical, financial, and governance dimensions of AI risk—translating between the subject-matter experts who reconstruct how AI systems behave and the boards and executives responsible for overseeing them.

Scott can be reached at
scott.womack@jsheld.com or
+1 678 971 7932.

> ***Disclosure:*** *Co-author James E. Malackowski serves as a member of the board of AIQA Global LLC, the creator of the AIQ Score™, referenced in this article as one example of an emerging standardized governance-scoring approach.*

Transparency Statement: Artificial intelligence tools were used to support the development of this thought leadership content, including research support, information synthesis, and editorial refinement. The analysis, insights, perspectives, and conclusions presented are those of the author(s) and reflect their professional expertise and judgment. The content was reviewed by the author(s) for accuracy, relevance, and consistency with the source materials. Artificial intelligence tools did not independently determine the content's conclusions or recommendations.