



The Liability Squeeze and the Governance Response:

How Documentation Becomes Leverage

AI Governance in Education Series

Memorandum No. 7

Ryan James Purdy

Purdy House Publishing & Consulting

January 2026

Working Paper

Abstract

Educational institutions face an emerging structural problem in AI adoption. Major technology vendors have structured their agreements to cap liability at nominal sums, while insurers are simultaneously excluding AI-related claims from standard coverage or conditioning specialty coverage on governance documentation most schools cannot produce. The result is an expanding zone of institutional exposure that neither vendors nor insurers will cover.

This memorandum documents the two-sided liability squeeze, examines the information asymmetry that sustains it, and argues that governance documentation provides leverage to change the dynamic. The analysis draws on vendor contract terms from Microsoft and Google, insurance market developments including ISO exclusion endorsements and specialty carrier programs, and emerging regulatory requirements across multiple jurisdictions.

Key findings include: vendor liability caps of \$5 to \$10 for free-tier services create exposure gaps of potentially millions of dollars; commercial AI exclusions are proliferating while education-specific exclusions have not yet materialized, creating a finite window for governance development; and governance documentation serves not merely as compliance artifact but as negotiating leverage in vendor contracts, insurance renewals, and board accountability.

The memorandum concludes that the structural conditions for the liability squeeze are in force, the forcing functions are activating, and institutions that build governance infrastructure proactively will be positioned to negotiate from strength rather than scramble under pressure.

Key Findings

1. The liability squeeze is structural, not rhetorical.

Vendor liability caps (Microsoft: \$5-\$10 for free services, 12 months fees for enterprise; Google: \$1,000 for free tier) combined with proliferating insurance exclusions create an expanding zone of uninsured institutional exposure. This is not a temporary market condition but a stable equilibrium that will persist until schools develop countervailing documentation capacity.

2. Commercial AI exclusions are in force; education-specific exclusions are not yet.

ISO endorsements CG 40 47, CG 40 48, and CG 35 08 allow carriers to exclude AI-related claims from general liability. Proprietary exclusions from Berkley, Hamilton, and others are already deployed. However, education-sector carriers (United Educators, CM Regent) have not announced AI-specific products or exclusions, creating a window for governance development before requirements tighten.

3. Governance-dependent underwriting is emerging.

Munich Re's aiSure program, Armilla AI's Lloyd's-backed coverage, and Beazley's approach all condition coverage on demonstrated governance. This model will extend to education. Organizations that build governance documentation now position for specialty coverage as it becomes available.

4. Documentation is leverage, not just compliance.

Schools with governance infrastructure can present vendor governance questionnaires, require contractual addenda, document evaluation criteria, and create evidence trails that matter in negotiations, insurance interactions, board accountability, and litigation. Schools without it accept default risk allocation.

5. The pre-certification window is finite.

Ohio (July 2026), Colorado (June 2026), and EU AI Act (August 2026) deadlines establish hard compliance requirements. Insurance market evolution will establish soft deadlines. Organizations acting now retain design flexibility; organizations waiting will inherit standards designed without their input.

6. Collective action creates upstream pressure.

Individual districts lack leverage. When multiple institutions adopt common governance frameworks and ask common vendor questions, the cumulative effect creates market pressure that reshapes vendor behavior. The pattern mirrors how SOC 2 became effectively mandatory for enterprise technology sales.

Contents

Abstract

Key Findings

Series Context

Scope, Method, and Limits

1. The Two-Sided Squeeze

1.1 The Contract Lever: Vendor Liability Caps

1.2 The Insurance Lever: Coverage Exclusions

2. The Asymmetry Problem

2.1 Vendor Information Advantage

2.2 Insurer Information Advantage

2.3 The Institutional Disadvantage

3. Governance as Leverage

3.1 Leverage in Vendor Relationships

3.2 Leverage in Insurance Negotiations

3.3 Leverage in Board Accountability

4. Structural Indicators: Where the Squeeze Manifests

5. The Pre-Certification Window

5.1 Why the Window Exists

5.2 Why the Window Will Close

6. The Feedback Loop

7. Implications

7.1 For Districts and Schools

7.2 For EdTech Vendors

7.3 For Insurers

8. Conclusion

Appendix A: Evidence Categories for Insurance Readiness

Appendix B: Vendor Governance Questions

References

About the Author

1. The Two-Sided Squeeze

This memorandum documents how vendor liability caps and insurance market evolution are creating structural exposure for educational institutions, and how governance documentation provides leverage to counter it.

Educational institutions face an emerging structural problem that few have yet recognized. The vendors supplying AI tools have structured their agreements to cap liability at nominal sums. Insurers are simultaneously excluding AI-related claims from standard coverage or conditioning specialty coverage on governance documentation most schools cannot produce. The result is an expanding zone of institutional exposure that neither vendors nor insurers will cover.

This section examines the two sides of the squeeze: the contract lever that limits vendor accountability, and the insurance lever that restricts institutional protection.

1.1 The Contract Lever: Vendor Liability Caps

Major technology vendors have structured their AI product agreements to minimize their own exposure while maximizing institutional adoption. The liability caps embedded in these agreements are not negotiating positions. They are standard terms that most schools accept without modification, often without full awareness of their implications.

Microsoft's Liability Structure

Microsoft's liability terms vary across agreement types, creating complexity that few school administrators fully understand. The Microsoft Software License Terms cap liability at \$5.00 USD for standalone software products. The Microsoft Services Agreement, which governs consumer cloud services, limits liability to fees paid during the month in which the breach occurred, or \$10.00 for free services. Enterprise and Education agreements executed through volume licensing typically cap liability at fees paid during the twelve months preceding the incident.^[1]

For schools using Microsoft 365 Education with Copilot features, the applicable cap depends on licensing tier and specific product terms. A district paying \$5,000 annually for Microsoft 365 Education licenses faces maximum vendor liability of \$5,000 for any AI-related harm. A district using free Copilot Chat features may be limited to \$5 or \$10 in recoverable damages regardless of actual harm caused.

Microsoft does provide one significant protection: the Copilot Copyright Commitment offers uncapped indemnification for intellectual property infringement claims arising from AI-generated content. However, this protection applies only to paid Copilot versions, requires use of built-in content filters and safety systems, and covers only copyright claims. It does not extend to claims based on accuracy, hallucinations, discriminatory outputs, privacy violations, or fitness for educational purpose.^[2] Put simply: the commitment protects Microsoft from copyright lawsuits, not schools from educational harm. The protection that exists is narrow; the exposure that remains is broad.

Google's Liability Structure

Google Workspace for Education terms create a particularly significant gap for free-tier users. The standard agreement caps Google's liability at the greater of \$1,000 USD or fees paid in the twelve months preceding the claim. For the many schools using Google Workspace for Education Fundamentals at no cost, this means maximum liability recovery of \$1,000 regardless of the scope of harm.^[3]

Google's generative AI indemnification follows a similar pattern to Microsoft's: it requires paid tier usage and excludes free services. The Google Workspace Service Specific Terms state that AI-related indemnification applies only to services "paid for by Customer and not subject to credits or free tier usage."^[4] Schools using Gemini features through free Google Workspace for Education Fundamentals may have no AI-specific indemnification whatsoever.

Both vendors explicitly disclaim warranties, presenting AI services "as is" and "with all faults." Both exclude consequential damages. Both require schools to assume responsibility for verifying AI outputs. The language varies; the effect is consistent.

The Scale of the Gap

To understand what these caps mean in practice, consider potential AI-related claims in educational contexts:

Scenario	Potential Damages	Vendor Cap (Free Tier)
AI admissions tool produces discriminatory outcomes affecting 500 applicants	\$500,000 - \$5,000,000+	\$1,000 (Google)
AI tutoring system provides harmful advice to student	\$100,000 - \$1,000,000+	\$5 - \$10 (Microsoft)
AI system exposes student data through prompt injection	\$200,000 - \$2,000,000+	\$1,000 (Google)
AI-generated content defames student or staff member	\$50,000 - \$500,000+	\$5 - \$10 (Microsoft)

The gap between potential damages and recoverable amounts is not a rounding error. It is orders of magnitude. Vendor liability caps ensure that virtually all financial exposure from AI-related harm remains with the institution.

The Opt-Out Illusion

One counterargument deserves examination: schools can disable AI features if they find the liability exposure unacceptable. This is technically true. Microsoft 365 administrators can disable Copilot through the Admin Center, license management, Group Policy, or

security group exclusions. Google administrators can toggle Gemini features on or off per organizational unit through granular Admin Console controls.^[5]

However, practical pressures create functional lock-in that makes opt-out unrealistic for most institutions. Default settings enable AI features. Vendor marketing promotes AI capabilities as competitive advantages. Students and teachers develop expectations based on AI availability. Peer institutions adopt AI tools, creating competitive pressure. Productivity benefits create dependency that makes removal disruptive.

Schools may have the technical ability to disable AI. They lack the practical ability to avoid it entirely while remaining competitive and meeting stakeholder expectations. The opt-out exists in theory; lock-in exists in practice. And the liability exposure exists regardless of whether schools consciously chose to accept it.

1.2 The Insurance Lever: Exclusions and Conditions

While vendors limit their liability through contract terms, insurers are simultaneously restricting coverage for AI-related claims. This restriction operates through two parallel mechanisms: excluding AI from standard coverage lines, and conditioning specialty coverage on governance documentation.

Standard Lines Are Opting Out

ISO (Insurance Services Office) and Verisk introduced AI exclusion endorsements for commercial general liability policies, with forms becoming available for market use in late 2025 and early 2026:

Endorsement	Effect
CG 40 47	Artificial Intelligence Exclusion (broad)
CG 40 48	Artificial Intelligence Limited Coverage
CG 35 08	AI-related exclusions for specific contexts

These endorsements allow carriers to explicitly exclude AI-related claims from general liability coverage that might otherwise apply. The exclusions address what insurers call "silent AI" exposure: the risk that existing policies could be interpreted to cover AI-related claims they were never priced to include.^[6]

Proprietary exclusions are also proliferating. Berkley Insurance introduced an "Absolute AI" exclusion for Directors and Officers, Errors and Omissions, and Fiduciary Liability policies. The exclusion applies to claims "based upon, arising out of, or attributable to any actual or alleged use, deployment, or development of Artificial Intelligence."^[7] Hamilton Insurance Group has implemented Generative AI Exclusions that specifically name ChatGPT, Bard, Midjourney, and DALL-E.^[8]

These exclusions are not education-specific. They apply across industries. But they create a coverage environment in which standard liability policies increasingly do not respond to AI-related claims.

Specialty Lines Are Conditioning on Governance

As standard coverage excludes AI risk, a specialty market is emerging to address it. However, this specialty market does not simply offer coverage to anyone willing to pay premium. It conditions coverage on demonstrated governance.

Munich Re's aiSure program provides AI performance guarantee insurance, but access requires passing technical due diligence. Munich Re uses external domain experts to assess AI systems before offering coverage. According to program documentation, approximately 90% of companies that pass assessment receive coverage offers. The inverse is equally significant: companies that cannot demonstrate adequate governance do not receive offers.^[9]

Armilla AI, the first Lloyd's coverholder dedicated solely to AI liability, launched affirmative AI Liability Insurance in April 2025. The product covers AI underperformance, hallucinations, and critical errors. Armilla's underwriting explicitly incorporates NIST AI Risk Management Framework alignment as an assessment criterion. Coverage depends on governance posture, not merely premium payment.^[10]

Beazley has not announced AI exclusions but takes a governance-dependent underwriting approach. Coverage availability and terms depend on demonstrated "good business practices" and documented oversight mechanisms. The company's Head of US Cyber and Tech Underwriting has indicated that underwriting increasingly examines how organizations govern AI use, not merely whether they use it.^[11]

The Education Gap

Despite these market developments, no specialty carrier has developed education-specific AI liability products. United Educators, the dominant education-sector insurer serving 1,600+ member institutions with coverage up to \$35 million, publishes AI risk guidance but has not announced AI-specific products or exclusions.^[12] CM Regent, another significant education carrier, shows no public AI coverage position.

Education-specific carriers appear aware of AI risk. They are advising member institutions on governance. But they have not yet built the underwriting infrastructure to price education AI risk, nor have they developed products that address education-specific exposures such as admissions algorithm bias, student data privacy, or academic integrity violations.

This creates a peculiar gap. Commercial AI exclusions are proliferating. Commercial AI specialty coverage is emerging but requires governance documentation. Education-specific carriers have not yet acted in either direction. The result is uncertainty: education institutions do not know whether their existing coverage will respond to AI claims, and cannot yet access specialty coverage designed for their sector.

1.3 The Uninsured Middle

The structural result of vendor liability caps plus insurance market bifurcation is an expanding zone of uninsured institutional exposure.

When an AI system causes harm in a school, the liability allocation operates as follows:

The vendor's exposure is capped. Microsoft, Google, and other major vendors have structured agreements that limit their liability to nominal sums ranging from \$5 to twelve months of fees, depending on product and licensing tier. Copyright indemnification exists for paid tiers but does not extend to most categories of educational AI harm.

Standard insurance may not respond. General liability policies increasingly exclude AI-related claims through ISO endorsements or proprietary exclusions. Even where exclusions have not been attached, coverage disputes are likely as insurers contest whether policies priced before widespread AI adoption should cover AI-related losses.

Specialty coverage may not be available. Governance-dependent underwriting means that organizations cannot simply purchase AI coverage. They must demonstrate governance maturity that most educational institutions have not developed. And education-specific specialty products do not yet exist.

The institution absorbs the residual. When vendor caps, insurance exclusions, and specialty coverage prerequisites all limit external risk transfer, the remaining exposure sits with the institution. This is not a theoretical concern. It is a structural condition produced by the documented intersection of vendor contract terms and insurance market evolution.

The squeeze is mechanical, not rhetorical. Contract lever plus insurance lever equals institutional exposure. In the absence of leverage, schools absorb risk. But governance, when documented, can begin to reverse that dynamic. The question is not whether this exposure exists. The question is what institutions can do about it.

2. The Asymmetry Problem

The liability squeeze persists because of a fundamental asymmetry: vendors and insurers have documentation, expertise, and leverage that schools lack. This section examines how information asymmetry enables risk transfer to institutions, and why the squeeze will persist until schools develop countervailing documentation capacity.

2.1 Schools Lack Documentation to Negotiate

When a vendor presents standard contract terms with \$1,000 liability caps, a school needs documentation to argue for better terms. Effective negotiation requires the ability to articulate what risks the school faces, what controls the school has implemented, what governance the school expects from vendors, and what alternative vendors exist if terms are unacceptable.

Most schools cannot produce this documentation because the underlying governance infrastructure does not exist.

The Consortium for School Networking (CoSN) 2024 EdTech Leadership Survey of 981 district technology leaders found that 54% of districts lack any separate AI use policy.^[13] Of those with policies, most address student acceptable use rather than institutional governance, vendor management, or accountability structures. The governance infrastructure that would support contract negotiation, specifically documented risk assessments, vendor evaluation criteria, oversight protocols, and evidence of implementation, remains absent in most districts.

This absence is not unique to AI governance. It reflects broader patterns in educational technology procurement. Schools typically lack dedicated legal resources to negotiate technology contracts. They lack technical expertise to evaluate vendor risk claims. They lack collective bargaining infrastructure to aggregate purchasing power across districts. And they lack the time and administrative capacity to conduct the due diligence that effective negotiation requires.

The result is acceptance of default terms. Vendors know that most schools will sign standard agreements without modification. They structure those agreements accordingly.

2.2 The Information Asymmetry Vendors Exploit

Vendor liability caps are aggressive precisely because vendors understand the asymmetry they face. Microsoft and Google employ teams of attorneys who draft contract language designed to minimize vendor exposure. They have data on how their AI systems perform, what failures occur, and what claims arise. They understand the regulatory landscape across jurisdictions. They know that individual schools lack the resources to contest standard terms.

This asymmetry operates at multiple levels:

Legal expertise asymmetry. Major vendors have dedicated legal teams that specialize in technology contracts and liability allocation. Most school districts rely on general counsel, if they have in-house counsel at all, who handle technology contracts alongside employment disputes, facilities issues, and board governance. The specialization gap means vendors understand contract implications that schools may not recognize.

Technical knowledge asymmetry. Vendors understand how their AI systems work, what failure modes exist, and what risks those failures create. Schools often adopt AI tools without detailed understanding of underlying models, training data, or technical limitations. They cannot effectively negotiate risk allocation for risks they do not fully understand.

Market information asymmetry. Vendors negotiate thousands of contracts and understand market standards, competitive positioning, and what terms other customers accept. Individual schools negotiate technology contracts infrequently and lack visibility into what terms are negotiable or what other districts have achieved.

Collective action asymmetry. Microsoft and Google are single entities that can establish uniform terms across all customers. Schools are fragmented across thousands of independent districts, each negotiating individually. Even when state education agencies establish master contracts, individual districts often lack awareness of available terms or capacity to invoke them.

The asymmetry is self-reinforcing. Vendors can impose aggressive liability caps because no individual school has leverage to demand otherwise. Schools accept aggressive caps because they lack information about alternatives and resources to pursue them. The pattern perpetuates across contract cycles.

2.3 Without Evidence, Schools Absorb Risk by Default

The asymmetry means that risk flows downhill to the least documented party.

Vendors disclaim liability through carefully drafted contract terms. They can point to specific language, negotiation history, and industry standards to defend their position. When disputes arise, vendors have documentation to support their liability limitations.

Insurers condition coverage through underwriting requirements. They can point to questionnaire responses, governance assessments, and policy language to support coverage determinations. When claims arise, insurers have documentation to evaluate whether coverage applies.

Schools often have neither. When disputes arise with vendors, schools may lack documentation of what representations were made, what risks were disclosed, or what alternatives were considered. When claims arise with insurers, schools may lack documentation of governance controls, oversight mechanisms, or incident response procedures.

In the absence of documentation, defaults prevail. Contract defaults favor vendors who drafted the terms. Insurance defaults favor carriers who wrote the policies. Schools absorb whatever exposure remains after vendors and insurers have protected their positions.

This is not a market failure that will self-correct through competitive dynamics. It is a stable equilibrium. Vendors benefit from the asymmetry and have no incentive to change it. Insurers benefit from clear underwriting criteria and have no incentive to insure undocumented risks. Schools bear the costs but lack the organized capacity to demand change.

The equilibrium persists until schools develop the documentation infrastructure to shift the dynamic. Without it, schools remain risk absorbers. With it, they can become risk negotiators.

3. Governance as Leverage

The preceding sections established the problem: vendors cap liability, insurers exclude or condition coverage, and schools lack the documentation to push back. This section examines the solution. Governance documentation is not merely a compliance artifact that satisfies external requirements. It is a negotiating tool that changes the power dynamic schools currently lose.

3.1 From Exposure to Leverage: What Documentation Unlocks

A school with documented AI governance occupies a fundamentally different position than a school without it. The difference is not abstract. It manifests in concrete interactions with vendors, insurers, boards, and regulators.

In vendor negotiations:

A school without governance documentation accepts default contract terms because it has no basis for demanding alternatives. It cannot articulate what risks concern it, what controls it expects, or what vendor commitments would address its exposure. The vendor presents standard terms. The school signs.

A school with governance documentation can present specific questions vendors must answer. It can require contractual addenda addressing identified gaps. It can create a paper trail showing good-faith oversight. It can justify rejection of non-compliant vendors to boards and stakeholders with documented rationale.

Consider the difference in practice. A district technology director reviewing a new AI tutoring platform receives a standard vendor agreement with a \$1,000 liability cap. Without governance infrastructure, the director has two options: accept the terms or reject the product. With governance infrastructure, the director has a third option: present the district's AI vendor governance questionnaire, document the vendor's responses, identify gaps between vendor commitments and district requirements, and negotiate specific contract language to address those gaps.

The vendor may refuse to modify terms. But the district now has documentation showing it asked the right questions, received inadequate answers, and made a risk-informed decision. That documentation matters when boards ask what due diligence occurred, when insurers ask what vendor oversight exists, and when incidents prompt questions about what the district knew and when.

In insurance interactions:

Insurance underwriters increasingly expect evidence of governance, not merely assertions of it. Willis Towers Watson guidance to professional liability underwriters notes that carriers now look for evidence of a "risk-aware culture" around AI, including documented policies on how AI is used, what oversight exists, and what controls are in place. Organizations that cannot demonstrate these elements face higher premiums, coverage exclusions, or claim denials.^[14]

A school without governance documentation cannot respond effectively to underwriting questionnaires. It may have informal practices but no evidence those practices exist. It may have policies but no records of implementation. When the questionnaire asks whether human oversight protocols exist for AI-assisted decisions, the school cannot point to documented procedures, training records, or oversight logs.

A school with governance documentation can respond to questionnaires with structured evidence. It can demonstrate the "risk-aware culture" underwriters seek. It can appeal exclusions or premium increases by pointing to documented controls. It can position for specialty coverage as governance-dependent products become available for the education sector.

The *Travelers v. International Control Services* case illustrates the stakes. In that matter, the insurer alleged material misrepresentation regarding multi-factor authentication controls. The insured had attested to MFA implementation on its insurance application. Post-breach investigation revealed MFA protected only the firewall, not other digital assets as represented. The matter concluded via stipulation declaring the policy void from inception.^[15]

The lesson extends beyond cybersecurity to AI governance. Self-attestation without documentation is vulnerable to post-incident scrutiny. Organizations that claim governance controls must be able to demonstrate those controls existed and operated as described. Documentation is not bureaucratic overhead. It is the evidence that makes attestations defensible.

In board and stakeholder accountability:

Every board will eventually ask the question that exposes governance gaps: "Can we prove any of this?"

A superintendent pointing to an AI policy document has answered only the first question. The follow-up questions are harder. Can we prove staff were trained on the policy? Can we prove vendors were evaluated against documented criteria? Can we prove human oversight occurred for high-stakes AI-assisted decisions? Can we prove incident response procedures were tested?

A school without governance documentation cannot answer these questions. The policy exists. Evidence of implementation does not. When the board asks what the district knew and what it did, the answer is: "We had a policy." That answer satisfies no one when an incident has occurred.

A school with governance documentation can answer with timestamped evidence. Training records show completion dates. Vendor evaluation forms show assessment criteria and findings. Oversight logs show review occurred before decisions were finalized. Incident response documentation shows procedures were followed.

The documentation does not prevent incidents. It demonstrates that the institution exercised reasonable care before incidents occurred. That evidence doesn't just check

a box. It protects funding, mitigates reputational damage, and keeps governance power in institutional hands.

In litigation and regulatory review:

When AI-related claims reach litigation or regulatory investigation, the central question is what the institution knew and what it did. Institutions with governance documentation can point to affirmative efforts: risk assessments conducted, questions asked, procedures implemented, training provided. Institutions without documentation must rely on testimony about informal practices that may or may not be credited.

The Houston Independent School District attempted to roll out AI-generated lesson content district-wide. The result was widely criticized as inadequate, with obvious quality issues and lack of human oversight. The initiative was halted after public backlash.^[16] Los Angeles USD and San Diego USD both experienced difficulties with AI procurements that lacked clear vendor roles, updated contracts, or transparency, leading to project failures and stakeholder complaints.^[17]

These cases did not reach formal litigation. But they illustrate the reputational and operational consequences of governance gaps. Institutions that could point to documented evaluation criteria, pilot testing protocols, and quality review procedures would have been better positioned to defend their decisions or catch problems earlier.

3.2 The Evidence Base Insurers Now Expect

The shift from self-attestation to documented evidence is not theoretical. It reflects concrete changes in how underwriters evaluate risk. Willis Towers Watson's guidance on AI risks in professional liability identifies specific evidence categories underwriters examine:^[14]

Evidence Category	What Underwriters Look For
Policy documentation	Written AI use policy addressing scope, permitted uses, prohibited uses
Oversight mechanisms	Defined review processes for AI outputs before consequential use
Training records	Documentation that staff understand AI limitations and responsibilities
Vendor governance	Assessment criteria, contract terms, ongoing monitoring
Incident response	Defined procedures for AI failures, escalation paths, notification protocols
Control monitoring	Evidence that stated controls actually operate as described

The guidance emphasizes that underwriters expect to see evidence of a "diligent corporate policy... outlining how and when AI is used, monitored, and controlled." Organizations that cannot produce such evidence face underwriting consequences.

This shift mirrors what occurred in cybersecurity over the past decade. Early cyber insurance relied heavily on self-reported security practices. As claims experience accumulated, underwriters learned that self-reports did not reliably predict actual security posture. The market evolved toward requiring documented evidence: vulnerability scan results, penetration test reports, security awareness training completion records, and third-party attestations like SOC 2.

AI governance is following the same trajectory. The underwriting questionnaires being deployed in 2025 and 2026 are the early stage of a longer evolution. Organizations that build documentation infrastructure now position themselves for a market that will increasingly demand it.

3.3 Upstream Pressure Through Collective Action

The leverage potential extends beyond individual negotiations. When multiple institutions adopt common governance frameworks and ask common questions, the cumulative effect creates upstream pressure on vendors.

Consider the dynamic if thirty school districts use the same governance framework to evaluate AI vendors. Each district sends the same vendor governance questionnaire. Each district documents vendor responses using the same criteria. Each district makes procurement decisions based on documented findings.

Vendors face a choice. They can refuse to answer governance questions and lose procurement decisions to competitors who engage. Or they can develop governance documentation that responds to the questions districts are asking.

The effect compounds over time. Vendors that invest in governance documentation can point to it as competitive differentiation. Vendors that refuse to engage find themselves excluded from governance-mature districts. The market begins to reward governance maturity rather than treating it as optional overhead.

The collective effect also reduces systemic risk. When districts screen vendors using common governance criteria, they filter out providers with weak controls before breaches occur. Risk pools and consortia benefit from fewer claims across member institutions, not just better documentation after incidents.

This dynamic already operates in cybersecurity. Enterprise buyers routinely require SOC 2 reports, penetration test results, and security questionnaire responses as procurement prerequisites. Vendors that cannot produce these artifacts lose deals. The requirement originated with sophisticated buyers, spread through market pressure, and became effectively mandatory for enterprise sales.

Education AI governance can follow the same path. The precondition is that enough institutions adopt common frameworks and ask common questions. Individual districts

lack leverage. Collective adoption of governance standards creates leverage that individual action cannot achieve.

Independent assessment frameworks become de facto procurement standards when adopted at scale. The question is whether education institutions will develop such frameworks proactively or inherit them from external sources under pressure.

4. Structural Indicators Matrix

The liability squeeze is real, but it is not uniform. Different pressure types are at different stages of realization. Understanding this gradient helps institutions prioritize action and calibrate urgency appropriately.

4.1 The Gradient of Pressure

Three categories distinguish the current state of various forcing functions:

In Force: These conditions are documented and operational. They apply now, regardless of whether institutions have recognized or responded to them.

Emerging: These conditions are activating. Early adopters are experiencing them. Broader application is imminent but not yet universal.

Forecasted: These conditions are anticipated based on documented trends. They have not yet materialized but evidence suggests they will within the 2026-2028 window.

4.2 Structural Indicators by Category

Indicator	Status	Evidence	Affects	Implication
Vendor liability caps	In Force	Microsoft \$5-\$10 (free) to 12-mo fees; Google \$1K-\$25K	All institutions	No vendor backstop
Vendor warranty disclaimers	In Force	"As is," "with all faults," no fitness warranty	All institutions	Vendors disclaim accuracy
Copyright indemnification gaps	In Force	Paid tiers only; excludes non-IP claims	Free-tier users	Most harm categories unprotected
Commercial AI exclusions (GL)	In Force	ISO CG 40 47/48, Berkley, Hamilton	Commercial policyholders	Standard coverage excludes AI
Education-specific AI exclusions	Forecasted	No announcements from UE, CM Regent	Education institutions	Window exists before carriers act
Governance-dependent underwriting	Emerging	Munich Re aiSure, Armilla, Beazley	Specialty coverage seekers	Coverage requires documentation
AI underwriting	Emerging	WTW guidance, broker	Renewwin	Evidence

questionnaires	ng	reports	g policyholders	demands increasing
State AI policy mandates	Emerging	Ohio (July 2026), Tennessee, California	Districts in mandating states	Hard deadlines approaching
EU AI Act obligations	In Force / Emerging	Prohibited (Feb 2025); High-risk (Aug 2026-27)	EU-exposed institutions	Phased requirements now
Colorado AI Act	Emerging	Effective June 30, 2026	Colorado deployers	Disclosure, risk management
Procurement governance requirements	Emerging	Post-PowerSchool demands for SOC 2	Vendors to enterprise districts	Governance becoming prerequisite
Education AI specialty products	Forecasted	No products announced	Education institutions	Market opportunity gap
Litigation precedent	Emerging	Hingham, Yale, IXL cases	All institutions	Liability allocation being tested
Insurer-required certifications	Forecasted	No education requirements yet	Future policyholders	ISO 42001 may become prerequisite

4.3 Reading the Matrix

Several patterns emerge from the matrix:

Vendor-side pressure is fully in force. Liability caps, warranty disclaimers, and indemnification gaps are not emerging trends. They are current contract terms that apply to every institution using major platforms. The exposure exists today.

Insurance-side pressure is bifurcated. Commercial AI exclusions are in force. Education-specific exclusions are not yet documented. Governance-dependent underwriting is emerging but not yet universal. The window for building governance before insurance requirements tighten remains open but is closing.

Regulatory pressure has hard dates. Ohio's July 2026 mandate, Colorado's June 2026 effective date, and the EU AI Act's August 2026 general applicability are not forecasts. They are published deadlines. Institutions subject to these requirements face compliance obligations on specific dates regardless of readiness.

Procurement pressure is accelerating post-breach. The PowerSchool breach created immediate procurement requirement changes. Schools now require SOC 2 compliance and cybersecurity insurance evidence from vendors. Similar dynamics will likely follow AI-related incidents, creating sudden procurement requirement shifts that unprepared vendors cannot meet.

The pre-certification window is finite. No education-specific AI certification exists today. This creates design flexibility for institutions building governance now. When certification requirements emerge, whether from insurers, regulators, or accreditors, institutions with established governance can adapt. Institutions starting from scratch will build under pressure.

4.4 Strategic Implications of the Gradient

The gradient suggests differentiated responses based on institutional position:

For institutions with no AI governance: The vendor exposure is immediate. Begin with AI inventory and policy development to establish baseline documentation. Prioritize evidence that responds to emerging insurance questionnaires.

For institutions with policies but limited documentation: The gap is implementation evidence, not policy existence. Focus on training records, vendor evaluation documentation, and oversight logs that demonstrate policy operates as described.

For institutions with documented governance: The opportunity is leverage. Use governance documentation in vendor negotiations, insurance renewals, and board reporting. Consider whether current documentation meets emerging regulatory requirements.

For EdTech vendors: Procurement pressure is emerging and will accelerate. Governance documentation is becoming a sales prerequisite for enterprise districts. Vendors that build documentation proactively gain competitive advantage.

The matrix is not static. Indicators currently marked "Emerging" will move to "In Force." Indicators marked "Forecasted" will begin to emerge. The trajectory is toward increased documentation requirements across all categories. The question is whether institutions build documentation infrastructure ahead of requirements or scramble to build it after requirements arrive.

5. The Pre-Certification Window

The structural indicators matrix reveals a critical pattern: education-specific AI governance requirements are forecasted but not yet in force. This creates a window of opportunity that will not remain open indefinitely.

5.1 No Certifying Body Exists for Education AI Governance

As documented in Memorandum 4 of this series, the infrastructure for formal AI governance certification in education does not yet exist. ISO/IEC 42001, the international standard for AI management systems, provides a framework but presents significant barriers to educational adoption:

Barrier	Impact
Cost	\$30,000 to \$108,000 for certification
Timeline	6 to 12 months for implementation and audit
Expertise	No education-specific auditor capacity has developed
Scope	Standard addresses enterprise AI systems, not educational contexts

No education-sector equivalent of healthcare's HIPAA compliance ecosystem or financial services' SOC 2 attestation infrastructure exists for AI governance. Schools that seek external validation of their AI governance have no sector-appropriate certification to pursue.

This absence creates both risk and opportunity. The risk is that institutions delay governance development while waiting for certification infrastructure that may take years to emerge. The opportunity is that institutions building governance now can shape how certification criteria eventually apply to their context rather than inheriting criteria designed without their input.

5.2 Three Paths Are Emerging

Memorandum 5 identified three paths emerging to fill the pre-certification gap:

Path A: Reputation-based assessors. Individual assessors and consulting firms build credibility through track record and published work. Quality varies significantly. No standardized criteria govern who can conduct assessments or what methodologies they must use. This path is already active, with assessment services available from multiple providers.

Path B: Voluntary certifying bodies. Industry associations or coalitions establish frameworks and credential assessors. This path is not yet operational for education AI governance, though analogues exist in adjacent domains such as student data privacy

(Student Data Privacy Consortium) and educational technology evaluation (Digital Promise).

Path C: Insurance-employed specialists. Carriers employ or contract domain experts for technical due diligence integrated into underwriting. Munich Re's aiSure program demonstrates this approach, using external specialists to assess AI systems before offering coverage. The model is operational in commercial contexts but has not yet extended to education.

Path C is particularly significant for the thesis of this memorandum. Munich Re's approach validates that insurers will invest in domain expertise to price AI risk they cannot otherwise measure. The company uses "different domain experts depending on the specialty being insured."^[18] Education is a specialty that lacks such experts positioned within the insurance ecosystem.

This gap represents both a market opportunity and a service need. Carriers seeking to underwrite education AI risk need domain experts who understand educational regulatory frameworks, operational constraints, and risk profiles. Educational institutions seeking specialty coverage need assessors who can translate their governance practices into evidence insurers recognize.

5.3 Design Flexibility vs. Inherited Standards

Organizations acting during the pre-certification window retain design flexibility that later adopters will not have.

Design flexibility means:

Institutions choose governance frameworks aligned with their values and operational context. They build relationships with assessors on terms they help define. They influence how criteria apply to educational settings through early adoption and feedback. They position for smooth transition as markets formalize because their infrastructure already exists.

Inherited standards mean:

Institutions adopt frameworks designed without their input. They build documentation under deadline pressure when requirements arrive. They engage assessors assigned by external parties rather than chosen through relationship. They react to standards rather than shape them.

The difference is not merely procedural. It affects whether governance serves institutional purposes or merely satisfies external requirements. Institutions that build governance proactively integrate it with educational mission. Institutions that build reactively treat it as compliance overhead.

The window is finite. Ohio's July 2026 policy mandate, Colorado's June 2026 AI Act effective date, and the EU AI Act's August 2026 high-risk system requirements establish hard deadlines for institutions subject to those jurisdictions. Insurance market evolution

will establish soft deadlines as governance-dependent underwriting becomes standard. The only choice is when to act, and whether to act from a position of control or in response to someone else's clock.

6. The Feedback Loop

The forcing functions examined in this memorandum do not operate in only one direction. Pressure flows downstream to schools, but governance-mature schools can also exert pressure upstream to vendors. This bidirectional dynamic creates opportunities for institutions that recognize it.

6.1 Downstream to Schools

Schools face converging pressures from multiple sources:

Regulatory deadlines. Ohio requires AI policies by July 2026. Tennessee has enacted similar requirements. California mandates AI literacy curricula. Colorado's AI Act takes effect June 2026. The EU AI Act classifies educational AI as high-risk with obligations phasing in through 2027. These deadlines are published and approaching.

Insurance evolution. Underwriting questionnaires increasingly include AI governance questions. Commercial AI exclusions are proliferating. Governance-dependent specialty coverage is emerging. Even where education-specific changes have not occurred, the broader market trajectory is clear.

Board accountability. High-profile AI failures in peer districts create board awareness. Houston ISD's AI content backlash, LAUSD's chatbot collapse, and the PowerSchool breach all generate questions boards ask their own administrations. "What are we doing about AI governance?" becomes a standing agenda item.

Stakeholder expectations. Parents expect schools to manage AI risks affecting their children. Teachers expect clear guidance on AI use in instruction. Students expect fair treatment by AI-assisted systems. Community members expect responsible stewardship of public resources. These expectations create accountability pressure independent of formal requirements.

The convergence matters. A school that addresses only regulatory compliance may still face insurance friction. A school that satisfies insurance requirements may still face board questions. Comprehensive governance addresses all channels simultaneously because the underlying documentation serves multiple purposes.

6.2 Upstream to Vendors

Schools are not merely recipients of vendor terms. They are customers whose collective purchasing decisions shape vendor behavior.

When governance-mature districts require vendor attestations as procurement prerequisites, vendors must respond or lose business. The PowerSchool breach demonstrated this dynamic in cybersecurity: schools nationwide now require SOC 2 Type II compliance and cybersecurity insurance evidence from vendors. Vendors that cannot produce these artifacts face procurement exclusion.

AI governance will follow the same pattern. Districts that require documented AI governance from vendors create market pressure. Vendors that invest in governance documentation gain competitive advantage. Vendors that refuse to engage lose procurement decisions to competitors who can answer governance questions.

The effect is multiplicative when districts coordinate. State education agencies that establish AI governance procurement criteria apply pressure across all districts in their jurisdiction. Purchasing consortia that include governance requirements in master contracts apply pressure across all member districts. Risk pools that condition membership on governance standards apply pressure across all participating institutions.

Individual districts lack leverage. Collective action creates leverage that reshapes vendor behavior.

6.3 Convergence: Insurance Requirements and Educational Duty-of-Care Align

A critical insight from this analysis: insurance requirements and educational duty-of-care are converging rather than conflicting.

Underwriters do not want compliance theatre. Theatre does not reduce claims. Underwriters want operational infrastructure that actually prevents incidents and contains damage when incidents occur.

What reduces claims? Incident response training that ensures staff know what to do when AI systems fail. Reporting pathways that surface problems before they escalate. Professional development that helps educators understand AI limitations. Named officers with bounded authority who can make decisions without delay. Human oversight protocols that catch errors before they cause harm. Vendor governance that screens out providers with inadequate controls.

These are the same elements that constitute genuine educational governance. A school with real incident response capability produces incident response logs as a byproduct. A school with real professional development produces training records as a byproduct. A school with real oversight protocols produces decision trails as a byproduct.

The documentation is not the point. It is the exhaust of functioning systems.

This convergence means schools building governance for educational reasons are simultaneously building insurance readiness. Schools building governance for insurance reasons are simultaneously improving educational practice. The forcing functions align rather than conflict. Institutions that recognize this alignment can pursue unified governance strategies rather than maintaining separate compliance tracks.

7. Implications

This memorandum has documented the liability squeeze, examined the asymmetry that sustains it, and argued that governance documentation provides leverage to change the dynamic. This section translates analysis into action for the four audiences this work addresses.

7.1 For School Districts and Boards

Immediate actions:

Conduct an AI inventory. What AI systems are in use across the district? What data do they process? What decisions do they influence? Many districts cannot answer these questions because no systematic inventory exists. The inventory is the foundation for all subsequent governance.

Review vendor contracts. What liability caps apply to AI-related claims? What indemnification exists or does not exist? What governance attestations were required during procurement? Most districts have not examined existing contracts through an AI governance lens.

Assess documentation state. Can the district demonstrate its governance posture with evidence? Do training records exist? Are vendor evaluations documented? Can human oversight be proven for high-stakes AI-assisted decisions? The gap between policy and evidence is where liability exposure concentrates.

Identify regulatory deadlines. What requirements apply by when? Ohio districts face July 2026. Colorado districts face June 2026. EU-exposed institutions face August 2026. Deadlines create urgency that general best practices do not.

Strategic positioning:

Build documentation infrastructure before external requirements dictate terms. The pre-certification window offers design flexibility. Use it.

Use governance assessment as a vendor negotiation tool. The third option identified in Section 3 requires documented governance criteria. Develop those criteria and apply them.

Prepare for insurance questionnaires with structured evidence. The evidence categories in Appendix A map to what underwriters increasingly expect. Build documentation that responds to those categories.

Position for specialty coverage as it becomes available. Education-specific AI products do not exist today. When they emerge, governance-mature institutions will be first in line.

Remember: the documentation is not the point. It is the exhaust of functioning systems. Build the systems. The evidence follows.

The board question:

Every board will eventually ask: "What did we know, and what did we do?"

The answer cannot be: "We trusted the vendor."

The answer must be: "Here is our documented governance framework, our risk assessment, our vendor evaluation, our oversight protocols, and our incident response procedures."

Build the documentation that makes that answer possible.

7.2 For EdTech Vendors

The procurement shift:

Governance documentation is becoming a sales prerequisite, not a post-sale compliance exercise. The PowerSchool breach accelerated this shift for cybersecurity. AI governance will follow.

Enterprise districts increasingly require evidence of vendor governance before procurement. Vendors that can answer governance questions win decisions. Vendors that cannot answer lose to competitors who can.

Strategic response:

Build governance documentation proactively. The questions districts will ask are predictable. Appendix B provides examples. Prepare responses before requests arrive.

Prepare attestations that respond to emerging frameworks. NIST AI RMF alignment, ISO 42001 readiness, and documented risk management processes all signal governance maturity. Develop documentation that demonstrates these capabilities.

Consider independent assessment as competitive differentiation. Third-party validation carries weight that self-attestation does not. In a market where most vendors lack governance documentation, assessment creates distinction.

Anticipate that procurement requirements will expand. Post-PowerSchool requirements focused on cybersecurity. Post-AI-incident requirements will focus on AI governance. Build infrastructure before incidents force reactive compliance.

7.3 For Insurance Carriers

The education opportunity:

Education is a high-risk sector under emerging AI regulation with no sector-specific assessment capacity. The market conditions that drove Munich Re's aiSure model apply to education, but no carrier has developed education-specific products.

The gap is substantial: 1,600+ institutions in United Educators alone; thousands more in state risk pools and commercial markets; 54% of districts lacking even basic AI policies; regulatory deadlines creating governance urgency; no specialty AI products addressing education-specific risks.

The need:

Carriers entering this market require domain experts who understand educational regulatory frameworks (FERPA, COPPA, state student privacy laws, emerging AI mandates), educational operational constraints (budget limitations, staffing challenges, academic calendar pressures), educational risk profiles (student data sensitivity, algorithmic decision-making in admissions and discipline, academic integrity), and insurance evidence requirements and underwriting criteria.

This combination rarely exists in a single professional background. The gap between educational expertise and insurance expertise is precisely where assessment services create value.

7.4 Positioning Statement

This memorandum series has documented: the governance gap in education AI (Memorandum 1); the insurance market's bifurcation between exclusions and governance-dependent underwriting (Memorandum 2); the translation problem between regulatory, institutional, and insurance frameworks (Memorandum 3); the structural inaccessibility of ISO 42001 certification for education (Memorandum 4); the case for independent assessment in the pre-certification period (Memorandum 5); the preliminary research framework for specialty carrier engagement (Memorandum 6); and the liability squeeze mechanism and the governance response (this Memorandum).

The analysis points toward a specific market need: education requires sector-specific AI governance assessment capacity that does not currently exist.

The author has developed operational governance frameworks for education, including assessment methodologies that translate between educational practice and insurance evidence requirements, documentation systems that produce the artifacts underwriters expect, and implementation tools that help institutions build governance infrastructure efficiently.

The author is exploring partnerships with specialty carriers seeking domain expertise for this underserved market, EdTech vendors seeking governance differentiation in procurement, and educational institutions seeking assessment services that address their specific context.

Inquiries may be directed to jamespurdy624@gmail.com or via LinkedIn at www.linkedin.com/in/purdyhouse.

8. Conclusion

8.1 The Squeeze Is Structural

The liability squeeze is not a temporary market condition that will self-correct. It is a structural result of how vendor contracts and insurance markets have evolved in response to AI risk. Vendors limit liability because they can. Insurers exclude or condition coverage because they must price risk they can measure. Schools absorb residual exposure because they lack the documentation to do otherwise.

8.2 Governance Is Leverage

The response to the squeeze is not merely compliance with external requirements. It is the development of documentation infrastructure that changes the negotiating position of educational institutions. Schools with governance evidence can push back against vendor terms, respond to insurer questionnaires, satisfy board accountability, and defend decisions under scrutiny. Schools without it accept default risk allocation.

8.3 The Window Is Finite

The pre-certification period offers design flexibility that will not last. Regulatory deadlines in 2026-2027 will transform voluntary governance into mandatory requirements. Insurance market formalization will establish evidence standards institutions must meet. Organizations building infrastructure now choose their approach. Organizations waiting will have approaches chosen for them.

8.4 The Machinery Is Built

The vendor liability caps are in force. The insurance exclusion endorsements are filed. The regulatory deadlines are published. The underwriting questionnaires are being deployed.

The lock hasn't fully clicked for education specifically. But the machinery is built, and it is running.

The question is not whether schools will need governance documentation. The question is whether they will build it proactively, as leverage, or reactively, under pressure.

Appendix A: Evidence Categories for Insurance Readiness

Based on emerging underwriter expectations and the evidence categories identified in Willis Towers Watson guidance, schools should prepare documentation in the following areas:

1. AI System Inventory

Catalog of AI tools in use across the institution; data types processed by each system; decision categories influenced by AI outputs; integration points with other institutional systems.

2. Risk Assessment Documentation

Formal assessment of AI-related risks by system; risk classification methodology; identified vulnerabilities and mitigation measures; periodic review and update records.

3. Governance Policy

Board-approved policy addressing AI use, oversight, and accountability; scope definitions (what AI uses are covered); permitted and prohibited use categories; policy communication and acknowledgment records.

4. Role Assignments

Named individuals with documented authority and responsibility; reporting relationships for AI governance; decision-making authority boundaries; succession and coverage arrangements.

5. Vendor Governance

Evaluation criteria for AI vendors; contract review checklist and findings; attestation records from vendors; ongoing monitoring procedures.

6. Human Oversight Protocols

Intervention points in AI-assisted processes; override authority and procedures; escalation paths for uncertain cases; documentation requirements for oversight decisions.

7. Incident Response Procedures

Defined processes for AI failures; escalation paths and notification requirements; communication templates and protocols; post-incident review procedures.

8. Training Records

Staff AI literacy training completion; role-specific governance training; refresh and update cycles; competency verification where applicable.

9. Stakeholder Communication

Parent notification regarding AI use; student guidelines and expectations; transparency measures for AI-assisted decisions; feedback and complaint mechanisms.

10. Continuous Improvement

Review cycles for governance documentation; update procedures when systems or risks change; maturity progression tracking; external assessment and benchmarking.

Appendix B: Vendor Governance Questions

Schools with documented governance frameworks can present these questions during procurement. Vendors that cannot answer with documentation represent unassessed risk.

AI System Identification

1. What AI or machine learning capabilities are embedded in your product?
2. Which features use generative AI, and which use other AI/ML approaches?
3. What third-party AI models or services does your product incorporate?

Data Governance

4. What data is processed by AI components, and for what purposes?
5. Is customer data used to train or improve AI models?
6. What data retention and deletion policies apply to AI-processed information?

Liability and Indemnification

7. What liability caps apply to AI-related claims under your standard terms?
8. What indemnification do you provide for AI-generated outputs?
9. Are AI-related claims subject to different terms than other product claims?

Governance and Certification

10. What certifications or third-party assessments have you obtained (ISO 42001, SOC 2, etc.)?
11. What internal governance processes exist for AI development and deployment?
12. How do you document and communicate AI system changes to customers?

Bias and Fairness

13. How do you test for bias in algorithmic outputs?
14. What documentation exists regarding bias testing methodology and results?
15. What remediation processes exist when bias is identified?

Human Oversight

16. What human oversight mechanisms are built into AI features?
17. Can customers configure oversight requirements for their implementation?

18. What logging and audit capabilities exist for AI-assisted decisions?

Incident Response

19. What incident response procedures exist for AI failures?

20. How are customers notified of AI-related incidents?

21. What support is available when AI outputs cause problems?

Operational Control

22. Can AI features be disabled without losing core functionality?

23. What configuration options exist for AI behavior?

24. How are AI feature changes communicated before deployment?

References

- [1] Microsoft Services Agreement, Section 13 (Limitation of Liability); Microsoft Software License Terms; Microsoft Product Terms (volume licensing). Liability structures vary by agreement type, with consumer services capped at monthly fees or \$10 for free services, software products at \$5, and enterprise agreements typically at 12 months' fees. Available at: <https://www.microsoft.com/en-us/servicesagreement>
- [2] Microsoft, "Microsoft announces new Copilot Copyright Commitment for customers," Microsoft On the Issues, September 7, 2023. The commitment covers paid Copilot versions with content filters enabled but explicitly excludes claims based on accuracy, fitness for purpose, or non-copyright harms. Available at: <https://blogs.microsoft.com/on-the-issues/2023/09/07/copilot-copyright-commitment-ai-legal-concerns/>
- [3] Google Workspace for Education Agreement, Section 12 (Limitation of Liability). Caps liability at the greater of \$1,000 or fees paid in prior 12 months. Available at: https://workspace.google.com/terms/education_terms.html
- [4] Google Workspace Service Specific Terms (February 2024), Section on Generative AI Features. AI-related indemnification requires paid tier and excludes free and credit-based usage. Available at: <https://workspace.google.com/terms/service-terms/>
- [5] Microsoft Learn, "Manage Copilot in Microsoft 365"; Google Workspace Admin Help, "Turn Gemini features on or off for users." Both platforms provide administrative controls to disable AI features. Available at: <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-overview> and <https://support.google.com/a/answer/13623623>
- [6] Insurance Services Office / Verisk, AI-related endorsements CG 40 47, CG 40 48, CG 35 08. Filed for use beginning late 2025 through early 2026 depending on jurisdiction. Available at: <https://www.verisk.com/>
- [7] Berkley Insurance, Absolute AI Exclusion for D&O, E&O, and Fiduciary Liability policies.
- [8] Hamilton Insurance Group, Generative AI Exclusion endorsement.
- [9] Munich Re, aiSure program documentation. The program uses domain experts for technical due diligence with approximately 90% of passing companies receiving coverage offers. Available at: <https://www.munichre.com/en/solutions/for-industry-clients/insure-ai.html>
- [10] PR Newswire, "Armilla Launches Affirmative AI Liability Insurance with Lloyd's Underwriter, Chaucer," April 2025. Coverage addresses AI underperformance, hallucinations, and critical errors with NIST AI RMF-aligned assessment. Available at: <https://www.prnewswire.com/>

[11] Industry commentary on Beazley's governance-dependent underwriting approach. No formal exclusions announced but coverage terms depend on demonstrated governance practices.

[12] United Educators, "Generative Artificial Intelligence Use Policy or Guidelines for Employees," risk management guidance. UE serves 1,600+ member institutions with coverage up to \$35 million but has not announced AI-specific products or exclusions. Available at: <https://www.ue.org/>

[13] Consortium for School Networking (CoSN), 2024 EdTech Leadership Survey Report, survey of 981 district technology leaders. Found 54% of districts lack separate AI use policy. Available at: <https://www.cosn.org/edtech-topics/edtech-leadership-survey/>

[14] Willis Towers Watson, "Navigating AI risks in Professional Liability," October 2023. Guidance for underwriters on assessing AI governance in professional liability contexts. Available at: <https://www.wtwco.com/>

[15] Travelers Property Casualty Company of America v. International Control Services Inc., No. 22-cv-2145 (C.D. Ill. 2022). Matter concluded via stipulation declaring the policy void from inception after allegations of material misrepresentation regarding MFA implementation.

[16] Education Week and local reporting on Houston Independent School District's AI-generated lesson content initiative and subsequent backlash, 2024-2025. Available at: <https://www.edweek.org/>

[17] Reporting on Los Angeles USD and San Diego USD AI procurement difficulties, including lack of clear vendor roles and contract terms, 2024-2025.

[18] Munich Re, aiSure program documentation. The program uses "different domain experts depending on the specialty being insured" for technical due diligence. Available at: <https://www.munichre.com/en/solutions/for-industry-clients/insure-ai.html>

About the Author

Ryan James Purdy is an independent researcher and writer focused on AI governance in education and other regulated sectors. He brings nearly 30 years of education experience across North America, Europe, and Asia to his current work on AI governance in educational settings. He is the author of three books on AI policy design and compliance in education, focused on practical and economical solutions for institutions navigating AI adoption.

The AI Governance in Education memoranda series represents an ongoing effort to analyze, predict, and prepare for changes in AI governance affecting educational institutions through 2026-2028. The series examines regulatory developments, insurance market evolution, and institutional readiness gaps to help schools, vendors, and policymakers anticipate requirements before they arrive.

About Purdy House Publishing & Consulting

Purdy House Publishing & Consulting is an independent research and consulting practice publishing working papers on AI governance in education and related regulated contexts. The AI Governance in Education series examines gaps between AI governance policy and institutional implementation.

Correspondence: jamespurdy624@gmail.com

LinkedIn: www.linkedin.com/in/purdyhouse