EXECUTIVE SUMMARY

The Insurability Problem in Healthcare AI

A Standards-Based Framework for Underwriting Risk Assessment

Jennifer Shannon, MD Glacis Technologies **Sarah Gebauer, MD** Validara Health

March 2026

The Problem

Healthcare AI is being deployed at scale without insurance coverage designed for its risks. A single AI failure in a clinical setting can simultaneously generate product liability claims against the vendor, malpractice claims against the supervising physician, and enterprise liability claims against the deploying health system. Existing insurance frameworks were not designed for this convergence.

3	43%	$1.5M	? yrs
Domains converge with a single AI failure	Of FDA-cleared AI devices recalled <1yr of clearance	Max HIPAA Tier 4 annual penalty	Until actuarial claims data stabilizes

Why Traditional Insurance Frameworks Fail

- **One event, three domains.** A single AI failure creates simultaneous claims across product liability, professional liability, and enterprise risk, each with different policyholders, carriers, claims logic, and actuarial traditions.
- **No actuarial basis.** Medical device coverage relies on regulatory classifications and claims history that don't exist for AI. The FDA's 510(k) pathway now spans scheduling tools and critical care decision support under the same classification.
- **Misaligned pricing.** Professional liability prices individual physician risk, not AI-amplified caseloads where a radiologist reads 500 scans daily instead of 80. The premium stays the same; the exposure does not.
- **Unbounded enterprise risk.** AI failures can be subtle, systemic, and population-scale. 300,000 prior authorization denials in two months. A hallucinated allergy propagating across a medical record for months.

> **Key Risk Illustration**
>
> A clinical documentation AI reported "95% accuracy" across 40,000 encounters/month. Stratifying errors by clinical consequence revealed 120 notes/month containing hallucinated allergies, invented symptoms, or omitted findings — concentrated in the most complex patients. The headline figure was meaningless for underwriting.

The Standards-Proof Framework

Rather than waiting for actuarial data that won't stabilize for a decade, the framework repurposes international medical device standards that healthcare AI companies already implement for regulatory clearance, converting compliance work into underwriting evidence.

Layer	Primary Domain	What It Assesses
Layer 1 Foundation	Product Liability	Risk management quality across ISO 14971, IEC 62304, ISO 13485, and AI-specific extensions. Scores implementation depth, not binary compliance.
Layer 2 Healthcare-Specific Validation	Professional Liability	Validation requirements by risk tier. **Tier 1** (direct clinical decision support): prospective studies with independent physician review and subgroup stratification. **Tier 2** (documentation, prior auth): harm-pathway-matched validation and minimum physician review time studies. **Tier 3** (administrative): harm pathway analysis before accepting the classification; the prior auth AI labeled "administrative" that carried direct clinical risk and embedded demographic bias is the reason this tier exists.
Layer 3 Continuous Operational Assurance	All Domains	Pre-deployment adversarial stress testing plus post-deployment tamper-evident, inference-level monitoring of drift from baseline intent. Attestation that safety controls actually executed.

Layers 1–2 tell the underwriter what should happen. Layer 3 tells the underwriter what did happen — the evidence infrastructure that makes the other two layers verifiable in operation.

Parametric Coverage: A New Mechanism

Layer 3's continuous operational evidence enables a coverage structure traditional indemnity cannot offer: **parametric triggers** that respond to verified performance failures in near-real time, before failures compound into patient harm.

Instead of indemnifying proven losses after months of discovery, parametric coverage pays a predetermined amount when a measurable threshold breach is verified against a tamper-evident audit trail:

- **Performance threshold breach** — hallucination rate exceeds 0.5% over a 30-day rolling window
- **Equity threshold breach** — authorization approval rates diverge >3% between demographic subgroups
- **Model drift beyond validated bounds** — diagnostic sensitivity drops below clinical validation threshold
- **Safety control execution failure** — required controls failed to execute for a defined percentage of inferences

When a trigger fires, two parallel tracks activate: an **operational response** (revert model, notify institution, conduct chart review) and a **financial payout** sized to cover immediate response costs within days, not months. Parametric coverage complements traditional indemnity; it does not replace it. Unobservable failures, monitoring infrastructure failure, and long-tail liability remain addressed through conventional coverage.

Open standards for runtime attestation infrastructure (such as OVERT, Observable Verification Evidence for Runtime Trust, overt.is) define the risk signal architecture, legal preservation requirements, and independent verification framework that make these parametric structures technically feasible.

Market Opportunity

The coverage gap is present in every health system that has already deployed AI tools.

- **Insurance as procurement gate.** Health systems are conditioning enterprise AI contracts on specified minimum coverage across cyber, technology E&O, professional, and general liability, tiered by clinical risk level.
- **No carrier addresses the convergence.** Munich Re's aiSure, Beazley/Chubb AI cyber, and Coalition/CFC tech E&O address adjacent risks but not the specific healthcare AI convergence of product, professional, and enterprise liability.
- **Regulatory tailwinds.** NAIC AI Model Bulletin adopted in 24+ states. Colorado AI Act (effective June 2026). EU AI Act high-risk classification. FDA AI/ML guidance expected 2026–2028. All create demand for exactly the structured risk evidence this framework generates.
- **First-mover advantages compound.** Early entrants accumulate claims experience, data advantage, and influence over emerging standards — advantages late entrants cannot replicate retroactively.

Key Recommendations

For Insurers

- Build clinical advisory capability now. Healthcare AI underwriting requires physician domain expertise, and takes 12–18 months to develop.
- Pilot with 3–5 companies across risk tiers to calibrate evidence requirements and pricing.
- Require tamper-evident continuous runtime assurance infrastructure, not just design documents and test reports.

For Healthcare AI Companies

- Stratify error rates by clinical consequence, not overall accuracy.
- Build inference-level audit infrastructure proactively. Companies with tamper-evident runtime logs have a fundamentally different claims defense posture.
- Validate in your deployment population, not your study population. FDA clearance establishes substantial equivalence under tested conditions; insurers underwrite the deployment context.

For Health Systems & Payors

- Require proof of insurance as a proxy for risk management maturity — and require the underwriting evidence, not just the policy.
- Require vendors to maintain tamper-evident audit logs accessible to the institution.
- Coordinate AI deployment decisions with malpractice carriers when physician workflows or caseloads change.
- Collect and preserve institutional performance data. It will form the actuarial foundation for future pricing.

For Clinicians

- Know what your AI tool was validated on and whether that population matches yours.
- Document your clinical reasoning independently of AI output. This is your primary protection.
- Treat AI-generated documentation as a draft, not a record.

We are in the interval between deployment and accountability. The frameworks built in this window (by carriers, by AI companies, by clinicians) will determine what accountability looks like, and how to insure it, when it arrives.

Disclosure: The authors have commercial interests in Glacis Technologies and Validara Health respectively. The framework is vendor-neutral and implementable through multiple providers. Full paper available upon request.