PART I — REINSURER ISSUE NOTE (CANONICAL PDF)

Systemic Dependency in Clinical Decision Risk

Implications for Reinsurance Aggregation and Accumulation Control

Author: Ayokunle Ogungbemi
Date: March 2026
Version: 1.0 — Canonical
Status: Discussion note for reinsurance consideration

Status Statement
This document is the canonical version of the paper. Any summaries, excerpts, or references should cite this PDF directly.

---

1. Purpose of this Note

This note is intended to raise an aggregation and accumulation question relevant to reinsurers exposed to clinical professional liability and adjacent lines. It does not propose a product, pricing model, or contractual structure. Its purpose is to identify a structural change in how clinical decision risk is generated and to consider the implications for correlation, accumulation, and ex-ante risk control. By framing shared clinical decision logic as a source of endogenous dependency with delayed discovery, the note positions systemic clinical decision risk as a question of risk representation and capital adequacy, rather than solely of liability attribution, ethics, or model governance.

---

2. Background: A Structural Change in Risk Generation

Clinical decision support systems, including AI-enabled tools, are now deployed at scale using shared models, shared training data, and shared workflow integration patterns. Unlike traditional clinical practice—where error arises from heterogeneous judgement and localised conditions—these systems introduce common decision logic across multiple insured entities.

As a result, error is no longer purely idiosyncratic. Under stress conditions such as data drift, operational pressure, or population shift, identical decision logic can deviate in a correlated manner across sites. Correlation, in this setting, is not incidental noise but an endogenous property of system architecture.

---

3. Why This Matters for Reinsurance

Reinsurance structures are designed to absorb tail risk that remains after diversification at the primary level. This presupposes that underlying losses are not perfectly correlated across cedants.

Where identical decision architectures are deployed widely, this assumption may fail. Loss emergence can become clustered across insureds exposed to the same underlying logic, producing accumulation that is not easily visible through conventional exposure metrics.

From a reinsurance perspective, this raises three concerns:

- **Hidden accumulation**: exposure to the same decision logic may be embedded across multiple cedants without being identified as a common driver.
- **Correlation under stress**: diversification may collapse precisely under the conditions that drive claims.
- **Delayed visibility**: loss emergence may lag underlying deviation, limiting the effectiveness of claim-led risk management.

---

4. Latency and Discovery as Accumulation Risks

A distinguishing feature of this risk class is latency. Correlated decision-output deviation may propagate across sites before downstream harm, claims, or litigation make the issue observable. By the time losses are reported, aggregation may already be embedded across reinsurance portfolios.

This challenges reliance on historical loss data and conventional exposure reporting. From an accumulation perspective, the ability to surface correlated deviation early becomes a risk control variable, not merely an operational detail.

---

5. Limits of Traditional Aggregation Treatment

Traditional catastrophe analogies are imperfect in this setting. Unlike natural perils, correlation here is driven by shared logic rather than shared geography or stochastic events. Under stress, correlation may approach unity, rendering probabilistic diversification ineffective.

As a result, the key question is not how to model the tail more accurately once it exists, but how to prevent correlated deviation from propagating undetected across insureds in the first place.

---

6. Reinsurance-Relevant Question

The aggregation question for reinsurers can be stated as follows:

How should reinsurance structures distinguish between clinical decision risks that remain structurally correlated under stress and those where dependency has been demonstrably reduced through governance, monitoring, and early discovery?

This is not a question about technology adoption or clinical practice. It is a question about how accumulation arises and whether it can be meaningfully damped ex-ante.

---

7. Implications for Underwriting and Portfolio Management

Without prescribing solutions, several implications follow:

- Exposure to shared decision architectures may warrant treatment as a distinct aggregation driver.
- Traditional cedant-level diversification assumptions may overstate risk dispersion.
- Evidence of early discovery and intervention capability may be materially relevant to accumulation control.
- Where such evidence is absent, reinsurers may be unintentionally writing correlated risk under independence assumptions.

---

8. Conclusion

As clinical decision architectures evolve, so too does the structure of associated reinsurance risk. This note suggests that aggregation and accumulation may increasingly be driven by shared logic rather than shared events, and that conventional diversification assumptions may not always hold.

The purpose of this note is not to propose a reinsurance solution, but to ensure that aggregation risk remains visible and governable as the underlying drivers of loss change.