

Systemic Dependency in Clinical Decision Architectures

Implications for Solvency UK Capital Frameworks

Author: Ayokunle Ogungbemi

Date: March 2026

Status: Discussion note for supervisory consideration

1. Purpose of this Note

This note raises a supervisory question concerning the treatment of clinical decision risk under Solvency UK capital frameworks. It does not propose a product, firm, or regulatory intervention. Its purpose is to identify a structural feature of emerging clinical AI deployment that may not be well captured by existing capital assumptions, and to ask how firms should evidence mitigation of that feature where it exists. 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 Change in Dependency Structure

Clinical decision support systems, including AI-enabled tools, are increasingly deployed at population scale using shared models, shared training data, and shared integration patterns. Unlike traditional clinical practice—where error arises from heterogeneous human judgement—these systems introduce common decision logic across sites.

Under stress conditions such as data drift, operational pressure, or population shift, this shared logic creates the possibility of correlated decision-output deviation across otherwise independent insured entities. The resulting dependency structure differs materially from the independence assumptions embedded in standard actuarial treatment of professional liability risk.

3. Why the Standard Formula May Misrepresent This Risk

The Solvency UK Standard Formula is designed to provide a conservative, generalised capital requirement under broad assumptions. One such assumption, implicit across multiple risk modules, is that losses arise independently across insureds, allowing diversification to operate.

Where identical decision logic is deployed at scale, correlation is not exogenous noise but an endogenous property of system architecture. In such settings, diversification does not reliably dampen tail risk. Instead, loss emergence may be clustered under specific stress conditions, with delayed visibility and ambiguous attribution.

KE

This raises a supervisory question: where dependency is structural rather than incidental, does the Standard Formula systematically misrepresent the risk profile?

4. Latency and Discovery as Prudential Concerns

A further 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 loss experience emerges, corrective intervention may be difficult or impossible.

From a prudential perspective, this challenges reliance on claim-led learning and raises the importance of ex-ante discovery mechanisms capable of surfacing correlated deviation early. The absence or presence of such mechanisms may materially affect the true risk profile.

5. Implications for Internal Models and Capital Assessment

Solvency UK permits firms to use full or partial internal models where these better reflect their risk profile than the Standard Formula. The framework already recognises that certain risks—because of structure, governance, or dependency—require bespoke modelling.

The question raised here is not whether firms should adopt any particular control or governance mechanism, but whether where firms can evidence reduction of structural dependency and earlier discovery of correlated deviation, capital requirements should reflect that reduced risk.

Conversely, where such evidence is absent, supervisors may reasonably question whether diversification assumptions remain appropriate.

6. A Supervisory Gap Question

The issue can be stated simply:

How should Solvency UK frameworks 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 AI ethics or patient safety per se. It is a question about risk representation, dependency, and capital adequacy in the presence of shared decision architectures.

7. Areas for Supervisory Consideration

Without proposing specific solutions, the following areas appear relevant to supervisory judgement:

- What forms of evidence are sufficient to demonstrate that correlated decision-output deviation is unlikely to propagate undetected?

- How should latency and discovery capability be reflected in capital assessment?
 - Under what conditions might internal models or partial internal models be more appropriate than the Standard Formula for this risk class?
 - How should ongoing monitoring and change control be treated where decision systems evolve over time?
-

8. Related Work and Positioning

A growing body of literature examines the implications of algorithmic decision-making for professional liability, patient safety, and insurance. Several strands are relevant but incomplete with respect to the question raised in this note.

First, clinical and medico-legal research has examined how clinical decision support systems may alter malpractice exposure through phenomena such as alert fatigue, automation bias, and altered standards of care. This work typically focuses on individual clinician liability, evidentiary challenges, and post-event attribution rather than on aggregation or capital representation.

Second, recent insurance and risk-management literature has begun to address “algorithmic” or “AI-enabled” liability more broadly, noting that a single model error can generate many similar claims and that traditional Errors & Omissions frameworks may be strained by scale and opacity. These discussions correctly identify aggregation potential but are generally framed at the level of product design, pricing, or governance taxonomy, rather than as a prudential capital problem.

Third, quantitative work on algorithmic insurance and liability modelling has highlighted that replacing heterogeneous human judgement with a single model changes the statistical structure of risk. This insight is often illustrated using medical imaging or diagnostic examples, but the analysis is typically directed toward contract design or expected loss estimation, rather than toward the implications for regulatory capital frameworks that rely on diversification assumptions.

Separately, Solvency II and Solvency UK guidance on the Standard Formula and internal models recognise that where a firm’s risk profile deviates materially from standardised assumptions, bespoke modelling may be appropriate. However, this guidance is generally technology-agnostic and does not explicitly address dependency structures arising from shared decision architectures deployed across multiple insured entities.

The contribution of this note is to connect these strands by framing shared clinical decision logic as a source of endogenous dependency with prudential significance. By emphasising correlation under stress, delayed discovery, and the limits of claim-led learning, 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. In doing so, it seeks to surface a supervisory and reinsurance-relevant gap that is not explicitly addressed in the existing literature.

9. Conclusion

As clinical decision architectures evolve, so too does the dependency structure of associated risks. This note suggests that existing capital frameworks may not always surface that dependency explicitly and invites supervisory consideration of how firms should evidence meaningful risk reduction where it exists.

The intention is not to advocate a regulatory outcome, but to ensure that capital treatment keeps pace with changes in how risk is generated, governed, and discovered.

End of note

KE