

Rejecting AI Legal Personhood: Targeted No-Fault Compensation for AI-Attributable Harm in Clinical Practice

A clinician-led, policy-oriented normative legal analysis for England and Wales

"A liability architecture that keeps the clinician, not the algorithm, at the centre of accountability"

Sin Ting CHAN

MBChB (CUHK), FRCR (Clinical Oncology), LLM (Healthcare Ethics and Law); Radiation Oncology Fellow, BC Cancer Surrey, British Columbia, Canada; Alumna, School of Social Sciences, University of Manchester, United Kingdom; Correspondence: stachanmedcareer@proton.me



Aim: To evaluate whether granting legal personhood to AI is necessary or desirable to resolve liability for harms in NHS clinical practice, and if not, to identify, through normative legal analysis, a defensible UK-specific alternative.

Methodology: a doctrinal analysis of UK case law and statute to identify the responsibility gap; a normative evaluation of two reform options (legal personhood; no-fault compensation); a comparative review of EU, US, and South Korean approaches; and analogical reasoning from the Vaccine Damage Payment Scheme as the domestic no-fault precedent.

I: Problem and Legal Gap

II: Core Proposal

III: Policy Implications and Feasibility Analysis

I: Problem and Legal Gap

1. The AI Responsibility Gap

AI is now embedded across NHS practice: Cytosponge for Barrett's screening [1], EchoGo for HFpEF [2], DERM for skin cancer [3], autonomous radiotherapy auto-contouring, and TORTUS ambient-voice AI scribes [4]. The NHS 10-Year Health Plan commits the NHS to becoming "the most AI-enabled healthcare system in the world" [5].

When AI-assisted care causes harm, existing legal frameworks struggle to allocate liability. This is the **responsibility gap**. (i.e. the responsibility gap is the doctrinal space in which an AI causes patient harm but no party (clinician, Trust, or manufacturer) can be held legally liable under existing UK law.)

AI System used

Harm

Who is liable?

The Responsibility Gap: no party liable under existing law

Technical drivers:

- Opacity — black-box decision-making is not traceable [6]
- Hacking vulnerability [7]
- Adaptivity — continual learning destabilises the regulatory "expected safety" benchmark
- Algorithmic bias — under-representative training data produces inequitable outcomes [8]

2. Why Traditional Legal Frameworks Fail

Fault-based frameworks face substantial obstacles in establishing **both** breach and causation.

- Breach.** A clinician acting on a Medicines and Healthcare products Regulatory Agency (MHRA)-registered AI within accepted clinical practice has not breached duty under *Bolam* [9]: the standard is defined by what peers do, not by whether the practice is itself reliable.
- Causation.** Opacity defeats the "but for" test [57]; the causal pathway is not retrievable [6].

Doctrine	Obstacle
Clinical negligence (<i>Bolam</i> [9])	No breach where clinician follows MHRA-approved AI within accepted clinical practice
Informed consent (<i>Montgomery</i> [11])	Patient-centred test for disclosure of material risks; potentially applies to disclosure of AI involvement and its limitations, but does not modify the standard for the underlying clinical decision
Strict product liability (Consumer Protection Act (CPA) 1987, ss 2, 3, 5)	Section 3 "expected safety" test is doctrinally challenging when applied to adaptive AI: the same input may produce different outputs over time as the system retrains
Vicarious liability	No predicate tort where clinician is not personally negligent
Duty of care (<i>Caparo</i> [10])	AI cannot satisfy 'fair, just and reasonable' limb without moral agency (see '4,' of this poster)

The clinician becomes a **'liability sink'**, absorbing residual risk from upstream components [12], and faces temporal moral luck: identical clinical conduct produces no liability if the AI happens to be correct, but potential negligence exposure if it is not [13, 14].

Finding: AI lacks the moral agency required for legal personhood. An alternative is needed to resolve the liability issue in medical litigation involving AI.

Potential Solution: A new, discrete no-fault statutory scheme administered by NHS Resolution provides a workable, NHS-compatible alternative.

5. The Analytical Conclusion

A new AI Harms expert panel within NHS Resolution, established by secondary legislation and structurally separate from Clinical Negligence Scheme for Trusts claims handling, would administer the scheme. It complements ordinary tort: tort rights are preserved, and NHS Resolution retains a right of recovery against any negligent party.

(All thresholds, tariffs, and operational parameters set out in this proposal are indicative only. Final values would be established through stakeholder consultation, government policy development, and parliamentary scrutiny of any implementing legislation.)

II: Core Proposal

6. Example of Existing No-Fault Model: Vaccine Damage Payment Scheme (VDPS)

Both VDPS and the new proposed schemes share core features:

- Harm may occur **without negligence**
- Strict **causal proof** standard, no fault requirement
- Public-confidence rationale, supporting safe technology adoption
- Tax-free statutory payment

The VDPS is administered by the NHS Business Services Authority [26, 36]. Its 60% threshold was set by the Pearson Commission [30] to focus the scheme on cases sufficiently grave that absence of compensation would constitute a public-confidence problem [37]. The same logic applies to AI-attributable harm.

Lessons from VDPS COVID surge: from 27 claims in 2019 to over 22,000 during the pandemic, with administrative cost of approximately £48 million over five years [38]. This argues for **rigorous tiering, transparent criteria, and in-house NHS Resolution** capacity.

7. Eligibility Criteria

The scheme applies where **three** conditions are met:

- Device.** An MHRA-registered AI as a Medical Device (AlaMD) [23] (or as defined by the National Commission's forthcoming high-risk classification [24]) was deployed in NHS care.
- Causation.** The AI's output **materially contributed** to the harm, on the balance of probabilities.
- Severity threshold.** The harm meets a 60% disablement threshold. Harms below this threshold revert to existing pathways.

Following the **Vaccine Damage Payment Scheme** (VDPS) precedent [25, 26], clinician conduct is not part of the eligibility inquiry. Causation, not fault, is the contested question [27, 28].

8. Scope: Examples of Which AI Systems Are Covered

Included (AlaMD with material clinical influence):	Excluded (causal contribution sufficiently attenuated):
- Autonomous and quasi-autonomous clinical decision systems - Adaptive radiotherapy plan generation (e.g., MR-Linac on-couch adaptation) - AI-driven triage executing clinical action without human review - Advisory diagnostic AI in radiology, pathology, and oncology where clinician oversight is mandated but materially constrained	- Administrative AI (scheduling, scribing such as TORTUS [4]) - Decision-support tools where clinicians retain substantial authority to amend AI outputs (e.g., current auto-contouring with mandatory expert review)

11. Operational Model — Claims Process

Step 1 — Trigger

Adverse outcome involving an eligible AlaMD. Voluntary patient referral; mandatory NHS Trust notification under the duty of candour and MHRA post-market surveillance regulations (June 2025) [29].

→

Step 2 — Causation Screen

A new, dedicated NHS Resolution AI Harms expert panel applies the material contribution standard. No requirement to prove fault.

→

Step 3 — Tiered Compensation Determination

Quantum determined administratively according to tiered structure.

→

Step 4 — Learning Loop

Every payment triggers structured incident review via NHS Resolution's Safety and Learning function. Findings to MHRA, Trust, and manufacturer.

Timeline: Aim to be finished within weeks to months, not years. AI Harms expert panel composition: clinical, technical, legal, and patient or lay membership; specialty representation matched to the harm category [58]

III: Policy Implications and Feasibility Analysis

12. Why This Matters for the NHS

The NAO (October 2025) records total NHS clinical-negligence liability of **£60 billion**, second-largest on the government balance sheet after nuclear decommissioning [34]. Annual settlement costs reached **£3.6 billion** in 2024–25 [34]. In low-value claims, legal costs average nearly four times damages [34]. Mental health and radiology have seen the largest percentage increases in claims [34]. As AI deployment in radiology accelerates, the responsibility gap will particularly affect the specialty where claims are already rising fastest.

A targeted scheme:

- Compensates a defined cohort that current frameworks fail
- Reduces legal-cost inefficiency through fixed-tariff determination
- Preserves the existing tort architecture for non-AI claims
- Aligns with the NHS 10-Year Plan [5] and the Topol Review [39]
- Embeds safety learning into compensation events

13. International Comparison: Regulatory Position

Jurisdiction	Approach	Status (2026)
UK	Principles-based, sector-led; MHRA AI Airlock [43]; National Commission Dec 2025 [24]	Commission recommendations due 2026
EU	AI Act 2024 — risk-tiered; high-risk medical AI requires conformity assessment [41]	7 May 2026 Omnibus: delays high-risk to Dec 2027, embedded AI to Aug 2028 [42]
USA	Fragmented state-level [46]; FDA Software-as-a-Medical-Device framework [51, 52]; CDS exemption test [44, 45]	250+ bills in 47+ states [46]
South Korea	AI Basic Act, in force 22 Jan 2026 [47]; risk-based, healthcare in high-impact category [53]	First generative-AI vision-language CXR reporting tool approved Apr 2026 [48]

The UK's principles-based approach suits iterative clinical AI but is inadequate for cross-sectoral liability allocation, hence the National Commission's December 2025 establishment with explicit remit on liability for AI-attributable harm [24]. No jurisdiction has resolved AI-specific liability through bespoke legislation. The EU has partially addressed defective-product liability through the recast Product Liability Directive [15]; New Zealand's accident-compensation framework [50] would absorb AI-attributable harm within its existing no-fault treatment-injury cover. Neither produces the targeted, tiered, supplementary architecture proposed here.

14. Layered Accountability Retained

The scheme does **not** insulate human or institutional actors:

- Clinicians — for selecting, reviewing, and overriding AI outputs
- NHS Trusts — for procurement, validation, and integration
- Manufacturers — for design, lifecycle, and post-market surveillance [29]

Concurrent tort claims are permitted. The model mirrors the Authorised Self-Driving Entity construct under the Automated Vehicles Act 2024 [40].

15. Political Feasibility and Statutory Vehicle

The proposal is deliverable through two plausible routes:

- Faster route:** secondary legislation under the Medicines and Medical Devices Act 2021, expanding NHS Resolution's remit
- More durable route:** primary legislation accompanying the National Commission's 2026 recommendations and the Lock KC review of clinical negligence costs [49]

The threshold-based design follows the VDPS architectural principle that has been politically sustainable since 1979. The current proposed scheme sits within a convergent policy window: NHS 10-Year Plan, National Commission recommendations, Lock Review, and EU AI Act timeline.

16. Safeguards Against Misuse

The scheme is not a blanket shield. Three structural safeguards apply:

- Audit obligation** — Department of Health and Social Care to designate an oversight body for periodic audit; reports published in the public domain
- Continuing duties retained** — manufacturers and Trusts remain bound by MHRA, CPA 1987, and Digital Technology Assessment Criteria framework obligations, consistent with the Patient Safety Commissioner's redress principles [59]
- Court access preserved** — claimants retain the right to litigate; the scheme is supplementary, not exclusionary, unlike New Zealand's accident-compensation model [50]

17. Key Message:

Legal personhood for AI is neither necessary nor desirable. On the analysis defended here, a new discrete no-fault statutory pathway administered by NHS Resolution, operationally distinct from its negligence indemnity functions, is a defensible response to the responsibility gap, protecting patients, preserving human accountability, and aligning with the NHS's ambition to lead in AI-enabled healthcare.

18. Contribution of This Analysis:

This analysis contributes to UK discourse on liability for AI-attributable harm. It defends a tiered statutory no-fault compensation framework, administered by NHS Resolution as a discrete pathway, threshold-restricted to 60% disablement on the VDPS model, and funded jointly through central exchequer and a hypothecated industry levy, as a coherent response to the responsibility gap. Offered for consideration by clinicians, policymakers, regulators, and scholars working on AI integration into NHS practice.

Conflict of Interest The author declares no conflicts of interest.

Acknowledgements This work draws on the author's LLM dissertation (University of Manchester, 2024) and reflects developments to May 2026.

Use of Artificial Intelligence Anthropic's Claude (Opus 4.7) was used during preparation of this poster to assist with structural drafting, literature signposting, citation formatting, and editorial review of the author's analytical content. The underlying legal-doctrinal analysis, clinical interpretation, ethical reasoning, and policy conclusions are the author's own and build on the author's LLM dissertation (University of Manchester, 2024). The author verified all citations and substantive claims against primary sources and assumes full responsibility for the content and accuracy of the work. No AI tool generated primary clinical data or original empirical findings.

References Full reference list available via QR code, or URL: <https://zenodo.org/records/20261589>

Global AI Conference 2026 — AI Governance Section — Poster Code [rcrglobalai2026.14401b2] – Submission Code 147