Medical Protection

Closing the AI liability gap:

AI, safety, and the case for legislative change

June 2026

Executive Summary

The law has always struggled to keep up with technological change. With AI, the pace of change is so rapid that this gap feels less like a step and more like a widening gulf.

A recent White Paper[1]- produced through a collaboration between the MPS Foundation, York University's Centre for Assuring Autonomy, and the Improvement Academy at the Bradford Institute for Health Research - highlights how clinicians could find themselves exposed when their decisions are influenced by AI recommender systems. Such systems analyse patient data and suggest personalised treatment plans, diagnoses, or medications.

There are also concerns about who might be held liable in the event of a claim relating to AI scribes, automated documentation assistants, triage algorithms, and other forms of clinical decision support. These all share a common feature: they shape clinical reasoning, records, and workflows without taking autonomous responsibility for the outcomes.

Under the current legislative framework, there is a risk that doctors could be held wholly liable if an AI suggestion turns out to be wrong and they have followed it. That's because the existing product liability regime was never designed with AI in mind.

Clinicians could be left in a bind. If they reject an AI output and things go wrong, there is a real and significant risk that they could face allegations of negligence if the AI output was deemed to be appropriate. But if they follow an inappropriate recommendation and the outcome is still poor, they risk becoming the 'mark' or the 'liability sink' - the obvious target for a damages claim. All of this despite having little control over, and perhaps only a limited grasp of, how the AI system reaches its conclusions. We think that's inequitable. When harm arises from a defective AI product, system, or algorithm, responsibility shouldn't fall squarely on the clinician's shoulders. Clinicians should be given the same clarity as they would in instances of defective physical medical equipment.

The law needs to evolve. Clinical negligence rules, product liability, and new regulatory frameworks must be updated to ensure accountability is fairly shared with developers, manufacturers, and suppliers where defects in the AI tool have led to demonstrable harm.

We urge the government to develop legislation that would classify AI systems as products subject to strict liability, taking inspiration from the EU's revised Product Liability Directive whilst being tailored to the UK context.

We see four clear benefits from reform:

1. **Unlocking AI's potential.** If doctors can trust the tools at their disposal and be clear about where liability sits if things go wrong, they're more likely to use them. That means faster uptake, greater efficiency, and better patient care - in line with the Government's own ambitions for AI in healthcare.

2. **Introduction of safer AI tools.** Clearer liability arrangements would align responsibility with control. In other high risk and highly regulated sectors such as pharmaceuticals, aviation, and medical devices, this alignment has historically driven higher safety standards. Applied to AI, it would create stronger incentives for developers to prioritise safer design, ongoing testing, and deployment, rather than externalising risk to users or downstream organisations.

3. **Fairer outcomes when things go wrong.** Such reform would strengthen the liability framework by ensuring that responsibility for harm caused by defective AI systems can be properly attributed to developers, manufacturers, importers and distributors, rather than defaulting to clinicians through

[1] The MPS Foundation (2025) *Avoiding the AI 'off-switch': Make AI work for clinicians, to unlock potential for patients*. Available at: The MPS Foundation

clinical negligence claims. This would support a fairer and more transparent allocation of responsibility across the AI supply chain and its end users, and without undue compromise or delay in the primary objective in compensation being paid to patients.

4. **Protection for the NHS and healthcare professionals against the cost of claims when harm arises from a defective AI product.** This would apply to the NHS - where taxpayers ultimately foot the bill - and to clinicians who indemnify or insure themselves against claims.

This paper aims to set out the challenges we expect clinicians will face, and the action policymakers can take now to make sure AI delivers benefits without leaving doctors unfairly exposed.

Case studies

The potential opportunities created by AI for advancements in medicine are exciting and only limited by one's imagination. It is however still important to consider what could go wrong, and the ways in which the use of AI could result in a claim.

Clinical negligence claims can take a long time to arise. While AI has been used in medicine for some time, it may take a number of years before we know the volume and nature of claims that will arise linked to this current period of rapid increase in AI usage. Included below are two hypothetical case studies. These illustrate how there is a real and significant risk that a claim could be brought against a doctor in relation to the use of AI tools. Under the current product liability framework in the UK, there is a risk that clinical negligence claims could be brought against the clinicians in these cases and that they would be held wholly liable.

Case study 1: inpatient care needed following a GP's decision to increase prescription dosage

A 69-year-old woman was on long term warfarin for atrial fibrillation, with her GP practice undertaking regular monitoring and dosing advice. The patient reported a recent short episode of diarrhoea and vomiting, which had subsequently settled.

Recent blood tests had been inputted into an AI system that assisted with warfarin dose titration, which inappropriately recommended an increase in the warfarin dose. The GP followed the recommendation and increased the warfarin dose.

One week later the patient was admitted to hospital with a gastrointestinal bleed requiring surgery and ICU stay and was found to have an INR significantly above the recommended range.

The patient later put in a claim against the GP alleging that the increased dose of warfarin was inappropriate and had resulted in the haemorrhage and hospital stay.

In this case, the AI-generated recommendation was unsafe and inappropriate. Although the GP retains clinical responsibility for prescribing, the AI developer/manufacturer should also bear liability due to the flawed dosing advice.

Case study 2: delayed cancer diagnosis following a flawed x-ray review

A 64-year-old man presented to their GP with a 3-month history of dry cough and was referred for a chest x-ray to exclude malignancy.

The x-ray was initially screened by AI and no abnormality highlighted.

Dr A, a consultant radiologist, reassured by the AI review, then reviewed the x-ray and reported no significant abnormality other than a prominent right pulmonary hilum most likely due to an unfolded aorta.

Due to the encouraging x-ray report and no recommendation for further imaging, the GP treated the patient for possible reflux. The patient's symptoms persisted and 6 months later an x-ray and CT demonstrated a squamous cell carcinoma of the lung with bony metastases. By then, the only treatment option was palliative, and the patient sadly died a few months later.

A claim was raised alleging that had the abnormality been identified at the earlier x-ray, the patient would have had an earlier CT/diagnosis, the option of chemotherapy with the potential for clinical response and tumour shrinkage, and avoidance of a significant period of pain and suffering. Subsequent review of the initial x-ray did reveal an abnormality that should have prompted further investigation.

In this case, the analysis by the AI was incorrect and therefore made no recommendations regarding diagnosis or treatment, so both the clinician (Dr A), and the AI developer/manufacturer should have had a liability to the patient.

These case studies illustrate a recurring problem: where an AI system contributes to harm, current UK law could leave the clinician as the primary focus of a claim, creating a real and significant risk that liability would fall disproportionately on them. This raises concerns about fairness and the long-term sustainability of AI adoption in healthcare. It seems unfair to make clinicians solely responsible given the limited control they have over the design, performance, and ongoing behaviour of these technologies.

Legislative context

Consumer Protection Act 1987

Product liability in Great Britain and Northern Ireland is governed by the Consumer Protection Act (CPA) 1987 and by the Consumer Protection (Northern Ireland) Order respectively, both of which implements the 1985 EU Product Liability Directive.

The Act was drafted for a world of static, physical products and does not readily accommodate AI systems. Some shortcomings of the CPA 1987 when considering AI include[2]:

- **It can be difficult for patients to bring a claim.** The CPA 1987 defines 'product' as primarily tangible goods and AI would not appear to readily qualify as product under this law, which inadvertently makes it hard for patients to bring claims against developers or suppliers of AI tools.

- **Clinicians have the potential to become 'liability sinks'.** The legislation is triggered if a product is 'defective'. With AI tools, it may be difficult to prove how or why a decision was defective. Clinicians are expected to exercise judgement when using AI tools but if a 'defect' occurs, the clinician could absorb

[2] Sproul, P. and Gibson, M. (2025) *Challenges for Consumers in the Existing UK Product Liability Framework – Part 1*. Available at: Lexology

the legal liability without even being aware that a defect occurred at the time. The CPA does not apply to software and so AI developers are 'protected' from CPA liability rules and thus having to explain any defect in their product. This 'loophole' clearly presents a risk to patient safety.

Product Regulation and Metrology Act 2025

The Product Regulation and Metrology Act (PRMA), passed in 2025, is intended to provide a mechanism for the UK Government to manage product safety while giving it legal powers to regulate risks arising from technological developments where those technologies form part of a product or influence how a product behaves. Whilst not explicit, the Act could potentially apply to AI-enabled tools. However, the PRMA falls short of what is required in several areas:

- **The Act defers substantive reform rather than delivering it.** It operates as an enabling framework, delegating broad powers to Ministers to introduce future secondary legislation on product safety and compliance, and with those decisions to be made through secondary legislation subject to more limited parliamentary oversight. While this allows flexibility as technologies evolve, it also postpones key policy decisions and provides limited certainty for sectors such as healthcare. In particular, the Act does not itself address how responsibility or liability should be allocated where AI-enabled products or systems cause harm.

- **AI systems aren't clearly defined as products.** Although Clause 2(2)[3] of the Act includes 'product components' and allows regulations to reference 'technical standards' which could be used to cover software, depending on how secondary legislation is drafted, it does not explicitly state AI systems as unique products. This creates uncertainty over whether decision-support systems in healthcare fall under the definition of a regulated product, particularly in the case of standalone software.

- **It isn't decisive on liability.** Clause 2(3)[4] of the Act does not establish rules for liability, i.e. who is legally responsible when AI outputs cause harm in healthcare. It defines who may be subject to product regulations but does not impose liability or define how liability will be allocated in the event of harm caused by AI-assisted decisions, like those found in the EU's revised Product Liability Directive. This leaves healthcare providers, developers, and insurers operating in a grey zone for now.

Medicines and medical devices do not appear to be the target of the Act. Its powers are framed so as not to displace existing specialist regulatory frameworks. Medicines and medical devices continue to be governed by comprehensive, sector-specific regimes overseen by the Medicines and Healthcare products Regulatory Agency, including the Human Medicines Regulations 2012 and the UK Medical Devices Regulations 2002.

[3] *Product Regulation and Metrology Act 2025*, s. 2(2). Available at: Legislation.gov.uk
[4] *Product Regulation and Metrology Act 2025*, s. 2(3). Available at: Legislation.gov.uk

The EU Product Liability Directive

The new EU Product Liability Directive 2024/22853 is replacing the Council Directive 85/374/EEC and modernises the European product liability framework to take account of emerging technologies, including artificial intelligence.

Key changes that the Directive introduces include:

- **AI 'products' subject to liability regime**. Software, including standalone AI systems and AI components embedded in devices, is expressly included within the definition of a "product." This ensures that AI systems are subject to similar liability regime as physical goods.

- **Defectiveness redefined.** The assessment of whether a product is defective expressly considers characteristics specific to AI, such as post-sale learning, updates, cybersecurity vulnerabilities, and the way a system interacts with users.

- **Broader scope of liable actors**. Liability may attach not only to manufacturers, but also to importers, software developers, fulfilment providers, online marketplaces, and those responsible for substantial post-sale modifications (for example, through software updates).

- **Ongoing responsibility.** The framework clarifies that liability may arise from events occurring after a product has been placed on the market, including unsafe updates or subsequent changes in the behaviour of AI systems that continue to learn.

- **No contractual limits on liability.** Economic operators cannot limit or exclude their liability through contracts or national legal provisions. Any contractual terms seeking to cap or waive liability are prohibited, ensuring that manufacturers cannot use contracts with patients, clinicians, or healthcare organisations to dilute their responsibility for harm caused by defective AI products.

- **Reduced burden of proof for claimants.** Where a claimant faces excessive difficulties in proving defectiveness or the causal link, in particular due to technical or scientific complexity, such as with AI systems, courts may presume defectiveness and/or causation, where it is at least probable. This reduces the evidential burden on claimants in complex cases.

In effect, for EU member states, the new Directive closes gaps that previously left AI outside the traditional product liability regime. It imposes continuing obligations on those who design, update, and distribute AI systems, while potentially making it easier for claimants to pursue redress where harm arises.

Countries within the EU will now have to transpose the provisions of the new Directive into national law by December 2026. With the UK no-longer being a member of the EU, the UK is not required to implement the Directive, and the government has not indicated that it is likely to introduce legislation along these lines.

It is important to note however, the unique position of Northern Ireland. Under the Northern Ireland Protocol (as amended by the Windsor Framework), EU product liability law continues to apply in Northern Ireland, and any updates would take effect automatically. This means the new Directive is expected to apply in Northern Ireland from 9 December 2026, even without formal transposition into Northern Ireland law by Westminster.

This may lead to legal divergence within the UK, with a different product liability framework operating in Northern Ireland compared to Great Britain. This could create uncertainty around applicability, jurisdiction and enforcement of liability rules, including how responsibility is allocated and how claims involving AI systems are brought and resolved in practice.

What needs to happen

We propose that the law of negligence, product liability, and emerging regulatory frameworks should evolve to ensure liability is properly allocated to developers, manufacturers, and suppliers where it can be shown that defects in AI have contributed to demonstrable harm.

Specifically, we urge the UK government to develop legislation that would classify AI systems as products subject to strict liability, taking inspiration from the EU's revised Product Liability Directive whilst being tailored to the UK context.

We see four clear benefits in making such a change:

1. **Unlocking AI's potential.** If healthcare professionals can trust the tools at their disposal and be clear about where liability sits if things go wrong, they're more likely to use them. That means faster uptake, greater efficiency, and better patient care - in line with the Government's own ambitions for AI in healthcare.

2. **Introduction of safer AI tools.** Clearer liability arrangements would align responsibility with control. In other high-risk and highly regulated sectors such as pharmaceuticals, aviation, and medical devices, this alignment has historically driven higher safety standards. Applied to AI, it would create stronger incentives for developers to prioritise safer design, ongoing testing, and deployment, rather than externalising risk to users or downstream organisations.

3. **Fairer outcomes when things go wrong.** Such reform would strengthen the liability framework by ensuring that responsibility for harm caused by defective AI systems can be properly attributed to developers, manufacturers, importers and distributors, rather than defaulting to clinicians through clinical negligence claims. This would support a fairer and more transparent allocation of responsibility across the AI supply chain and its end users, and without undue compromise or delay in the primary objective in compensation being paid to patients.

4. **Protection for the NHS and healthcare professionals against the cost of claims when harm arises from a defective AI product.** This would apply to the NHS - where taxpayers ultimately foot the bill - and to clinicians in private practice who indemnify or insure themselves against claims.

There does seem to be a growing acknowledgement of the need for action to address this issue.

The MHRA is carrying out a review of the regulation of AI in healthcare, including a public call for evidence that seeks views on accountability and how responsibility and liability should be allocated between manufacturers, healthcare organisations and clinicians when AI tools are used.

The Law Commission has also confirmed it is undertaking a review of the product liability regime, including its application to digital products and emerging technologies such as AI. A formal public consultation on proposals for reform is expected in the second half of 2026, with further recommendations to follow.

MPS will engage with both reviews and continue to raise the need for action with policy makers more broadly.

Further information

AI: The Safer Practice Framework

Medical Protection has developed a two-part 'AI safer practice framework' that can help medical professionals to navigate the challenges that might be presented when using AI in practice when it comes to consent, accountability, and clinical judgement. It can also help guide medical professionals in making informed decisions and ensuring records support them in case of any medicolegal challenge.

Further information is available on our website: AI Safer in Practice

Avoiding the AI 'Off-Switch': Make AI Work for Clinicians, to Deliver for Patients

We have been proud to support the publication of a groundbreaking white paper in 2025, developed in collaboration between York University's Centre for Assuring Autonomy, the MPS Foundation, and the Improvement Academy hosted at the Bradford Institute for Health Research

The white paper sets out a range of recommendations for Government, AI developers, and regulators to urgently consider, including:

- Urgent reform to product liability, due to significant difficulties in applying the current product liability regime to an AI tool.
- AI tools to provide information rather than recommendations to clinicians. This will reduce the potential risk to both clinicians and patients, prior to product liability reform.
- Clinicians to be fully involved in the design and development of the AI tools they will be using to ensure they are usable, useful, and safe.

The paper is available on the MPS Foundation website: Avoiding the AI 'off-switch': Make AI work for clinicians, to unlock potential for patients.

About MPS

Medical Protection is a registered trademark and a trading name of The Medical Protection Society Limited ("MPS"). MPS is the world's leading member-owned, not-for-profit protection organisation for doctors, dentists, and healthcare professionals.

Founded in 1892, we exist to protect and support the professional interests of more than300,000 members around the world, helping them to understand and navigate the ongoing challenges of modern practice. Today we proudly support members in the United Kingdom, Ireland, South Africa, New Zealand, Australia, Hong Kong, Singapore, Malaysia, the Caribbean and Bermuda.

Our in-house experts can assist members with the wide range of legal and ethical problems that can arise from their professional practice. This can include assistance with clinical negligence claims, complaints, medical and dental council inquiries, legal and ethical dilemmas, disciplinary procedures, inquests and fatal accident inquiries.

Our philosophy is to support safe practice in medicine and dentistry by helping to avert problems in the first place. We do this by promoting risk management through our workshops, e-learning, clinical risk assessments, publications, conferences, lectures and presentations.

We also actively campaign for regulatory and legal reforms that benefit members and the wider healthcare professions.

We are not an insurance company, and all benefits of MPS membership are discretionary, as set out in the Memorandum and Articles of Association.