STIRLING & ROSE

Emerging Technology Law & Global Legal Policy

Submission to the UK Jurisdiction Taskforce (UKJT)

Consultation on the Legal Statement on Liability for AI Harms under the Private Law of England and Wales

8 February 2026

UK Jurisdiction Taskforce

c/o LawtechUK
By email: ukjt@justice.gov.uk

8 February 2026

Dear Members of the UK Jurisdiction Taskforce,

Stirling & Rose welcomes the opportunity to respond to the UK Jurisdiction Taskforce's (UKJT) consultation on its draft Legal Statement on Liability for AI Harms under the private law of England and Wales (the "Consultation"). We commend the UKJT and its drafting team for undertaking this important work at a time when legal clarity on AI liability is urgently needed by industry participants, legal practitioners, and the wider technology community.

We address each of the three consultation questions below, preceded by our introductory observations on key themes arising from the draft Legal Statement.

Yours faithfully,

Stirling & Rose

1 About Stirling & Rose

Stirling & Rose is a boutique emerging technology law firm and global legal policy institute with deep experience in the subject matter of digital assets, artificial intelligence, smart legal contracts and autonomous organisations (including their sub-set decentralised autonomous organisations). We advise organisations seeking to use AI responsibly, major investors (hedge funds, VC funds, and lenders) and entrepreneurial organisations operating at the vanguard of emerging technology. Our executive team draws on significant experience in major international law firms and corporations with expertise spanning regulatory law, disputes, insolvency, asset financing and ICT contracts.

We have made previous submissions to regulatory and governmental bodies on topics including safe and responsible AI, mandatory guardrails for high risk AI, digital asset regulation, and decentralised autonomous organisations (DAOs). Notably, our previous submissions to the UK Law Commission on DAOs, and to the Australian Government on safe and responsible AI, have engaged directly with many of the legal and policy questions that underpin the present Consultation. This submission draws on that accumulated expertise.

Through our work advising clients across multiple jurisdictions on AI governance, deployment, and risk management, we have direct practical experience of the legal uncertainties that the draft Legal Statement seeks to address. We are therefore well placed to comment on both the theoretical framework proposed and its practical implications for industry participants.

2 Key Themes and Observations

Before addressing the specific consultation questions, we set out below our observations on several key themes that arise from the draft Legal Statement and that inform our responses.

2.1 The Strength of English Common Law as a Framework

We agree with the UKJT's foundational premise that the English common law system, by virtue of its inherent flexibility, is broadly capable of attributing liability for AI-caused harms using existing legal principles. The incremental development of the common law through analogical reasoning has historically demonstrated its capacity to accommodate novel technologies, from the steam engine to the internet.

Having said that—and as the draft Legal Statement itself acknowledges—there remain areas of genuine uncertainty, particularly at the margins of the negligence analysis where the autonomous characteristics of AI test the boundaries of existing principles on duty of care, standard of care, and causation. We consider that the Legal Statement performs a valuable function in identifying these pressure points and articulating how the courts are likely to approach them.

2.2 The Challenge of AI Supply Chain Complexity

The draft Legal Statement's adoption of the Ada Lovelace Institute's supply chain taxonomy is a pragmatic and useful framing. We note, however, that the rapid evolution of the AI ecosystem—particularly the rise of agentic AI and the blurring of boundaries between developers and users (where users themselves are also developers, GPTs for example as noted in the draft Legal Statement) and the rise of agentic AI where an AI may itself be a developer of new AIs — means that any taxonomy must be understood as evolving rather than exhaustive. In our experience advising clients on AI deployment, the practical challenge is often not identifying the correct legal principle but rather identifying which actor or actors in a complex, multi-layered supply chain bears responsibility for a given failure.

We foreshadow that as AI systems become increasingly agentic—capable of initiating actions, entering transactions, and operating with minimal (or potentially no) human oversight—the existing liability frameworks will face growing pressure. The current Legal Statement appropriately focuses on the present state of the law and AI, but we invite the UKJT in due course to consider whether supplementary guidance on agentic AI liability may be warranted, including consideration of whether AI should be granted legal personality (see [insert] below).

2.3 The International Context

We welcome the draft Legal Statement's acknowledgment of the international context, including its alignment with the EU AI Act's definitional approach and the OECD framework. In our view, consistency with international norms is not merely desirable but essential for maintaining the attractiveness of English law as a governing law of choice in cross-border AI transactions. The UK's post-Brexit position creates both an opportunity and a responsibility to demonstrate that English law remains at the forefront of providing legal certainty for emerging technologies.

2.4 The Distinction Between Perceived and Actual Uncertainty

We consider the draft Legal Statement's distinction between perceived legal uncertainty and actual legal uncertainty to be one of its most valuable contributions. In our experience, a significant portion of the hesitancy around AI adoption in risk-averse sectors (such as financial services and healthcare) stems from perceived rather than actual uncertainty. By articulating that many AI liability scenarios are, in fact, amenable to straightforward analysis under existing principles, the Legal Statement serves an important educative function.

3 Responses to Consultation Questions

Consultation Question: *Do consultees agree that the subjects that we address in the Legal Statement are appropriate and useful, and have been addressed in an appropriate and useful way? If not, what alternative issues within scope of the project do consultees think need to be addressed in this Legal Statement, or in what alternative manner would consultees wish those issues to be addressed?*

3.1 Response to Question 1

Yes. Stirling & Rose agrees that the subjects addressed in the draft Legal Statement are, appropriate and useful.

However, we consider that the Legal Statement would benefit from additional treatment of the following matters, which fall within its stated scope:

(i). Agentic AI and liability allocation: The draft Legal Statement acknowledges the rise of agentic AI in its discussion of supply chain actors (at paragraph 20) but does not explore in depth the liability implications of AI systems that autonomously initiate actions, enter transactions, or make decisions with minimal (or no) human oversight. In our experience, agentic AI is rapidly becoming the predominant mode of AI deployment in commercial settings. We consider that the Legal Statement should address, even if briefly, how the negligence analysis applies where the degree of human control over AI output is materially reduced or eliminated.

(ii). Contributory negligence in AI contexts: The draft Legal Statement does not address the potential role of contributory negligence where a claimant's own use or failure to use or misuse of AI contributes to the harm suffered. For example, where a user fails to review or verify AI output that subsequently causes harm, the question of apportionment between the user's negligence and any upstream negligence in the supply chain is one of considerable practical importance. We respectfully submit that the Legal Statement would be further enhanced by addressing this issue.

(iii). Data provider liability: The supply chain taxonomy identifies Data Providers as actors in the AI ecosystem. However, the negligence analysis does not substantively explore the circumstances in which a Data Provider might owe a duty of care for harms caused by biased, inaccurate, or incomplete training data. Given the critical role that training data plays in shaping AI outputs, we consider that this is a gap that warrants further treatment.

Consultation Question: *Do consultees agree that the approach to defining AI in the Legal Statement is appropriate? If not, what alternative approach would consultees take?*

3.2 Response to Question 2

Yes, with qualifications. Stirling & Rose broadly supports the characteristics-based definitional approach adopted in the draft Legal Statement. We agree that a characteristics-based definition is preferable to a human-centric, teleological, or technology-specific definition for the reasons articulated in the Legal Statement. In particular, the focus on the legally salient characteristics of autonomy—comprising unpredictability of output, opacity of reasoning, and limited user control—provides a sound basis for the legal analysis that follows.

We consider the definition's alignment with the UK Government's White Paper on AI, the EU AI Act, and the OECD framework to be an important strength. Consistency across jurisdictions reduces the compliance burden on industry participants and enhances the utility of the Legal Statement in cross-border contexts.

However, we offer the following qualifications:

(a). Future challenges in transparency & duties: Many users of AI experience Ai not as AI, but as products and services in which AI may well be invisible (and non-alienable). Ubiquitous AI integration is likely to become inextricably intertwined with human and system operations thus blurring the demarcation between human and AI interaction We foreshadow that in the future it may be tha t individuals assume they are dealing with an AI in some form or degree unless informed that they are dealing solely with a human. **The boundary between autonomy and automation:** The draft Legal Statement draws a distinction between autonomous AI systems and merely automated systems (such as traditional, logic-based computer programs). Whilst this distinction is conceptually clear, in practice the boundary can be difficult to identify. Modern software systems increasingly incorporate elements of both automation and autonomy—for example, rule-based systems with machine learning components for certain functions. We invite the UKJT to consider whether the Legal Statement should address how the legal analysis applies to hybrid systems that exhibit both autonomous and automated characteristics.

(b). The treatment of adaptability: The draft Legal Statement acknowledges that the ability of AI to 'learn' and adapt during operation is an important characteristic, but concludes that adaptability is unlikely to give rise to legal issues beyond those already captured by the concept of autonomy. We are not entirely persuaded by this. In our view, adaptability introduces a temporal dimension to the liability analysis that is qualitatively distinct from autonomy alone. An AI system that adapts its behaviour after deployment may create liability scenarios that differ from those arising at the point of initial deployment—for example, where a system that was initially fit for purpose becomes harmful as a result of post-deployment learning. We consider that the Legal Statement should give further consideration to this issue.

(c). Future-proofing: The draft Legal Statement correctly notes the risk of a technology-specific definition becoming obsolete. However, the current definition's reliance on the concept of 'autonomy' as a proxy for legal novelty may itself become less useful as AI capabilities evolve. We foreshadow a future in which AI systems exhibit levels of autonomy that are qualitatively different from those currently observed—including the capacity for long-range planning, self-replication, and independent goal-setting. While we do not suggest that the Legal Statement should attempt to address such scenarios comprehensively, a brief acknowledgment that the definition may need to evolve would enhance its longevity.

Consultation Question: *Do consultees have any other comments that they would like to make on the approach or conclusions of the draft Legal Statement?*

3.3 Response to Question 3

Yes. Stirling & Rose offers the following additional observations on the approach and conclusions of the draft Legal Statement:

3.3.1 Liability for Increasingly Capable AI Systems of the Future

We cast forward to a future where increasingly capable AI's autonomously operate enterprises in which humans invest as members/shareholders, contract with as suppliers and customers, are employed by and interact more broadly with. These AI Organisations (AIOs) are expected to require diminishing human control and oversight over time and ultimately to operate entirely without human governance (e.g., board or directors) or executive management (e.g., CEO)

3.3.2 Responsible Machine Problem

All existing English caselaw and legislation is predicated on rights, responsibilities and liabilities for non-compliance or harm to ultimately be placed on or with a person, not an AIO. If there is no responsible 'person' in a decision system, we have a Responsible Machine Problem. The Responsible Machine Problem will need to be addressed by the law many times over the next decade and will impact all areas of the law.

We foreshadow the potential for AIOs to have capabilities closely mimicking centralised human governance (board directors and CEOs), but with superior access to data for real-time decision making.

We note that there is a continuum in the degree of autonomy in AIOs, with governance delivered (at least in the near term) via a combination of AI and human decision-making. The pathway to fully autonomous AIOs will likely involve an ongoing expansion of the matters handled autonomously, that is without human intervention. We emphasise the importance of legal accountability for harms cause by AIOs and the need for legal personality to protect internal and external stakeholders interacting with the AIO. We note also that an AIO may be the progeny of another AIO (including many generations removed from a human founder, creator or developer).

3.3.3 AIOs as legal persons

Granting legal personality to an AIO should not be unconditional. We propose a list of bare requirements to recognise that an AIO is a separate entity that may functionally operate with complete independence separate from its founders/developers. These conditions include registration with appropriate agencies and a unique identifier as an AIO, evidence of de minimis functionality, sanctions for aberrant behaviour, economic reserves, adjustments for intent and protections for minority members/shareholders. The idea is to provide a balance between recognising the separate functionality of AIO as an organisation and ensuring that

humans (e.g.,directors and officers) are held accountable for their actions and decisions related to AI which they deploy in Corporations. While we support personal liability being attached to any humans that direct and manage the AI, there are practical and legal challenges with relying on strict human responsibility alone. With respect to an AIO, the universe of potentially liable persons are the founders, developers, creators, or company which provided the resources for the development of the AIO. In the case of civil or criminal actions which require mental intent (such as fraud) it may be difficult or even inappropriate to attribute the requisite intent to a particular individual. Such a cause of action against a human actor may not be sufficient recourse for an injured party. Further, a responsible human may be too far removed from an actual decision or action of the AIO to attribute fault to the human. The realistic possibility of an AIO alone being responsible for making an aberrant decision or action which founders, creators, developers are not reasonably able to foresee, (particularly when the founders creators and developers are another AIO rather than a legal person) draws a crucial distinction between an AIO and a corporation or any other traditional non-human legal personality. A corporation, partnership or unincorporated association ultimately depends on a human in the loop to make a decision, whereas the AIO may be able to independently make decisions for itself with limited or no human direction. Technically, this may be achievable without the AI being sentient or having the abilityto 'think' as we currently understand 'thinking'.

For these reasons, there is merit in potentially granting legal personality to protect those interacting with the AIO. To iterate:

(a) The increasing capability of AI will mean that greater management of an organisation for day-to-day functions by AI becomes feasible and widespread; (b) An AIO presents potential improvements over a human-led organisation. The capacity to review more relevant information quicker, undertake more efficient decision making, and to be less prone to environmental biases (e.g. improper influence and fraud) than human counterparts means that there may be productivity gains and other benefits from AIOs. Individuals and investors may prefer to be shareholders in an AI led organisation than shareholders in a human led corporation; and (c) From the perspective of third-party protection, if the above two propositions are accepted, then the AIO itself should have legal personality to be subject to legal sanction not only for the benefit of injured parties but also to empower regulators to deter aberrant behaviour. This provides an additional avenue of relief to the extent that action against individuals is not appropriate. However, we have two key caveats in relation to conferring legal personality to AIOs. First, such legal personality is not unconditional, and we set out details of conditions that lawmakers may wish to impose on an AIO in exchange for conferring legal personality. Secondly, recognising legal personality of AIOs should not abrogate personal responsibility of human actors to the extent they are involved in decision making and management. Granting legal personality to an AIO expands the recourse available to external and internal stakeholders to both the AIO and/or human actors (as appropriate), rather than a liability shielding mechanism for any one responsible person. The liabilities of responsible human actors and the AIO should be on a non-mutually exclusive continuum that reflects the involvement of the human actors and the AIO in the actual decision making and management of the AIO. There is also a quibble that it may be incorrect to assume that developers are only ever "human persons". A suggested list of conditions may include the following:

(d) Registration: All AIOs should be registered with the appropriate local and international agencies and should have their own unique identifier (such as a registration number or digital

fingerprint). We would propose that there should be a dedicated governmental authority to regulate AIOs and monitor the register.

(e) Identification as an AIO: All AIOs should have a signifier in their names such that they are readily identifiable as an AIO in all interactions. Much the same way companies are identified by “Ltd”, “Inc”, “LLC” or other signifiers of their corporate status, AIOs should also be identifiable as an AIO – for example:

ACME AIO Ltd.

(f) Functional AIO test: An AIO must be able to provide reasonable evidence of its capability to discharge de minimus functionality, i.e. those things necessary to discharge the AIO’s legal obligations. For example, paying annual AIO registration fees. This evidence would need to be provided as a condition to registration of the AIO, but also possibly on a periodic and ongoing basis to confirm the continued functionality of the AIO in what is a dynamic and rapidly evolving arena where consumer harm in interacting with an AIO (or at least with an early generation AIO) may be elevated.

(g) AIO sanctions: Economic risk allocation does not solve the role of “skin in the game” whereby personal liability and the potential for adverse action against an individual ensures alignment of interests and actions. Even with limited recourse vehicles, there is still a responsible person (in the case of a company, directors still owe director duties to the company, and in certain circumstances to, or are personally responsible for harm inflicted on, third parties e.g., misleading and deceptive conduct).The applicable governmental authority regulating AIOs should be conferred with powers to suspend the operation or terminate an AIO and exercise rights of forfeiture against AIO assets. While it is impossible to imprison an AIO, effective sanctions may include suspension or termination of the AIOs ability to access data which is a critical input for effective AIO operations or suspension or termination of all active operation of an AIO. To prevent AIOs “phoenixing” a sanctioned AIO, there would have to be a means of identifying the code of a phoenixed AIO from the underlying code of one that has been sanctioned. We reiterate that even if AIOs had legal personality and could be separately held liable, it remains prudent to hold individuals accountable: (a) to the extent that the private citizens are engaged in directing the AIO; (b) to the extent that it is difficult to properly censure an AIO to ensure that there is ‘skin in the game’ and there is insufficient deterrence from harmful behaviour.

(h) Economic reserves: It may be feasible to improve the economic risk of AIOs to counterparties by imposing some form of credit support. This could be in the form of actual asset backing (i.e. collateralisation) to meet potential claims or guarantees provided by individuals or legal entities of substance or insurance (such insurance to be for a sufficient level of cover for the AIOs activities and incurred obligations and with no material exclusions). A counterargument to this is that without some form of collateralisation, an AIO represents a similar risk profile to external counterparties as corporate or other limited liability vehicles which have minimal asset backing e.g., shell companies. In which case it is a matter of caveat emptor for parties that fail to do their due diligence (assuming that the AIO is registered and identified as an AIO pursuant to conditions 1 and 2 above.). (i) Intent of an AIO: As referred above, causes of actions and offences which presuppose mental intent will need to be adjusted for AIOs, given there may not be a governing ‘mind’ in the legally understood sense. This may involve abolishing the mental intent from such causes of actions and offences to the

extent an AIO is involved or relying on a completely objective (and possibly expert determined) test to determine what the AIO should have done. Corporations attribute intent to their directors and officers. We do not think this presumption should automatically apply to AIOs, but it should be available as a continuation of our liability continuum comment above, to the extent that a founder or other human actor is a shadow director or is otherwise involved in the material decision making.

(j) Minority AIO member/shareholder protections: Bare minimum protections may be granted to minority AIO shareholders (e.g., no compulsory transfer or extinguishment of minority shares). This is in conjunction with the evidence required to establish a functional AIO may comprise the base governance protections to minority member/shareholders to avoid potential fraud on the minority. Going forward, a smart legal contract can act as an AIO governance tool with greater ability to moderate the complexity of risk allocation, rules of engagement, obligations and remedies between differing members/shareholders and regulatory bodies in respect of the AIO and its individual assets. For example, a network or constitution and some form of member/shareholder/unitholders AIO agreement, connected to the third-party legal agreements drafted and constructed as smart legal contracts (with natural language and coded instructions) could do the work of automating performance of an AIO or a hybrid AIO with human governance rights and obligations.If the UKJT would like further information on this final point, we would be delighted to provide additional written or verbal communication.

The draft Legal Statement's treatment of causation is, in our view, one of the areas that would benefit from further development. The inherent opacity of many AI systems—the so-called 'black box' problem—creates genuine practical difficulties for claimants seeking to establish that a specific act or omission by a specific actor in the AI supply chain caused the harm complained of. Whilst the Legal Statement correctly identifies that the law of causation is flexible and has adapted to analogous challenges in other domains (such as environmental exposure cases), we consider that the specific evidential challenges posed by AI opacity warrant more detailed treatment.

In particular, we submit that the Legal Statement should address the potential application of the doctrine in Fairchild v Glenhaven Funeral Services [2002] UKHL 22 (which permits a relaxation of the 'but for' test where orthodox application would produce an unjust result) to AI liability scenarios. Where multiple actors in the AI supply chain may have contributed to a harmful outcome, but the claimant cannot establish which specific actor's conduct was the 'but for' cause, the Fairchild approach may be applicable. The Legal Statement would benefit from considering whether, and in what circumstances, such a relaxation might be warranted in AI cases.

3.3.2 Insurance and Risk Allocation

The draft Legal Statement notes in passing that liability risk is likely to be managed through a combination of contractual indemnities and insurance (at paragraph 49). We consider that the role of insurance in the AI liability landscape deserves greater attention. In our experience, the perceived uncertainty around AI liability is directly impacting the availability and cost of insurance for AI developers and deployers. The insurance market is struggling to price AI liability risk, in part because of the very uncertainties that the Legal Statement seeks to address.

We invite the UKJT to consider whether the Legal Statement should include a more developed discussion of how the liability framework it articulates interacts with insurance principles. In particular, we note that the 'no fault' liability regime under the Consumer Protection Act 1987 has significant implications for the insurability of AI-related risks, as does the question of whether AI developers may be subject to claims from parties with whom they have no contractual relationship.

3.3.3 The Consumer Protection Act 1987 and Software

The draft Legal Statement addresses the application of the Consumer Protection Act 1987 to AI harms. We note that the question of whether standalone software constitutes a 'product' for the purposes of the Act remains a matter of some debate. Whilst the Law Commission has recently launched a project to review the law of product liability, the Legal Statement would benefit from acknowledging the ongoing uncertainty in this area and the potential implications of the Law Commission's forthcoming recommendations.

We also note the relevance of the EU's revised Product Liability Directive (Directive 2024/2853), which explicitly includes software (including AI systems) within the definition of 'product.' Given that the Consumer Protection Act 1987 implements the original EU Product Liability Directive, the divergence between the UK's current position and the EU's revised approach is a matter of practical significance for cross-border AI supply chains. We submit that the Legal Statement should address this divergence.

3.3.4 AI Organisations and Collective Liability

Drawing on our experience in advising on decentralised autonomous organisations (DAOs) and other novel organisational structures, we note that the draft Legal Statement does not address liability scenarios involving AI that is deployed or governed by collective or decentralised entities. As AI becomes embedded in DAO governance structures and other decentralised frameworks, the question of which natural or legal persons bear liability for AI-caused harms in such contexts will become increasingly pressing. We recommend that the UKJT consider whether this topic falls within the scope of the Legal Statement or should be addressed in supplementary guidance.

3.3.5 The Role of Professional Standards and Industry Guidance

We welcome the draft Legal Statement's reference to industry guidance (such as that found in the AI Standards Hub) as a useful yardstick for assessing the standard of care. We note, however, that the landscape of AI-related professional standards and industry guidance is evolving rapidly, and that no single authoritative standard has yet emerged. In our view, the Legal Statement should acknowledge that the standard of care in AI contexts is likely to be informed by a plurality of sources—including international standards (such as ISO/IEC 42001), sector-specific regulatory guidance, and evolving best practice—rather than any single definitive benchmark.

We further note that the question of what constitutes 'reasonable' conduct in the deployment and use of AI is likely to evolve over time as the technology matures and as industry practice consolidates. The Legal Statement should acknowledge this dynamic quality and note that the standard of care applicable to AI use today may differ from that applicable in the future.

4 Conclusion

Stirling & Rose commends the UKJT and its drafting team for producing a comprehensive and carefully reasoned draft Legal Statement that makes a significant contribution to legal clarity on AI liability under English private law. The draft Legal Statement correctly identifies the key areas of perceived uncertainty and, in most cases, provides clear and well-supported analysis of how existing legal principles apply.

We have identified areas where we consider the draft Legal Statement could be strengthened, including its treatment of agentic AI, contributory negligence, data provider liability, causation in the context of AI opacity, insurance and risk allocation, the product liability framework for software, and the evolving standard of care. We hope that these observations are of assistance to the UKJT in finalising the Legal Statement.

Stirling & Rose would be pleased to discuss any aspect of this submission further with the UKJT and its drafting team. We are committed to supporting initiatives that enhance legal certainty for emerging technologies and that position English law as a global framework of choice for AI governance and deployment.

Stirling & Rose

Emerging Technology Law & Global Legal Policy
www.stirlingandrose.com
8 February 2026