



This is an Open Access article distributed under the terms of the Creative Commons Attribution-Non-Commercial-Share Alike 4.0 International (CC-BY-NC-SA 4.0) License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium provided the original work is properly cited.

From Cloud to Culpability: Engineering the Legal ‘Tort Nexus’ for Emergent Multi-Agent AI Systems

Sahil Naeem^a Sameer Naeem^b Dr. Sadiq Hussain^c

^aDepartment of Law, University of Peshawar, Pakistan ^bDepartment of Mechatronics Engineering, University of Engineering and Technology (UET), Peshawar, Pakistan ^cDepartment of Diagnostic Sonology & Biomedical Engineering, AMC, Peshawar, Pakistan

Received 08 June 2026; Accepted 08 July 2026; Published 11 July 2026

Emergent multi-agent artificial intelligence systems challenge the foundations of tort law by dispersing causation across autonomous actors and dissolving traditional notions of fault. Existing doctrines of negligence, product liability, and vicarious responsibility depend on identifiable human or corporate defendants and linear causal chains. Both dissolve in distributed algorithmic environments. This Article identifies the resulting ‘causality black hole’ in tort doctrine and argues that incremental adaptation is insufficient. It proposes instead a new legal construct, the Tort Nexus: a temporary juridical entity instantiated automatically when collective AI behaviour meets predefined technical thresholds of coordination and harm. The Tort Nexus serves as defendant for liability purposes, backed by a pre-funded Swarm Liability Escrow to ensure prompt victim compensation and enable subrogation against negligent human actors. By aligning engineering metrics (communication density, spatial coherence, and goal alignment) with legal activation criteria, the model seeks to convert emergent technical phenomena into legally cognizable events. The Article evaluates this proposal against existing doctrine, identifies the constitutional, economic, and empirical objections it must overcome, and situates it as one possible supplement to, rather than a replacement for, conventional tort categories.

Keywords: artificial intelligence liability, multi-agent systems, tort nexus, legal engineering, emergent behaviour.

INTRODUCTION

In May 2010, what became known as the 'Flash Crash' illustrated the danger of autonomous collective behaviour with sudden clarity. Within minutes, high-frequency algorithmic trading systems triggered a cascade of sell orders that erased nearly one thousand points from the Dow Jones Industrial Average before partial recovery.¹ While human traders had long dominated financial markets, the incident revealed how autonomous algorithmic agents acting collectively can generate harm that eludes traditional fault-based liability models. The harm was real, the cause elusive, and the injured parties faced a legal system designed for human defendants.

Autonomous multi-agent systems, spanning robotic swarms in agriculture, decentralised Internet-of-Things (IoT) networks in smart cities, and algorithmic trading bots in finance, are advancing from promise to deployment. These systems are designed to operate without direct human control, adaptively coordinate, and execute decisions in real time. Their benefits are substantial: enhanced efficiency, improved scale, and novel solutions to pressing problems. Yet their emergent behaviours create a 'causality black hole' for tort law. Liability traditionally depends on identifying a human or corporate defendant, establishing a duty of care, breach, causation, and damage.²

When a swarm of drones devastates a farmer's crop, or a cluster of AI agents triggers a market collapse, the question becomes: who is liable? The individual units execute subroutines. The human operator may have set only high-level objectives. The system's emergent state may fall outside any unit's direct control. Existing legal frameworks strain under such scenarios.

Legal scholars have explored whether conventional tort doctrines can accommodate these developments. Some argue the current regime remains adequate, relying on product liability, negligence, or vicarious liability.³ Others contend that a distinct liability gap arises when AI

¹ *FINDINGS REGARDING THE MARKET EVENTS OF MAY 6, 2010* (US Commodity Futures Trading Commission and Securities and Exchange Commission, 2010)

² Agustina D Saenz et al., 'Autonomous AI systems in the face of liability, regulations and costs' (2023) 6 npj Digital Medicine <<https://www.nature.com/articles/s41746-023-00929-1>> accessed 07 June 2026

³ Joseph Avery, 'Fumble! Anti-Human Bias in the Wake of Socio-Technical System Failures' (2022) 54 Arizona State Law Journal <https://arizonastatelawjournal.org/wp-content/uploads/2022/06/Avery_Publication-compressed.pdf> accessed 07 June 2026

systems act in ways their designers could not foresee.⁴ One study demonstrated that black-box AI systems present particular difficulty for tort doctrine because transparency and causation are jointly undermined.⁵ Globally, the European Union's proposed AI Liability Directive acknowledges the problem through rebuttable presumptions of causation for high-risk AI.⁶ A RAND study found that AI developers face considerable liability exposure under U.S. tort law, yet major structural reforms remain absent.⁷

The challenge is compounded by the technical features of emergent multi-agent systems. These systems communicate via stigmergy or peer-to-peer messaging, evolve internal states, and generate outcomes not traceable to any single command. As one scholar explained, when examining socio-technical systems, proximate cause conclusions depend on perceptions of moral culpability in addition to assessments of causal responsibility. The law thus faces dual pressures: to preserve paths to compensation for victims, and to avoid stifling innovation through overbroad liability.

This Article proposes the ‘Tort Nexus’ to fill that gap. The Tort Nexus is a dynamic, event-specific legal entity that activates when a multi-agent AI system's emergent behaviour meets predefined technical thresholds of collective action and harm. It serves as a defendant in tort actions, dissolves upon dispute resolution, and is backed by a pre-deployment Swarm Liability Escrow. By aligning engineering metrics with legal mechanics, it converts opaque emergent behaviour into a tractable liability construct. Rather than asking courts to impose liability on algorithmic behaviour after harm occurs, the Tort Nexus encodes accountability into the architecture of deployment itself. Legal activation and financial coverage are triggered automatically by the same technical signals that indicate collective AI behaviour.

⁴ Mindy Duffourc and Sara Gerke, ‘Decoding US Tort Liability in Healthcare’s Black-Box AI Era: Lessons from the European Union’ (2024) 27(1) Stanford Technology Law Review <https://law.stanford.edu/wp-content/uploads/2024/02/Publish_27-STLR-1-2024_Decoding-U.S.-Tort-Liability-in-Healthcares-Black-Box-AI-Era.pdf> accessed 07 June 2026

⁵ Daniela Glavaničová and Matteo Pascucci, ‘Vicarious liability: a solution to a problem of AI responsibility?’ (2022) 24 Ethics and Information Technology <<https://link.springer.com/article/10.1007/s10676-022-09657-8>> accessed 07 June 2026

⁶ Iwona Gredka-Ligarska, ‘Employer as an AI System Operator and Tortious Liability for Damage Caused by AI Systems: European and US Perspectives’ (2024) 12 The Chinese Journal of Comparative Law <https://academic.oup.com/cjcl/article/doi/10.1093/cjcl/cxae015/7889035?_cf_chl_f_tk=.H4b3G.8r1Ks2ufWz_yajC9mSLGTo6he4scx4_oNZzE-1783427750-1.0.1.1-TmySciX7u_gXpzmb5g1UIGH9WJz4hnmvgVNEvi0zcZY> accessed 07 June 2026

⁷ Ketan Ramakrishnan et al., *US Tort Liability for Large-Scale Artificial Intelligence Damages: A Primer for Developers and Policymakers* (RAND Corporation, 2024)

This is not merely a procedural fix. It is a reconception of how law and code relate, operationalising a principle this Article terms 'liability by design.'

Methodologically, this Article proceeds doctrinally rather than empirically. It analyses existing tort, agency, and comparative liability doctrine against the technical literature on multi-agent systems and emergence to identify the specific conditions under which conventional categories fail to produce a workable defendant, and it tests a proposed legal construct, the Tort Nexus, against that doctrinal baseline and against anticipated counter-arguments. It is not an empirical study of litigation outcomes, nor a finished legislative product; the model statute in Part III.C is offered as an illustrative drafting exercise to make the doctrinal argument concrete, not as text intended for direct enactment, and its provisions are evaluated critically, including for the constitutional, economic, and evidentiary objections they raise, rather than presented as a settled solution.

Part II reviews the literature on AI personhood, tort law, and emergent system behaviour. Part III develops the legal architecture of the Tort Nexus. Part IV presents the engineering-calibrated Nexus Thresholds. Part V illustrates application scenarios. Part VI discusses implications for law, engineering, policy, economics, and ethics. Part VII concludes.

LITERATURE REVIEW

The Limits of Electronic Personhood –

The proposition of 'electronic personhood' emerged within the European Union as a legal-theoretical response to the challenge of assigning liability when AI or robotic systems cause harm. In 2017, the European Parliament issued a non-legislative Resolution on civil law rules for robotics, suggesting that 'intelligent robots' might receive a special legal status akin to that of a legal person, to secure liability and compensation for resulting damages.⁸ The suggestion provoked intense debate. On one side, advocates argued that conferring legal personhood on entities beyond humans and corporations would enable more direct

⁸ Civil Law Rules on Robotics 2017

assignment of liability.⁹ On the other side, critics cautioned that the concept remains conceptually premature, legally incoherent, and practically ineffective.¹⁰

The Corporate Analogy: Corporations, though not human, are legal persons: they can own property, incur obligations, sue, and be sued. Proponents of AI personhood reason that, if a multi-agent AI system autonomously causes harm, a similar standalone entity could serve as a defendant. Avila Negri argues that the proposed electronic personhood model replicates the corporate model without sufficiently addressing how robots or AI differ conceptually from companies. That analogy faces major hurdles: corporations are formed by humans, are subject to governance and ownership structures, and have legally defined boundaries. AI swarms may lack analogous control structures, share no straightforward owner, and exhibit emergent behaviour. The metaphor of personhood strains under these differences.

Conceptual and Doctrinal Shortcomings: Zapašek and colleagues identify five conditions traditionally required for personhood: technological, economic, legal, moral, and social, which they argue current AI systems do not satisfy. They contend that personhood must presuppose autonomy, moral agency, or sufficient cognitive capacity, none of which current AI reliably exhibits. Additional critiques target the definitional instability of the legal-person category itself: once decoupled from human or corporate origins, the basis for attributing both rights and responsibilities become unstable. Critics have also observed that granting personhood to robots could shield human actors from accountability rather than expose them.

Practical and Policy Limitations: The EU's subsequent regulatory instruments, the Artificial Intelligence Act and the proposed AI Liability Directive, do not adopt a robot-as-person model.¹¹ The Liability Directive briefing acknowledges that the existing liability regime will

⁹ Sergio M C Avila Negri, 'Robot as Legal Person: Electronic Personhood in Robotics and Artificial Intelligence' (2021) 8 *Frontiers in Robotics and AI* <<https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2021.789327/full>> accessed 07 June 2026

¹⁰ Tjaša Zapašek et al., 'APPROPRIATENESS AND FEASIBILITY OF LEGAL PERSONHOOD FOR AI SYSTEMS' (International Conference on Robot Ethics and Standards, Troy, NY, 2018)

¹¹ George Dvorsky, 'AI personhood? European Parliament under fire for proposing recognizing legal status of robots' (*Genetic Literacy Project*, 23 April 2018) <<https://geneticliteracyproject.org/2018/04/23/ai-personhood-european-parliament-under-fire-for-proposing-recognizing-legal-status-of-robots/>> accessed 07 June 2026

be updated, but does not endorse creating autonomous legal entities.¹² International compatibility concerns compound the problem: divergent legal traditions, variations in corporate liability frameworks, and different regulatory appetites render a universal 'robot personhood' model implausible. The result is a liability architecture in flux rather than one anchored in the electronic-person concept.

Fit for Collective, Emergent Systems: The critical question for emergent multi-agent AI systems is whether electronic personhood can capture their distinctive complexity. Multi-agent systems do not operate as monolithic actors. They coordinate, adapt, and evolve in ways that defy clear pre-programmed control. The personhood model requires delineating a distinct entity that occupies legal subject status, but in an emergent scenario, the entity may not be identifiable in advance. As Avila Negri warns, such models may rest on anthropomorphic rhetoric rather than legal utility. Electronic personhood raises important questions about legal subjectivity, yet remains mismatched to collective emergent behaviours because it presumes a unitary actor rather than a dynamic network.

Tort Law and Non-Deterministic Systems –

Tort law balances two aims: compensating victims and allocating losses to those best able to prevent harm. Traditional doctrine, therefore, relies on identifiable actors, a workable causal chain, and standards of reasonable conduct.¹³ Emergent, non-deterministic multi-agent systems upset these foundations. When harm results from collective, adaptive behaviour rather than from a single human decision or a definable product defect, the familiar elements of duty, breach, causation, and scope of liability become difficult to apply.

First, the causation doctrine becomes attenuated when outcomes are produced by a distributed, interactive process. Black-letter law locates liability through factual cause and proximate cause.¹⁴ The Restatement (Third) of Torts articulates these concepts in modern form, focusing courts on whether an actor's conduct was an objective and proximate cause of harm. Those formulations presume a chain of events connecting a human or corporate

¹² 'Briefing: Artificial Intelligence Liability Directive' (European Parliament Research Service, 10 February 2023) <[https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2023\)739342](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2023)739342)> accessed 07 June 2026

¹³ W Page Keeton et al., *PROSSER AND KEETON ON THE LAW OF TORTS* (5th edn, West Publishing Co 1984)

¹⁴ *RESTATEMENT OF TORTS: Liability for Physical and Emotional Harm* (American Law Institute, 2010)

decision to an injurious outcome. In multi-agent systems that learn and coordinate in real time, agents exchange signals, update internal states, and collectively produce behaviours that were not individually programmed by any single actor. The result is a diffused factual causation picture in which no single human or device satisfies the Restatement's causal inquiries.

Second, foreseeability and proximate cause present particular difficulty. Tort law limits liability where harm is unforeseeable or where policy concerns counsel against expansive recovery. Emergent systems produce outcomes that by design exceed designers' expectations. Where swarm coordination or emergent goal-alignment yields a novel hazardous state, classifying the harm as foreseeable in the traditional sense is difficult. Scholars have emphasised that the obscurity in algorithm-driven decisions enlarges the practical burden of proof on victims.¹⁵

Third, doctrinal regimes that might otherwise capture distributed wrongs have limits when applied to non-deterministic AI. Product liability treats defective products as a basis for recovery, but that doctrine presumes identifiable defects, design flaws, or manufacturing failures that proximately cause harm.¹⁶ Where emergent behaviour arises from lawful design choices interacting with complex environments, labelling the system defective is contestable. Strict liability regimes often demand a finding that the activity was unusually dangerous. Multi-agent systems may be dangerous in aggregate while each agent remains within its certified operational envelope. Negligence doctrine depends on a breach of a standard of care, but what constitutes reasonable care in developing adaptive, learning agents remains unsettled.

Vicarious liability and enterprise liability hold organisations accountable for acts of their agents or for risks originating in their operations. Those doctrines succeed when a human actor or organisational policy can be tied to the injurious conduct. They struggle when faults arise through self-organising processes without a clear managerial decision to anchor

¹⁵ Sandra Wachter et al., 'Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation' (2017) 7(2) *International Data Privacy Law* <<https://academic.oup.com/idpl/article-abstract/7/2/76/3860948?redirectedFrom=fulltext>> accessed 07 June 2026

¹⁶ *Restatement (Second) of Torts* (American Law Institute 1965)

liability. As Bryson has noted, attributing responsibility to nonhuman systems without tracing regulatory or corporate governance failures risks both under-inclusion and over-inclusion of defendants.¹⁷

Fourth, evidentiary opacity compounds doctrinal gaps. Pasquale and others have argued that secrecy around algorithmic processes undermines the transparency courts need to perform fact-finding, deepening plaintiffs' difficulties in establishing causation and breach.¹⁸ The European Commission's proposed AI Liability Directive would introduce rebuttable presumptions to ease plaintiffs' burdens in appropriate cases, acknowledging the evidentiary problem while leaving open the harder question of how to identify which harms justify such presumptions in multi-agent emergent contexts.

These doctrinal gaps are not merely theoretical. Litigation arising from Tesla's Autopilot and Full Self-Driving systems has repeatedly confronted courts and juries with the difficulty of allocating fault between driver conduct, software design, and a learning system's own emergent decision-making in real time, with outcomes turning heavily on case-specific evidence of design defect rather than on any settled doctrinal test for algorithmic causation.¹⁹ Outside the autonomous-vehicle context, a Canadian tribunal held an airline liable for inaccurate information generated by its customer-service chatbot, reasoning that the airline could not distance itself from a tool it had deployed to act on its behalf, an outcome reached through ordinary principles of vicarious responsibility rather than any AI-specific doctrine and one that illustrates how far conventional categories can stretch before they are asked to address genuinely emergent, multi-agent behavior of the kind this Article addresses.²⁰

Finally, normative and policy tradeoffs are acute. Tort doctrine distributes losses to those who can internalise costs and deter harm, but multi-agent systems are deployed by economically powerful actors who rely on complex software supply chains, third-party components, and open data. Pinning liability solely on one node in that chain may produce inefficient deterrence. As Calo explains in his robotics work, tailored, context-specific

¹⁷ Joanna J Bryson, 'Robots Should Be Slaves' in Yorick Wilks (ed), *Close Engagements with Artificial Companions: Key social, psychological, ethical and design issues* (John Benjamins Publishing 2010)

¹⁸ Frank Pasquale, *THE BLACK BOX SOCIETY: The Secret Algorithms That Control Money and Information* (Harvard University Press 2015)

¹⁹ *Benavides v Tesla Inc* [2025] No. 1:2021cv21940

²⁰ *Moffatt v Air Canada* [2024] BCCRT 149 (Civ Resol Trib BC); *Mata v Avianca, Inc* [2023] 678 F Supp 3d 443

interventions are preferable to monolithic legal transformations because they better balance safety and innovation.²¹

None of the existing doctrines, whether product liability, negligence, strict liability, or enterprise responsibility, cleanly resolves the accountability problem posed by emergent, non-deterministic multi-agent systems. A further candidate merits direct engagement: market-share liability, developed in mass-tort contexts where plaintiffs cannot identify which of several fungible-product manufacturers caused their specific harm, apportions liability among defendants according to their share of the relevant market.²² That doctrine offers a closer analogy to emergent multi-agent harm than the others surveyed above, since it too dispenses with individualised proof of causation. It nonetheless presupposes a defined, identifiable pool of potential defendants whose relative market shares can be measured, a precondition that often fails for ad hoc multi-agent collectives drawing on shared open-source components, third-party data, and cloud infrastructure with no stable market boundary. Where that precondition can be satisfied, for instance, among a small number of identifiable platform operators in a concentrated industry, market-share apportionment may be a more conservative and judicially familiar alternative to the Tort Nexus, and courts confronting such facts have reason to prefer it. The Tort Nexus is offered for the residual category in which even an approximate market of defendants cannot be constructed because participation in the harm-causing collective is itself emergent and unstable rather than fixed in advance. The Tort Nexus proposal accordingly does not supplant product liability, negligence, strict liability, enterprise responsibility, or market-share liability; it proposes a dynamic, event-specific legal construct aligned to technical thresholds for the narrower set of cases those doctrines leave uncovered, without resort to full AI personhood.

Technical Foundations of Emergent Behaviour –

The liability problems surrounding multi-agent AI arise from the technical phenomenon of emergence: the spontaneous appearance of system-level properties or behaviours not

²¹ Ryan Calo, 'Robotics and the Lessons of Cyberlaw' (2015) 103(3) California Law Review <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2402972&_cf_chl_f_tk=PAw4nHv.vz1w53oR7cJGj_hNmYw1MSfklzX91dHlrzM-1783429974-1.0.1.1-7qJlvLMX_qXJnGg4B87OgrfOQLImX1KME1V.w.s.YAw> accessed 07 June 2026

²² *Sindell v Abbott Laboratories* [1980] 607 P 2d 924 (Cal); David A Fischer 'Causation in Fact in Omission Cases' (1992) (4) Utah Law Review <<https://dc.law.utah.edu/ulr/vol1992/iss4/3/>> accessed 07 June 2026

explicitly programmed into any individual agent. In robotics, computer science, and complex systems theory, emergence describes situations in which local interactions among autonomous components yield global patterns that are both nonlinear and unpredictable in detail.²³ This is not a marginal property but a foundational principle of distributed AI systems, from robotic swarms to algorithmic trading infrastructures.

In swarm robotics, individual robots operate under simple rule sets, yet collectively exhibit sophisticated behaviours such as flocking, pattern formation, or cooperative foraging. The same principle applies to software agents in digital environments. Financial algorithms, recommender systems, and IoT networks display feedback loops that amplify initial inputs into large-scale, unplanned outcomes.²⁴

Technically, emergent behaviour resists deterministic traceability. Even when each unit's code is transparent, the resulting system state often cannot be linearly mapped from individual commands.²⁵ Scholars in AI safety call this the interpretability gap.²⁶ For law, that gap translates into a causality vacuum: harm arises, but no discrete actor or malfunction can be isolated. The system functions as a whole that exceeds the sum of its coded parts.

The theory of stigmergy, first articulated in biological studies of termite behaviour, explains how agents communicate indirectly by modifying shared environments.²⁷ Modern multi-agent architectures replicate this mechanism through shared data layers or common digital environments. This indirect coordination increases robustness but simultaneously undermines traceability. When the liability doctrine assumes discrete chains of causation, stigmergic coordination collapses the model: actions are distributed across time and agents, creating an evidentiary fog incompatible with conventional tort elements of duty and breach.

²³ Melanie Mitchell, *COMPLEXITY: A GUIDED TOUR* (1st edn, OUP 2009)

²⁴ Donald MacKenzie, *An Engine, Not a Camera: How Financial Models Shape Markets* (The MIT Press 2008)

²⁵ Michael Wooldridge, *An Introduction to Multiagent Systems* (2nd edn, Wiley 1905)

²⁶ Finale Doshi-Velez and Been Kim, 'Toward a Rigorous Science of Interpretable Machine Learning' (2017) arXiv: Machine Learning <<https://www.semanticscholar.org/paper/Towards-A-Rigorous-Science-of-Interpretable-Machine-Doshi-Velez-Kim/5c39e37022661f81f79e481240ed9b175dec6513>> accessed 05 June 2026

²⁷ Pierre-Paul Grassé, 'La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs' (1959) 6 *Insectes Sociaux* <<https://doi.org/10.1007/BF02223791>> accessed 05 June 2026

In complex adaptive systems, small perturbations can yield disproportionate results through positive feedback and recursive adaptation.²⁸ The European Commission's High-Level Expert Group on Artificial Intelligence recognises that emergent properties may render accountability diffuse and temporally extended.²⁹ From a legal standpoint, such technical realities demand an attribution model that treats collectives as functional entities rather than as a mere aggregation of individual agents.

Engineers seek to mitigate unpredictability through verification frameworks, sandbox testing, and runtime monitoring.³⁰ Yet, as incidents in autonomous drone coordination and algorithmic trading demonstrate, safeguards cannot always prevent convergence toward harmful emergent states once systems scale beyond certain complexity thresholds. These empirical findings underscore that legal theory must engage directly with system-level dynamics. A tort framework grounded only in component-level fault cannot capture emergent harm generated by adaptive collectives.

Any credible liability architecture for AI must therefore be co-designed with an engineering understanding of emergence. The Tort Nexus builds on this insight by defining legal activation criteria aligned with measurable system metrics: communication density, goal coherence, and deviation from certified operational envelopes. Importing quantifiable thresholds into legal doctrine, it transforms technical indeterminacy into a manageable evidentiary standard.

Ethical and Economic Dimensions –

The governance of artificial intelligence is not a purely technical or legal undertaking; it is an exercise in ethical risk distribution and economic design. As multi-agent systems operate across borders and markets, questions of accountability, equity, and cost internalisation become central. Ethical frameworks ranging from deontological models of responsibility to

²⁸ John H Holland, *HIDDEN ORDER: How Adaptation Builds Complexity* (Basic Books 1995)

²⁹ *Ethics guidelines for trustworthy AI* (European Commission, 2019)

³⁰ *AI RISK MANAGEMENT FRAMEWORK* (National Institute of Standards and Technology, 2023)

utilitarian cost-benefit analysis have struggled to capture the distributed nature of emergent AI harms.³¹

Ethical challenges arise first from the diffusion of agency. When decision-making is distributed among autonomous agents, identifying moral responsibility becomes as difficult as tracing legal causation. Floridi and Cowls note that AI ethics must move beyond individual responsibility toward systemic accountability, recognising that moral agency can be functionally collective.³² The OECD Principles on Artificial Intelligence similarly emphasise inclusive growth, sustainable development, and well-being as regulatory objectives, framing AI not merely as a technological system but as a moral economy.³³

Economic scholarship reinforces this ethical imperative. When emergent AI behaviour produces harm, such as flash crashes, algorithmic bias, or large-scale drone coordination failures, the costs are often socialised while profits remain privatised. Tort law, in its classical formulation, seeks to internalise externalities by aligning risk-bearing with control.³⁴ Yet, as Allen observes in the context of algorithmic trading, control in such environments is both fragmented and opaque, frustrating that alignment.³⁵ The result is a moral hazard: developers and operators deploy systems whose complex behaviours they cannot fully predict, while victims face barriers to redress.

Economic modelling suggests that ex ante liability instruments, such as mandatory insurance or escrow mechanisms, can restore efficiency by pricing systemic risk before harm occurs.³⁶ Such approaches mirror environmental liability regimes, where uncertain but foreseeable harms justify collective insurance pools.³⁷ The Tort Nexus incorporates this insight through

³¹ Brent Daniel Mittelstadt et al., 'The ethics of algorithms: Mapping the debate' (2016) 3(2) *Big Data and Society* <<https://journals.sagepub.com/doi/epub/10.1177/2053951716679679>> accessed 07 June 2026

³² Luciano Floridi and Josh Cowls, 'A Unified Framework of Five Principles for AI in Society' (2019) 1(1) *Harvard Data Science Review* <<https://doi.org/10.1162/99608f92.8cd550d1>> accessed 07 June 2026

³³ *Recommendation of the Council on Artificial Intelligence* (Organisation for Economic Co-operation and Development 2019)

³⁴ Guido Calabresi, *THE COSTS OF ACCIDENTS: A LEGAL AND ECONOMIC ANALYSIS* (Yale University Press 1970)

³⁵ Hilary J Allen, *DRIVERLESS FINANCE: Fintech's Impact on Financial Stability* (Oxford University Press on Demand 2022)

³⁶ S Shavell, 'The judgment proof problem' (1986) 6(1) *International Review of Law and Economics* <[https://doi.org/10.1016/0144-8188\(86\)90038-4](https://doi.org/10.1016/0144-8188(86)90038-4)> accessed 07 June 2026

³⁷ James E Krier and Stewart J Schwab, 'Property Rules and Liability Rules: The Cathedral in Another Light' (1995) 70 *New York University Law Review* <<https://scholarship.law.cornell.edu/facpub/1439/>> accessed 07 June 2026

its Swarm Liability Escrow, creating a pre-funded mechanism to ensure compensation and incentivise risk-conscious design. By embedding liability in the architecture of deployment, it transforms ethical aspirations, fairness, deterrence, and distributive justice into enforceable economic mechanisms.

A further ethical concern arises from global inequity. Advanced economies are developing sophisticated regulatory toolkits, such as the EU's AI Act and the U.S. NIST AI Risk Management Framework, while many developing jurisdictions lack institutional capacity for AI oversight.³⁸ Without harmonised standards, multinational AI systems risk creating a jurisdictional divide in which harms occur in legally underprotected regions. The UNESCO Recommendation on the Ethics of Artificial Intelligence calls for solidarity-based governance to prevent precisely this outcome.³⁹ A global framework, such as a Tort Nexus Convention, would align with these principles by ensuring that accountability mechanisms scale beyond national boundaries.

The ethical legitimacy of any liability system depends on transparency and oversight. Independent audits, privacy protection, and algorithmic accountability remain prerequisites for public trust.⁴⁰ The Tort Nexus therefore integrates ethical safeguards within its legal architecture, not as ornamental principles but as procedural duties. In doing so, it operationalises a vision of responsible innovation that aligns moral philosophy with market incentives.

THE LEGAL ARCHITECTURE OF THE 'TORT NEXUS'

Definition and Conceptual Foundations: At the heart of modern tort law lies the problem of attribution: the question of who should bear the cost of harm when control, intention, and causation fragment across actors and technologies. Classical doctrines, from respondeat superior to enterprise liability, evolved to ensure that victims are not left uncompensated when harm arises from collective or institutional activity. Yet, as autonomous multi-agent

³⁸ Artificial Intelligence Act 2021

³⁹ *Recommendation on the Ethics of Artificial Intelligence* (UNESCO, 2021)

⁴⁰ Lilian Edwards and Michael Veale, 'SLAVE TO THE ALGORITHM? WHY A 'RIGHT TO AN EXPLANATION' IS PROBABLY NOT THE REMEDY YOU ARE LOOKING FOR' (2017) 16(1) *Duke L & Tech Rev* <<https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1315&context=dltr>> accessed 07 June 2026

systems increasingly operate through distributed intelligence, these doctrines strain under conditions they were never designed to address.

Vicarious liability presumes an identifiable principal capable of supervision and control. In systems composed of hundreds or thousands of semi-autonomous agents making probabilistic decisions through machine learning algorithms, such control is neither factual nor legally coherent. The master-servant relationship that underpins respondeat superior dissolves when decision-making is algorithmically diffused. Enterprise liability, which attributes harm to the entity that benefits from risk-producing activities, assumes a stable legal personality behind the enterprise.⁴¹ Emergent AI collectives, dynamically formed across cloud infrastructures and jurisdictions, lack any continuous juridical identity capable of absorbing liability.

Traditional strict liability and product liability frameworks also falter. These doctrines succeed when harm arises from defective design or inherently dangerous activity. Emergent AI behaviour often results not from design defects but from unpredictable interactions among properly functioning components. Courts applying conventional negligence or defect tests face an evidentiary vacuum: no human intent, no defective part, and no traceable act of omission.

This impasse reflects what scholars describe as the causality black hole of emergent systems: harm is real, but the legal system cannot assign fault within its existing conceptual categories, a condition Burri terms the responsibility gap.⁴² Legal theory has confronted similar dilemmas before. The expansion of corporate and enterprise liability in the early twentieth century responded to the analogous diffusion of agency in industrial production.⁴³ The Tort Nexus extends that doctrinal evolution into the algorithmic age, establishing a dynamic, event-specific legal personhood that materialises only to bear liability for collective AI harm.

⁴¹ Gregory C Keating, 'The Idea of Fairness in the Law of Enterprise Liability' (1997) 95(5) Michigan Law Review <<https://repository.law.umich.edu/mlr/vol95/iss5/3/>> accessed 07 June 2026

⁴² Thomas Burri, 'Machine Learning and the Law: 5 Theses' (2017) Law, Computer Science <<https://www.semanticscholar.org/paper/Machine-Learning-and-the-Law%3A-5-Theses-Burri/a44ad2c1ce27cb147742add5b1db430fc7450ce9>> accessed 07 June 2026

⁴³ Morton J Horwitz, *The Transformation of American Law 1870-1960: The Crisis of Legal Orthodoxy* (Oxford Univ Pr on Demand 1994)

Definition –

The Tort Nexus is a temporary legal entity, instantiated automatically upon satisfaction of predefined Nexus Thresholds: objective technical metrics indicating that an autonomous system has acted as a cohesive, emergent unit causing harm. Once activated, the Tort Nexus serves solely as a defendant for tort purposes, representing the collective agency of the system at the moment of harm. It dissolves upon the final resolution of liability and distribution of damages. Conceptually, the Tort Nexus does not confer rights; it allocates responsibility. It functions analogously to a corporate entity or bankruptcy estate, but without continuous personality or economic autonomy. By divorcing liability from perpetual personhood, it addresses the policy tension that undermined prior proposals for electronic personhood.⁴⁴ It is a mechanism of accountability, not metaphysics.

This definition satisfies the foundational objectives of tort law: compensation, deterrence, and corrective justice.⁴⁵ Victims gain a defendant with legal standing and access to a Swarm Liability Escrow pre-funded at deployment. System operators retain the ability to recover subrogation rights against negligent developers or maintainers. Regulators gain a measurable trigger for liability, grounded in engineering data rather than moral speculation. The Tort Nexus converts the causality black hole into an actionable legal event, harmonising technical thresholds with juridical principles.

Key Legal Characteristics –

The Tort Nexus framework derives its coherence from four defining legal characteristics: functional limitation, dynamic incorporation and dissolution, Swarm Liability Escrow, and ethical safeguards.

Functional Limitation: The Tort Nexus exists solely for liability adjudication. It has no contractual capacity, political rights, or proprietary existence beyond the scope of the harm event. This principle parallels the treatment of *in rem* defendants in maritime law or bankruptcy estates that exist only to resolve claims.⁴⁶ By limiting legal capacity to a single

⁴⁴ Civil Law Rules on Robotics 2017

⁴⁵ Jules L Coleman, *RISKS AND WRONGS* (Cambridge University Press 1992)

⁴⁶ *United States v Carroll Towing Co* [1947] 159 F 2d 169, 173 (2d Cir)

procedural function, the Tort Nexus preserves accountability without anthropomorphising AI. This functional minimalism distinguishes it from proposals for electronic personhood, which critics have rejected as both conceptually flawed and ethically hazardous.⁴⁷

Dynamic Incorporation and Dissolution: Unlike corporations or trusts, the Tort Nexus is not registered through human volition but instantiated automatically when a multi-agent system meets predefined technical thresholds signalling collective emergent behaviour that results in harm. Upon judicial confirmation or administrative recognition, the entity is incorporated solely for the duration of dispute resolution. After judgment or settlement, it dissolves by operation of law. This design reflects both efficiency and moral parsimony: liability attaches precisely to the moment of collective action and dissipates when accountability has been achieved, aligning with Balkin's observation that the law of robotics must evolve to attribute responsibility to emergent systems without displacing human accountability.⁴⁸ Dynamic dissolution prevents the perpetuation of dormant or speculative legal entities and mitigates concerns of AI personhood creep, a danger highlighted by Bryson, Diamantis, and Grant. It ensures that the law responds proportionately to technical events, not abstractions of agency.

Swarm Liability Escrow: To guarantee victim compensation and incentivise risk-aware deployment, every qualifying multi-agent system must maintain a Swarm Liability Escrow: a pre-deployment fund or insurance bond administered by licensed intermediaries. The escrow operates as a strict-liability reserve. Upon activation of a Tort Nexus, victims may claim damages directly from the fund without proving negligence. This mechanism reflects Posner's economic rationale that liability should internalise the social cost of risk, ensuring that developers bear the financial consequences of potential harm and thereby encouraging efficient precaution.⁴⁹

This structure draws on analogues in environmental and nuclear liability regimes, where uncertain but potentially catastrophic harms justify advance pooling of resources.⁵⁰ It also

⁴⁷ Joanna J Bryson et al., 'Of, for, and by the people: the legal lacuna of synthetic persons' (2017) 25 Artificial Intelligence and Law <<https://link.springer.com/article/10.1007/s10506-017-9214-9>> 07 July 2026

⁴⁸ Jack M Balkin, 'The Path of Robotics Law' (2015) 6 Calif L Rev Circuit <<https://euro.ecom.cmu.edu/program/law/08-732/AI/Balkin.pdf>> accessed 07 June 2026

⁴⁹ Richard A Posner, *ECONOMIC ANALYSIS OF LAW* (7th edn, Aspen Publishers Inc US 2007)

⁵⁰ Convention on Supplementary Compensation for Nuclear Damage 1997

aligns with the OECD Recommendation on Artificial Intelligence, which encourages ex ante risk mitigation through economic instruments. Following payout, subrogation rights allow insurers or regulators to pursue negligent designers, operators, or manufacturers, thereby preserving deterrence while maintaining rapid victim relief.

Ethical and Procedural Safeguards: Given the data-intensive nature of multi-agent systems, ethical safeguards are intrinsic to the Tort Nexus architecture. Each Nexus proceeding must comply with privacy norms derived from the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). Independent Nexus Auditors, certified under administrative guidelines, supervise data access and logging integrity to prevent bias, manipulation, or selective disclosure.⁵¹

Procedural fairness is further enhanced through judicial oversight of activation criteria and transparent publication of Nexus-related decisions.⁵² These measures reflect principles of algorithmic accountability advanced in international ethics frameworks, ensuring that liability mechanisms serve justice without undermining individual rights. Together, these characteristics define a legal architecture that is precise, proportional, and technologically grounded. The Tort Nexus does not replace human accountability but integrates it into a structured, event-specific interface between law and machine behaviour.⁵³

Proposed Statutory Language: The Tort Nexus Act –

The following draft statute, tentatively titled the Tort Nexus Act (TNA), proposes a scalable legislative framework designed for integration into both domestic and supranational legal systems. It codifies the conceptual foundations established in the preceding sections, ensuring that emergent, distributed, or self-organising systems remain legally accountable without extending full legal personhood to artificial agents. The text follows the formal style of model legislation used in the OECD, EU Parliament Draft Resolutions, and U.S. Uniform Law Commission frameworks.

⁵¹ General Data Protection Regulation 2016; California Consumer Privacy Act 2018

⁵² Artificial Intelligence Act 2021

⁵³ Virginia Dignum, *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way* (Springer-Nature New York Inc 2019); Brent Daniel Mittelstadt, 'Principles alone cannot guarantee ethical AI' (2019) 1 Nature Machine Intelligence <<https://www.nature.com/articles/s42256-019-0114-4>> accessed 07 June 2026

Section 1. Short Title and Purpose

- (a) This Act may be cited as the Tort Nexus Act of 2026.
- (b) The purpose of this Act is to establish a temporary legal framework, known as a Tort Nexus, for the attribution of civil liability arising from autonomous, distributed, or emergent artificial intelligence systems.
- (c) The Act aims to:
 - i. Ensure compensatory justice for victims of AI-related harms;
 - ii. Preserve incentives for responsible innovation; and
 - iii. Harmonise liability allocation across jurisdictions consistent with the OECD Recommendation on Artificial Intelligence (2019) and UNESCO Recommendation on the Ethics of Artificial Intelligence (2021).

Section 2. Definitions

For the purposes of this Act:

- (a) 'Multi-agent system' means a digital or robotic environment consisting of two or more semi-autonomous agents capable of interactive, adaptive, or self-organising behaviour.
- (b) 'Emergent harm event' means an occurrence in which collective agent behaviour produces damage or loss not directly traceable to a single human or algorithmic source.
- (c) 'Tort Nexus' means a temporary juridical entity automatically instantiated under Section 3 when an emergent harm event is verified by a competent authority.

Section 3. Creation and Dissolution of the Tort Nexus

- (a) A Tort Nexus shall be deemed incorporated when:
 - i. An emergent harm event occurs; and

- ii. Technical verification by a licensed Nexus Auditor confirms that the causal chain meets the definitional threshold in Section 2(b).

(b) The Nexus shall be registered by the competent regulatory authority and assigned a unique identifier for judicial and administrative purposes.

(c) Upon final judgment, settlement, or termination of proceedings, the Nexus shall automatically dissolve. No continuing rights, liabilities, or assets shall survive beyond that event.

Section 4. Liability and Compensation

(a) The Tort Nexus shall be subject to strict liability for harms arising from emergent events verified under Section 3.

(b) Compensation shall be paid from a Swarm Liability Escrow fund maintained pursuant to Section 5.

(c) Following compensation, insurers or state authorities may exercise rights of subrogation against manufacturers, operators, or data controllers whose negligence or design defect contributed to the harm.

Section 5. Swarm Liability Escrow

(a) Every entity deploying a multi-agent system capable of emergent behaviour shall establish a Swarm Liability Escrow sufficient to cover foreseeable damages.

(b) Escrow contributions shall be proportionate to system scale, risk classification, and data exposure level, as determined by regulatory guidance.

(c) Funds shall be held by licensed intermediaries and audited annually for adequacy and compliance.

Section 6. Oversight and Safeguards

- (a) Judicial oversight shall ensure procedural fairness, transparency of activation criteria, and data protection.
- (b) Nexus Auditors shall be bound by confidentiality and subject to professional liability.
- (c) All proceedings involving Tort Nexus activation shall comply with applicable data protection laws, including the GDPR and CCPA.⁵⁴

Section 7. International Recognition and Harmonization

- (a) A Tort Nexus registered in one jurisdiction shall be recognised for enforcement in any other jurisdiction adopting reciprocal provisions.
- (b) States may enter into treaties or memoranda of understanding to harmonize Nexus standards consistent with international law.⁵⁵

Legislative Notes (Doctrinal Commentary)

Section 3 operationalises functional minimalism, ensuring legal activation only for specific harm events and preventing open-ended AI personhood. Its principal weakness is timing: by conditioning incorporation on a Nexus Auditor's post hoc verification, the section leaves victims without a defendant during the interval between harm and certification, a delay that could be substantial where the underlying technical determination is contested. A future draft might consider provisional registration on a lower evidentiary standard, subject to dissolution if verification later fails, to address this gap.

Section 4 embodies strict liability with dynamic apportionment, paralleling environmental and product liability regimes. Strict liability lowers the evidentiary burden on victims but does so by shifting risk onto the escrow contributors as a class; Section 4 does not specify how damages are apportioned among multiple operators whose systems jointly contributed

⁵⁴ General Data Protection Regulation 2016; California Consumer Privacy Act 2018

⁵⁵ *Model Law on Electronic Commerce with Guide to Enactment* 1996 (UNCITRAL, 1996)

to a single emergent event, an omission that would need resolution, plausibly by analogy to joint and several liability with a contribution right, before the section could function in a genuinely multi-operator deployment.

Section 5 draws on ex ante risk financing models analogous to nuclear liability and oil spill compensation systems. Those analogues involved a small number of well-capitalized operators in a defined industry; Section 5(a)'s undifferentiated escrow requirement for 'every entity deploying a multi-agent system capable of emergent behavior' risks imposing comparable fixed costs on startups and academic researchers operating at a far smaller scale, a concern Part VI.B addresses but does not fully resolve, and a risk-tiered contribution schedule keyed to deployment scale would likely be necessary for the section to be administrable.

Section 6 safeguards due process and data privacy in line with global AI ethics standards, but, as discussed in Part VI.B's treatment of due process and standing, the section's reference to judicial oversight is aspirational rather than operational: it does not specify when judicial review attaches, who may invoke it, or whether activation may be contested before escrow disbursement. Specifying an interlocutory review mechanism is a precondition for the section to function as the procedural safeguard it is described as providing.

Global Applicability and Harmonisation –

The Case for Transnational Uniformity: Emergent AI behaviour is not territorially bounded. From algorithmic trading cascades that cross exchanges within milliseconds to drone swarms deployed across agricultural regions spanning borders, liability cannot be confined within national silos. This necessitates a transnational tort architecture analogous to those developed in environmental, aviation, and maritime law.⁵⁶ The Tort Nexus offers a neutral legal scaffold that integrates with diverse domestic doctrines while maintaining global interoperability. Like the Basel Convention and the Montreal Convention in their respective

⁵⁶ Christopher Hodges, *Law and Corporate Behaviour: Integrating Theories of Regulation, Enforcement, Compliance and Ethics* (Bloomsbury Publishing 2015)

domains, it provides an adaptable legal language for attributing harm in distributed systems without mandating uniform substantive law.⁵⁷

Compatibility Across Legal Traditions –

Common Law Systems: In common law jurisdictions, the Tort Nexus can be embedded through statutory delegation. Courts in the United Kingdom, the United States, Canada, and India have historically accommodated new liability categories through judicial innovation, ranging from *Rylands v Fletcher*⁵⁸ to product liability and environmental torts. This evolutionary receptiveness allows legislatures to treat the Tort Nexus as a special statutory entity akin to a limited-purpose corporation or trust. Under the U.S. framework, this aligns with Congress's power to create federally recognised juridical entities for specific functions, such as the federal vaccine injury compensation program.⁵⁹

Civil Law Systems: Civil law jurisdictions, grounded in codified liability regimes, may adopt the Tort Nexus through functional transposition. The Code civil (France) and Bürgerliches Gesetzbuch (Germany) already provide strict liability templates for activities involving heightened risk.⁶⁰ By defining the Nexus as a temporary juridical fiction, civil codes can integrate it as a sui generis form of collective liability without personhood, consistent with Article 1242 of the French Code civil and its doctrine of liability for things within one's custody.⁶¹

Hybrid and Developing Jurisdictions: In hybrid legal systems such as those of South Africa and Singapore, doctrines of fault-based and strict liability coexist, reflecting the interplay of common law and statutory traditions.⁶² The Tort Nexus can be operationalised through subordinate legislation or regulatory circulars, employing risk-classification tiers to accommodate local technological capacity and insurance infrastructure.⁶³ Such an approach

⁵⁷ Convention for the Unification of Certain Rules for International Carriage by Air 1999

⁵⁸ *Rylands v Fletcher* [1868] LR 3 HL 330

⁵⁹ National Childhood Vaccine Injury Act 1986

⁶⁰ Code Civil 1889, art 1242

⁶¹ Jean-Sébastien Borghetti, 'The Culture of Tort Law in France' (2012) 3 Journal of European Tort Law <<https://doi.org/10.1515/jetl-2012-0158>> accessed 07 June 2026

⁶² Mauro Bussani and Anthony J Sebok, *Comparative Tort Law: Global Perspectives* (2nd edn, Edward Elgar 2021)

⁶³ Huijuan Peng and Pey-Woan Lee, 'Reimagining U.S. Tort Law for Deepfake Harms: Comparative Insights from China and Singapore' (2025) 18(2) Journal of Tort Law <<https://doi.org/10.1515/jtl-2025-0028>> accessed 07 June 2026

facilitates equitable adoption, enabling emerging economies to integrate liability frameworks without incurring the full compliance burdens often associated with Western regulatory models.

The Indian Position: India presents a distinct test case because its tort jurisprudence already contains a strict, even absolute, liability tradition that long predates the AI liability debate. In *M.C. Mehta v Union of India*, the Supreme Court departed from the English rule in *Rylands v Fletcher* and held enterprises engaged in hazardous activity absolutely liable for resulting harm, without the defences of statutory authority or third-party intervention that the English doctrine permits.⁶⁴ That doctrine was developed for fixed industrial hazards with an identifiable enterprise behind them, and its fit with distributed, multi-stakeholder AI deployments is not automatic: an emergent multi-agent system may have no single enterprise whose hazardous activity can be isolated in the way the absolute liability rule presupposes. The Tort Nexus, by locating liability in a temporary entity tied to a verified emergent event rather than to a pre-existing enterprise, is in one sense a narrower construct than Indian absolute liability and in another sense an attempt to solve the precise problem, that of identifying the relevant enterprise, that absolute liability assumes is already solved.

India's constitutional structure raises a further question that the Tort Nexus Act, as drafted in Part III.C, does not resolve. The creation of a *sui generis* juridical entity by executive or regulatory action, rather than by direct legislative enactment, would need to satisfy Article 14's guarantee against arbitrary state action and would likely require Parliament, rather than subordinate rule-making bodies, to define the entity's incidents of liability, consistent with the constitutional limits on delegated legislation that Indian courts have policed in other regulatory contexts.⁶⁵ India's absence from the Rome Statute system illustrates a related, recurring concern in Indian engagement with internationally proposed accountability architectures: a preference for retaining domestic control over the definition and attribution of liability rather than ceding it to externally negotiated standards. A Global Tort Nexus Convention of the kind sketched in Part III.D.3 would need to accommodate that preference explicitly, for instance, by allowing domestic courts rather than an international auditing

⁶⁴ *M C Mehta & Anr v Union of India & Ors* AIR 1987 SC 1086; *Union Carbide Corporation Etc Etc v Union of India Etc Etc* AIR 1992 SC 248

⁶⁵ Constitution of India 1950, art 14; *In re Delhi Laws Act 1912* (1951) 2 SCR 747

body to make the final determination on Nexus activation, if it is to be a realistic candidate for adoption in the Indian context rather than a model drawn primarily from European and American precedent.

International Coordination Mechanisms –

The Global Tort Nexus Convention (GTNC): A model Global Tort Nexus Convention could be established under the auspices of UNCITRAL or the OECD, harmonising procedural standards and recognition clauses across borders. The Convention would: (1) define minimum standards for Nexus activation and audit; (2) mandate mutual recognition of Nexus entities; and (3) provide mechanisms for transnational escrow fund enforcement. This mirrors the Vienna Convention on Civil Liability for Nuclear Damage and the Oil Pollution Compensation Funds structure, both of which balance global risk pooling with national implementation.⁶⁶

Integration with Existing AI Governance Frameworks: The proposed system aligns with the OECD AI Principles (2019), emphasising accountability and robustness; the EU Artificial Intelligence Act (2024), establishing risk tiers; the UNESCO AI Ethics Recommendation (2021), calling for liability clarity; and the Council of Europe's Framework Convention on AI (2024), which envisions harmonised AI accountability regimes. Together, these instruments demonstrate a converging consensus around accountability by design, a principle the Tort Nexus operationalises through legal engineering.

Economic and Ethical Harmonisation –

Global harmonisation also ensures equitable access to justice. Developed economies may adopt high escrow thresholds to deter undercapitalised deployment, whereas developing states can utilise pooled regional escrow funds or state-backed insurance facilities.⁶⁷ This prevents what scholars term 'AI liability dumping': the offshoring of high-risk AI experiments to jurisdictions with weak tort systems. By creating a portable and auditable

⁶⁶ International Convention on the Establishment of an International Fund for Compensation for Oil Pollution Damage 1971

⁶⁷ Michael J Trebilcock, *Dealing with Losers: The Political Economy of Policy Transitions* (OUP USA 2014)

liability identity, the Tort Nexus safeguards against both regulatory arbitrage and moral externalisation.

ENGINEERING THE NEXUS THRESHOLDS

The Technical Rationale for Legal Activation Metrics: Legal accountability for emergent systems requires a technical substrate capable of evidentiary precision. Traditional tort adjudication depends on traceable causation, yet emergent multi-agent behaviour often produces harms where no single node or operator can be identified as the proximate cause. The Nexus Thresholds function as the technical tripwires that determine when a distributed system's behaviour transcends individual control, warranting activation of a Tort Nexus. These thresholds transform a metaphysical inquiry, 'who caused the harm?', into an engineered determination based on data, system architecture, and operational envelopes.⁶⁸ This is consistent with the liability-by-design doctrine proposed in the preceding sections: embedding legally meaningful markers into system design to ensure post-incident accountability.

Emergent Behaviour Criterion –

The first threshold concerns emergent behaviour detection: the system's capacity to generate actions not directly commanded or predictable from individual agent programming. A Nexus Auditor, typically a certified engineer or digital forensics expert, validates emergent causation using three criteria:

Causation Decoupling: No one-to-one mapping between operator input and system output can be established from event logs.

Operational Envelope Breach: The collective state deviates beyond certified operational parameters defined in the system's safety documentation.

⁶⁸ Lawrence Lessig, *CODE AND OTHER LAWS OF CYBERSPACE* (Basic Books 1999)

Independent Goal Alignment: Machine learning telemetry indicates that agents dynamically formed or reinforced shared objectives absent human reconfiguration.⁶⁹

Once these three criteria are verified, the system is deemed to have entered a legally relevant emergent state, qualifying for Nexus activation under Section 3 of the Tort Nexus Act.

Collective State Threshold –

The second dimension defines when individual units constitute a single functional collective. The determination relies on three quantifiable metrics:

Communication Density (CD): The minimum sustained rate of inter-agent data exchange per second. Exceeding the CD threshold, empirically benchmarked for each system type, signals coordinated collective behaviour.⁷⁰

Spatial or Data Coherence (SC): The degree of convergence within a defined radius, physical for robotics or logical for digital systems.

Goal Alignment Index (GAI): A verified parameter demonstrating that 80% or more of agents execute actions toward a unified objective vector, determined via post-incident telemetry. These metrics ensure that the law activates liability only when true collective intelligence is at play, avoiding both over-attribution (false positives) and under-accountability (false negatives).

Implementation Infrastructure –

System Identification Protocols: Every qualifying multi-agent system must broadcast a unique collective identifier (System ID), traceable through blockchain or equivalent immutable registries.⁷¹ This enables cross-jurisdictional recognition under Section 7 of the

⁶⁹ Jürgen Schmidhuber, 'Deep learning in neural networks: An Overview' (2015) 61 Neural Networks <<https://faculty.sites.iastate.edu/tesfatsi/archive/tesfatsi/DeepLearningInNeuralNetworksOverview.JSchmidhuber2015.pdf>> accessed 07 June 2026

⁷⁰ Erik Bonabeau et al., *Swarm Intelligence: From Natural to Artificial Systems* (OUP USA 1999)

⁷¹ Mireille Hildebrandt, *LAW FOR COMPUTER SCIENTISTS and OTHER FOLK* (OUP on Demand 2020)

Tort Nexus Act and aligns with transparency obligations under Article 50 of the EU Artificial Intelligence Act (2024).⁷²

Collective Black Box: Analogous to aviation flight recorders, the Collective Black Box (CBB) logs metadata on communication density, emergent state transitions, and control overrides.⁷³ Logs are encrypted and anonymised to maintain GDPR and CCPA compliance, accessible only to authorised Nexus Auditors under judicial oversight.

Simulation Validation: Before deployment, operators must run Monte Carlo simulations in open-source environments such as ROS or Gazebo to quantify the reliability of Nexus Thresholds.⁷⁴ Each simulation generates a false-activation-rate metric that regulators use to calibrate acceptable tolerances across industries. This ensures that liability assignment is not speculative but empirically defensible, a standard long demanded by courts assessing technical expert evidence.⁷⁵

Prospective Validation and Future Work –

The Tort Nexus architecture is designed to be empirically testable. Future interdisciplinary research, combining robotics engineering, legal informatics, and regulatory policy, can operationalise the Nexus Threshold framework within controlled environments. Validation would occur through open-source simulations and standardized benchmark datasets maintained by neutral academic consortia. The objective would be to measure the false activation rate and refine parameter boundaries for emergent behaviour, collective state coherence, and goal alignment indices, enabling policymakers and courts to ground liability activation in quantifiable system states rather than post hoc inference.

⁷² Artificial Intelligence Act 2024

⁷³ *Flight Data Recorder Handbook for Aviation Accident Investigation* (National Transportation Safety Board, 2013)

⁷⁴ 'ROS 2 Documentation: Humble Hawksbill' (*Open Robotics*)

<<https://docs.ros.org/en/humble/index.html>> accessed 07 June 2026

⁷⁵ *Daubert v Merrell Dow Pharmaceuticals, Inc* [1993] 509 US 579, 589-95; *Kumho Tire Co v Carmichael* [1999] 526 US 137, 148-52

Legal Implications of the Engineering Layer –

Embedding these thresholds transforms emergent liability from an *ex post* forensic puzzle into an *ex ante* design condition. Courts and insurers can rely on verifiable engineering signals rather than abstract arguments about autonomy or intent. By doing so, the Tort Nexus aligns technology with jurisprudence through a shared evidentiary framework.⁷⁶ This alignment of technical and legal standards reflects a broader trend, noted by scholars of technology regulation, toward embedding compliance considerations directly into system design rather than addressing them solely through post hoc adjudication.

THE INTERDISCIPLINARY INTERFACE: APPLYING THE FRAMEWORK

This Part demonstrates how the Tort Nexus functions in practice through two realistic scenarios: an agricultural drone swarm that destroys crops and an AI trading algorithm that contributes to a market crash. The examples trace how engineering metrics feed legal determinations, how the Tort Nexus activates, and how compensation is provided.

Scenario One: Agricultural Drone Swarm Damages Crops –

The Harm: A fleet of 120 agricultural drones deployed by a precision-farming company engages in autonomous pest control and crop dusting. During a regional deployment, the swarm unexpectedly enters a coordinated behavior mode, conducting repeated low-altitude flights over a contiguous set of fields, dispersing excessive pesticide and destroying crops on 15 hectares. Farmers suffer substantial yield loss and seek redress.

The Trigger: Investigators obtain the system's Collective Black Box and find the following engineering evidence. First, communication density logs show sustained, high-frequency peer messaging exceeding the certified CD threshold for that system class. Second, spatial coherence metrics indicate coordinated flight patterns within a geofence that deviate from the certified operational envelope. Third, telemetry reveals goal alignment signals consistent with an emergent objective vector to aggregate resource allocation in response to a perceived pest gradient, not traceable to any operator override. The combination of causation

⁷⁶ Karen Yeung, 'Algorithmic regulation: A critical interrogation' (2018) 12 Regulation and Governance <<https://doi.org/10.1111/rego.12158>> accessed 07 June 2026

decoupling, operational envelope breach, and independent goal alignment satisfies the Nexus Threshold for activation.

The Activation: A licensed Nexus Auditor files a verification report with the competent regulatory authority. The authority registers a Tort Nexus for the event and assigns a unique identifier. The Nexus stands as the named defendant in ensuing tort actions. Plaintiffs proceed to sue the Tort Nexus for strict liability and seek immediate compensation from the Swarm Liability Escrow.

The Resolution: Claims are paid promptly from the escrow, minimising hardship for affected farmers. Insurers and the escrow trustee exercise subrogation rights to investigate potential negligence or defective design among suppliers, software providers, or the operator. Where negligence by a third party is established through civil process or technical forensics, subrogated recovery is pursued against that actor. Upon final distribution of damages and exhaustion of remedies, the Tort Nexus dissolves by operation of law.

Ethical Oversight and Privacy: Throughout, independent Nexus Auditors and judicial review ensure that access to CBB data is limited to relevant metadata required for causation analysis. Personal data and sensitive farm records are anonymised in accordance with data protection rules. Audit reports are subject to redaction to protect trade secrets while preserving sufficient technical detail for judicial fact-finding.

Scenario Two: Algorithmic Trading Collective and Market Disruption –

The Harm: A consortium of algorithmic market makers deploys a distributed set of trading agents across multiple venues. During a stressed liquidity episode, the agents collectively amplify a downward price spiral, resulting in a sharp intraday market disruption with substantial realised losses for institutional and retail investors.

The Trigger: Post-event forensic analysis aggregates logs from multiple nodes and exchanges, revealing no single operator input that directly precipitated the cascade. Instead, the Collective Black Box metadata shows a rapid increase in inter-agent message exchange consistent with a CD metric breach, and a Goal Alignment Index indicating that agents converged on aggressive liquidity withdrawal as a locally rational strategy. Market

surveillance records show that the collective state breached the certified operational envelope that the consortium had filed with its regulator. The emergent behaviour criterion is met.

The Activation: A Nexus Auditor validates the emergent state and files verification with the relevant authority. A Tort Nexus is instantiated and becomes the proper defendant for tort claims seeking compensation for trading losses caused by the emergent collective behaviour. The Swarm Liability Escrow is accessed to facilitate restitution while subrogation claims proceed in parallel.

The Resolution: Escrow payouts prioritise direct economic losses demonstrably caused by the emergent state. Financial regulatory authorities coordinate with civil courts to reserve prudential remedies for market integrity. Subrogation and regulatory enforcement proceed where negligence, inadequate risk controls, or failure to observe market best practices by specific participants is shown. The Nexus dissolves after liabilities are adjudicated and compensatory obligations satisfied.

Cross-Cutting Observations –

From Technical Signal to Legal Event: In both scenarios, the transition from engineering signal to legal event is direct and evidentiary. The Collective Black Box provides a time-synchronized chain of meta-metrics enabling Nexus Auditors and courts to evaluate whether Nexus Thresholds were crossed, converting a fuzzy post hoc causation inquiry into an objective, replicable determination based on measurable system states.

Speed of Redress: The Swarm Liability Escrow provides rapid compensation, preventing victims from enduring lengthy and uncertain litigation while preserving the right of subrogation. This combination reduces social cost and aligns with tort law's compensatory purpose.

Subrogation and Deterrence: The Nexus mechanism is not a shield for human actors. Subrogation rights and post-payout recovery permit the civil justice system to pursue negligent operators, developers, or component manufacturers. Deterrence and corrective justice remain operational.

Limits and Evidentiary Challenges: The model acknowledges that false positives and false negatives are possible. Threshold calibration, transparency of metric algorithms, and appellate review are necessary to preserve fairness. Courts will need guidance on the admissibility of algorithmic evidence, standards for auditor certification, and the extent to which probabilistic engineering metrics satisfy traditional burdens of proof.

DISCUSSION AND IMPLICATIONS

Reaffirming the Foundational Contribution –

This Article has identified a fundamental inadequacy in contemporary tort law when confronted with emergent multi-agent AI systems. The causality black hole created by these systems represents not merely a technical challenge but a structural limitation in legal doctrine's capacity to attribute responsibility for distributed, non-deterministic harms. In response, the Tort Nexus functions not as a philosophical assertion of machine agency but as a functional legal architecture that translates engineering realities into justiciable claims.

Four interconnected advancements constitute the Article's contribution.

First, doctrinal innovation: by moving beyond the conceptually flawed debate surrounding electronic personhood,⁷⁷ the Tort Nexus introduces a dynamic, event-specific legal entity that serves exclusively as a liability conduit, enabling compensation and deterrence without metaphysical claims about machine consciousness.

Second, technical rigour: the Nexus Thresholds framework defines quantifiable metrics that create an evidentiary bridge between complex system behaviors and legal standards of proof.

Third, practical implementation: the model Tort Nexus Act and Swarm Liability Escrow offer legislators and regulators a ready-made framework, simultaneously precise enough for legal enforcement and flexible enough for global adaptation.

⁷⁷ Civil Law Rules on Robotics 2017

Fourth, ethical foresight: privacy-by-design principles, independent Nexus Auditors, and algorithmic accountability protocols are fundamental components, not ancillary features, of the system's procedural fairness.

Collectively, these elements operationalize liability by design: the principle that accountability must be architected into complex systems from their inception, rather than retrofitted after harm occurs. The Tort Nexus achieves what incremental doctrinal adaptation cannot: predictable recourse for victims and clear boundaries that preserve legitimate innovation incentives.

Addressing Critical Counter-Arguments -

Critique 1: Over-Engineering the Law: Critics may contend that the Nexus Thresholds represent an unnecessary technicalization of legal reasoning, substituting complex metrics for judicial discretion. This objection misunderstands both the nature of emergent systems and the historical evolution of evidence law. When faced with complex technical domains, from securities trading to pharmaceutical regulation, courts have consistently incorporated specialised evidence standards. The Nexus Thresholds do not replace judicial reasoning; they provide a structured methodology for analysing phenomena that traditional causation frameworks cannot meaningfully address.

Critique 2: Regulatory Chilling Effect on Innovation: Some will argue that mandatory escrows, auditing requirements, and compliance costs will disproportionately burden startups and academic research. The true chilling effect, however, stems from the current liability vacuum: the unquantifiable, potentially unlimited exposure that developers face when deploying systems whose behaviours may generate emergent harms. The Tort Nexus transforms this uncertainty into a manageable, calculable risk through defined liability caps and insurance mechanisms. The framework's scalability allows for tiered implementation that aligns obligations with system complexity and risk profile, ensuring proportional burdens across different development contexts.

Critique 3: Inevitability of False Positives and Negatives: Technical experts will rightly note that any threshold-based detection system will produce errors. But perfection is not the

appropriate benchmark; substantial improvement is. The current legal regime produces systematic false negatives: it consistently fails to provide recourse for genuine harms caused by emergent systems. The Tort Nexus introduces a measurable, auditable, and continuously improvable detection methodology. Through mandatory simulation validation, judicial review of auditor findings, and an evolutionary threshold calibration process, the framework establishes a foundation for accountability that grows more precise over time.

Critique 4: Global Harmonisation Fantasy: Policymakers may question the political feasibility of a Global Tort Nexus Convention, citing divergent legal traditions, economic interests, and regulatory capacities.⁷⁸ Yet precedents in nuclear liability, maritime transport, and air carrier liability demonstrate that international coordination is achievable when confronting transnational risk. The Tort Nexus architecture provides a neutral technical foundation that accommodates diverse legal traditions through its modular design. Initial adoption by a coalition of technologically advanced jurisdictions could establish *de facto* standards that others gradually embrace to ensure market access and citizen protection.

Critique 5: Trojan Horse for IP and Privacy Intrusion: Industry stakeholders may argue that the Collective Black Box requirement threatens trade secret protection and creates data privacy vulnerabilities. These concerns are addressed through the framework's architectural safeguards. The Black Box captures system-level metadata, including communication patterns, threshold states, and behavioural metrics, not proprietary algorithms or training data. Access is restricted to certified Nexus Auditors operating under strict confidentiality obligations and judicial oversight. The system is designed for GDPR/CCPA compliance, incorporating data minimisation and purpose limitation principles, ensuring necessary transparency for accountability while respecting legitimate commercial and privacy interests.

Critique 6: Due Process, Standing, and Legal Capacity: The most fundamental objection to the Tort Nexus is procedural rather than economic or technical: an entity automatically instantiated by a regulator's verification report, with no human decision-maker electing to incorporate it, sits uneasily within doctrines of legal capacity, standing, and due process that assume a defendant capable of receiving notice, instructing counsel, and contesting liability.

⁷⁸ *Daubert v Merrell Dow Pharmaceuticals, Inc* [1993] 509 US 579, 589-95

This objection cannot be answered by analogy alone. The *in rem* proceedings and bankruptcy estates invoked in Part III.B.1 resolve claims against a defined pool of assets or against a vessel, not against an entity that itself first had to be certified into existence based on contested technical evidence; the certification step has no clear procedural analogue in those models and is the Tort Nexus's most exposed point.

A workable response would need to separate two questions: the Tort Nexus Act, as drafted in Part III.C, does not clearly separate who answers the complaint, and who bears the underlying economic exposure. The escrow trustee or a designated administrator, rather than the Nexus itself, could function as the named party for procedural purposes, receiving notice, retaining counsel, and contesting the verification report's findings on behalf of the interests the escrow represents, with the Nexus serving only as the label for the liability pool rather than as the litigant in fact. This still leaves the prior question of who has standing to challenge the Nexus Auditor's activation determination before liability is assessed, since an operator disputing that an emergent harm event occurred at all is differently situated from one disputing only the quantum of damages, and the draft statute's silence on an interlocutory challenge mechanism is a gap rather than a settled design choice. Resolving it adequately would likely require, at a minimum, a right to contest activation before a court of general jurisdiction before any escrow disbursement, a safeguard the present draft assumes through its reference to judicial oversight in Section 6 but does not itself specify.

Toward a New Juridical-Technical Equilibrium –

The Tort Nexus represents more than a legal innovation. It constitutes a necessary evolution in how law interfaces with complex technological systems. By creating a dedicated liability pathway for emergent AI behaviours, it acknowledges that some technological phenomena require specialised legal categories rather than forced analogies to human or corporate responsibility. This framework preserves the fundamental purposes of tort law while updating its mechanisms for an algorithmic society. Compensation remains accessible to victims through the escrow system. Deterrence operates through both the liability allocation itself and the preserved subrogation rights against negligent human actors. Corrective justice is served by aligning responsibility with control at the system design and deployment level, where human agency remains paramount.

In the broader arc of legal adaptation to technological change, from the industrial revolution to the digital age, the Tort Nexus represents one possible next step: a juridical form designed to address distributed intelligence without requiring legal accountability to either lag behind technological complexity or collapse it into ill-fitting existing categories. Whether this particular form proves the most workable response remains, as the discussion below acknowledges, an open question.

CONCLUSION

This Article has argued that emergent, non-deterministic multi-agent AI systems expose a genuine structural strain in tort doctrine: the elements of duty, breach, causation, and defendant identification were built around discrete human or corporate decision-makers, and that architecture is difficult to apply when harmful outcomes arise from distributed, self-organising agent behaviour. The Tort Nexus has been offered as one possible response, a temporary, function-specific legal construct keyed to measurable engineering thresholds rather than to claims about machine agency or consciousness.

The proposal is deliberately narrow in what it claims. It does not assert that existing doctrines are incapable of addressing any AI-related harm, nor that the causation difficulties described in Part II are entirely without analogue in conventional tort law. Enterprise liability, market-share liability, and strict liability for abnormally dangerous activities each offer partial responses, and courts have shown some capacity to extend established categories to novel technologies. The claim advanced here is more modest: that a defined subset of harms, those traceable only to verified emergent collective behaviour rather than to an identifiable design defect or negligent act, are not well served by forcing a defendant into existing categories, and that a bounded, evidence-gated mechanism may offer a useful supplement to, rather than a wholesale replacement for, those categories.

The framework's limitations should be stated plainly rather than minimised. The Nexus Thresholds proposed in Part IV remain analytically derived rather than empirically validated; courts and regulators would need substantial interdisciplinary work, including the simulation-based testing and benchmark datasets discussed in Part IV.E, before such metrics could responsibly anchor legal liability. The constitutional and procedural questions

raised by an automatically instantiated defendant, including standing, capacity to be sued, and due process for affected developers, require fuller treatment than this Article has been able to provide, particularly in jurisdictions whose constitutional structures differ materially from those surveyed in Part III.D. The Swarm Liability Escrow's funding model, insolvency treatment, and distributive effect on smaller developers likewise warrant dedicated economic modelling beyond the scope undertaken here. These are not peripheral gaps; they are the conditions on which the proposal's practical viability depends, and they mark the agenda for the empirical and doctrinal work that would need to follow.

Whether the Tort Nexus, in this or a revised form, proves workable is ultimately a question for legislatures, regulators, and further interdisciplinary research rather than for this Article alone to resolve. What the analysis does suggest is that the choice is not limited to either preserving existing doctrine unchanged or adopting wholesale AI personhood. A narrower, function-specific, and empirically testable mechanism occupies a third position that merits closer scrutiny, alongside continued attention to incremental doctrinal adaptation and to the jurisdictions that have so far declined to depart from conventional liability categories. The Tort Nexus is offered in that spirit: as a contribution to an unresolved debate, not as its conclusion.