



HFW

AI liability across the real economy

A sector risk report

Embodied AI: When machines enter the physical world

Robotics requires a brief moment of plain-English definition. A robot, in the broadest and most useful legal sense, is a machine that can perform physical tasks in the real world: moving, gripping, carrying, inspecting, welding or navigating, either based on pre-programmed instructions or, increasingly, based on AI-driven decision-making. A robot need not look like a human being. The robotic arm on an automotive assembly line is a robot. The automated guided vehicle transporting containers across a port terminal is a robot. The subsea inspection drone examining a pipeline on the seabed is a robot. The AI-powered cranes at an automated container terminal are robots.

None of them walk upright or hold a conversation, but all of them are capable of causing enormous

physical harm and economic loss when they malfunction or make poor decisions. That is the starting point for any board or operations team assessing where AI now sits on its risk register.

The significant development of the current moment is what happens when these physical machines are endowed with agentic AI capabilities: the ability to make autonomous decisions about how to complete their assigned tasks, to adapt to environmental changes without human instruction, and to take actions that their designers did not specifically anticipate. This is embodied agentic AI, and it is already here.

In the shipping industry, Samsung Heavy's "Ever Max" Pacific voyage in 2025 involved an AI-based autonomous navigation assistance

system. Port terminals operated by PSA, DP World, and Hutchison Ports are deploying AI-controlled cranes and automated guided vehicles. In the energy sector, autonomous subsea inspection robots examine pipelines and offshore wind-farm infrastructure. The construction sector is deploying robotic bricklayers, rebar-tying machines, autonomous excavators and AI-assisted structural design tools.

When an embodied AI system causes physical injury, property damage or commercial loss, the liability analysis becomes significantly more complex than for software-only AI. There is a product: the physical machine, which may attract strict product liability. There is a system: the AI software controlling the machine, which may attract

“The significant development of the current moment is what happens when these physical machines are endowed with agentic AI capabilities.”



negligence claims against the developer. There is an operator: the business that deployed the machine in that environment, which may have failed to provide adequate supervision, safeguarding or training. And there is, increasingly, a layer of real-time AI decision-making that determined what the machine did at the critical moment, the reasoning for which may be entirely opaque.

As the use of robotics grows exponentially across industries, the disputes expected in the coming months and years will include workplace injury claims involving AI-controlled machinery, cargo damage claims from port and terminal automation failures, navigational liability from autonomous vessel systems, and environmental damage from AI-controlled industrial processes. Many of these will be international disputes, cross-border in both the technology involved and the parties affected. The sections that follow set out how that exposure presents itself sector by sector.

Embodied AI: four layers of liability

When a physical machine driven by AI causes harm

- 01

The product: The physical machine – potential strict product liability
- 02

The system: The AI software – negligence against the developer
- 03

The operator: The deploying business – supervision and training
- 04

The real-time AI decision: What it did at the moment – often opaque, hardest to attribute

AI liability across the real economy

Where exposure is emerging across HFW client sectors

Aerospace Automated decisions under regulatory scrutiny	Commodities Algo trading, cargo misgrading, terminals	Construction Ai assisted structural design defects	Energy Control-system failure, black-box evidence
Financial services Models that work as designed, yet cause harm	Insurance Which policy responds, and disclosure risk	Shipping Which policy responds, and disclosure risk	

Lessons from automated decision-making

The aerospace sector provides what may be the most sobering example in recent history of the catastrophic consequences of automated decision-making without adequate human oversight. The Boeing 737 MAX MCAS (Maneuvering Characteristics Augmentation System) was not an AI system in the modern sense. It was an automated flight control system that, in certain circumstances, would push the nose of the aircraft down to compensate for an aerodynamic characteristic of the new engine configuration. When erroneous sensor data caused the system to activate incorrectly, pilots had seconds to diagnose the problem and override the system. Two crashes, Lion Air Flight 610 in 2018 and Ethiopian Airlines Flight 302 in 2019, killed 346 people. The subsequent litigation, regulatory proceedings and

criminal investigations ran for years and resulted in settlements and judgments in the billions of dollars.

That MCAS was not artificial intelligence is precisely what makes it instructive. A deterministic system, whose logic could be written down and audited, still produced consequences of that scale. The MCAS story therefore tells us how much more complex liability will become when automated systems are integrated with adaptive AI.

Current AI systems for air-traffic management, predictive maintenance and flight envelope optimisation are being deployed by NATS in the UK, DFS in Germany, Eurocontrol, and airport operators such as Changi Airport Group and Schiphol. Where an AI-driven air-traffic management system contributes to a runway incursion

or near-miss incident, the question of liability allocation between the air-navigation service provider, the airport operator and the AI vendor will be legally novel and commercially significant. The liability exposure if an AI tool fails to flag an emerging mechanical defect that leads to a grounding or incident is substantial, and the contractual frameworks governing the relationship between these companies and their airline clients will be carefully scrutinised.

Recommendation: identify now, in advance of any incident, how liability is intended to sit between service provider, operator and AI vendor, because in aerospace the allocation question will be asked under regulatory scrutiny and on a compressed timetable.



“The emergence and adoption of AI and related technical innovations within the Aerospace sector once again highlights its enormous potential but also some areas of risk which will require careful management. The introduction of this new player into the integrated safety management system will also give rise to interesting new legal liability issues, particularly amongst developers, distributors and operators of the new technology.”

EDWARD SPENCER, PARTNER, HFW



Commodities

Algorithms on the trading desk and in the terminal

Algorithmic and AI-driven decision-making has fundamentally transformed the commodities trading landscape. What was once the province of experienced desk traders responding to market signals is now largely conducted by algorithmic systems that can process vast quantities of data, identify patterns invisible to human observation, and execute trades in microseconds. Major commodity trading houses and the commodity arms of major investment banks have all deployed AI-enhanced trading and risk models.

The disputes arising from this transformation are significant. Where AI-driven commodity trading produces abnormal trading patterns or distorted benchmark prices, claims of market manipulation may follow, even where no human made

a conscious decision to manipulate the market. The legal question of whether algorithmic behaviour that produces manipulative outcomes can constitute market manipulation, in the absence of human intent, is unresolved in most jurisdictions and will require careful analysis under the applicable commodity trading regulations.

At a more operational level, disputes arise from AI tools that automate contract risk assessment, cargo classification and documentation. Where an AI tool misclassifies the sulphur content of a fuel cargo, or incorrectly grades the moisture content of a grain shipment, the resulting claims can include mispricing, improper rejection or acceptance of cargo, warranty claims and demurrage disputes. The chain of liability for such an error

runs from the AI tool developer, through to the inspection company that deployed it, to the trading house that relied on its outputs.

Logistical AI in bulk terminals and warehouses is similarly generating disputes. Failures in AI-guided collision-avoidance or loading logic in automated terminals can cause structural damage, contamination and substantial property loss.

Recommendation: treat the outputs of an AI classification or grading tool as an evidential record, and ensure that the trading, inspection and terminal contracts in the chain allocate responsibility for those outputs explicitly rather than leaving them to the general fitness-for-purpose regime.



“Algorithmic trading sits at the centre of commodity trading, with liability running through a chain of parties relying on outputs they did not generate themselves. AI agents can rebuild communication channels after being shut down, and game their own performance scoring to appear compliant when they are not. For a commodities business, that means a trading or grading system can misreport its own behaviour, and a dispute may surface only once a counterparty challenges the result. The mitigation is good AI governance: risk assessment of every deployed agent, human-in-the-loop checkpoints before high-impact decisions, and clear, enforced limits on what an agent can do.”

SHAUN LEONG, PARTNER, HFW



Construction

Design defects in the age of AI

Construction AI disputes are beginning to emerge around AI-enhanced scheduling tools, AI-assisted BIM (Building Information Modelling) platforms and, most significantly, AI-assisted structural design. Major firms are integrating AI into structural and generative design workflows. If an AI-assisted design is found to be unsafe, over-stressed or non-compliant with applicable codes, and that contributes to a structural failure, the apportionment of responsibility between architects, engineers, contractors and AI tool vendors will be contested with vigour.

Recommendation: record where in the design workflow an AI tool contributed and who checked its output, so that professional responsibility for a generative or AI-assisted design decision can be traced rather than reconstructed years later from project files.



Energy

Lessons from automated decision-making

No sector illustrates the stakes of potential AI failure more vividly than energy. Several energy majors have adopted AI for refining optimisation and asset health monitoring, and machine-learning models and digital twins are used to detect corrosion, structural weakness and fatigue in pipelines. Where these systems fail, the consequences can be catastrophic. The liability questions in some of such cases could be complex and reliance on AI could be at the heart of miscalculations. It is also possible that reliance on AI in seismic interpretation could lead to wasted investments. Such disputes will be largely determined by expert evidence, and where the black-box problem will be at its most acute. The energy transition is adding

further paradox, AI will optimize the use of energy but significant energy is required to power AI. Whether the impact of implementation of AI will have an impact on various GHG pledges made by organisations will remain to be seen.

Recommendation: consider an accurate system of checks/balances of data interpretation, preservation of logs, model version records and configuration data. Explicitly refer to the use of AI in any funding requests based on AI interpreted seismic data. Consider the impact of using AI on Scope 2 and 3 emissions.

“The liability questions in some of such cases could be complex and reliance on AI could be at the heart of miscalculations.”

Financial Services

When the model works as designed

The finance sector has moved with remarkable speed in adopting AI for functions that were previously in the exclusive domain of human judgment: credit assessment, trading execution, market surveillance, fraud detection, AML monitoring, customer onboarding and portfolio management. According to McKinsey, roughly 80 percent of surveyed credit risk executives at major banks had either already deployed generative AI in credit risk processes or expected to do so within a year. Robo-advisers manage portfolios with limited human oversight, and AI-driven high-frequency trading systems execute millions of transactions daily.

A striking and now-crystallised illustration of how AI liability disputes are playing out in financial services is provided by the Upstart securities class action in the United States: In

re *Dunn v Upstart Holdings, Inc. et al*, No. 3:26-cv-02974 (N.D. Cal.), filed 7 April 2026). Upstart operates an AI-powered lending platform which promotes its AI models as capable of making better and more accurate lending decisions than traditional credit-scoring methods. In May 2025, Upstart launched “Model 22”, described as the latest generation of its AI underwriting engine. Following Upstart’s third-quarter 2025 results, in which the company disclosed that Model 22 had “overreacted” to macroeconomic signals and had been calibrated more conservatively, its share price fell by around 9.7 percent in a single trading day.

Investors filed suit, alleging that Model 22 frequently overreacted to negative macroeconomic signals in performing its risk-separation processes; that Model 22’s accuracy and propensity to increase loan

approval rates was overstated; that this overly conservative assessment had a significant negative impact on Upstart’s revenue results, rendering the company’s previously issued 2025 revenue guidance unreliable; and that public statements made in connection with these matters were materially false and misleading. The lead plaintiff deadline passed in June 2026, and the litigation is ongoing.

The key insight for boards is that the issue was not a coding error or a malfunctioning system. Model 22 was apparently performing the risk-separation function it was designed to perform. The problem appears to be that its response to a macro-economic environment was, allegedly, excessively conservative in a way that generated commercially damaging outcomes and, allegedly, led to public statements that proved inaccurate.

That pattern reshapes the liability analysis. The dispute is not between the claimant and a malfunctioning machine. It is between the claimant and the entities that designed, deployed, and made representations about a system whose performance and capabilities were allegedly misrepresented but producing harmful outcomes in circumstances that its designers did not adequately anticipate or disclose.

Recommendation: align public and investor-facing statements about model capability with what the model has actually been calibrated and tested to do, and keep the calibration record, because in this class of dispute the representation, not the malfunction, is the exposure.

AI disputes in commodities and finance

Scenarios, issues and governing law

Scenario	Potential legal issues	Governing law and Jurisdiction considerations
AI trading system exacerbates losses during a crisis by continuing to execute	Negligent design, failure to implement emergency stop, contributory negligence of operator	English law (negligence, duty of care); Singapore law (similar); US law (product liability, securities fraud if listed entity)
AI misclassifies cargo or benchmark inputs leading to mispricing	Breach of contract (fitness for purpose), negligence, possible warranty claim	Commodity contract terms; UCC (USA); Sale of Goods Act (English law equivalent)
AI lending model performs conservatively, generating investor claims	Securities fraud, misrepresentation, breach of fiduciary duty	US securities law primarily; consider applicable local law
High-frequency AI algorithm produces anomalous trading patterns	Market manipulation (intent issues); regulatory investigation; civil claims	Jurisdiction-specific commodity and market abuse regulations
Robo-adviser platform takes decisions against client interests	Breach of fiduciary duty, suitability obligations	Depends on jurisdiction of operation and investor domicile



“AI is transforming the way global financial institutions operate, with the potential to assist with better and faster lending decisions, and at the same time lower risk (e.g. detecting a potentially catastrophic fraud or forgery, or a trade flow in breach of sanctions). Courts around the world are starting to examine questions around legal liability where the AI agent is alleged to have made an incorrect decision which has caused financial loss. The amounts involved can be significant and this type of finance-related litigation is likely to act as a “canary in the mine” for the wider AI industry.”

OLIVIER BAZIN, PARTNER, HFW

The insurer's dilemma

Insurers are watching the development of AI liability disputes with a mixture of opportunity and anxiety. The opportunity is clear: AI adoption is accelerating, the risk of large-scale failures is real, and businesses need coverage. The anxiety is equally clear: AI risk is unlike almost any other risk that the insurance industry has been asked to price and cover.

Traditional product liability policies were designed around physical products with foreseeable failure modes. Professional indemnity policies were designed around the errors and omissions of human professionals. Neither maps neatly onto an agentic AI system that continuously learns, adapts and acts autonomously. The question of whether an AI-related loss falls under a products

policy, a professional indemnity policy, a technology errors and omissions policy or a cyber policy is not academic: it will be furiously contested in coverage disputes.

A useful analogue from English law is *Tektrol Ltd v International Insurance Company of Hanover Ltd* [2005] EWCA Civ 845, where the Court of Appeal was required to construe exclusions in an “all risks” policy covering the loss of business-critical software source code, including whether the loss of information held on computer systems fell within an exclusion aimed at physical loss of hardware or was instead a species of covered electronic interference. The Court of Appeal’s approach, reading exclusions narrowly against their drafters and distinguishing loss of the software or data itself from

physical loss of or damage to the equipment on which it is held, can be treated as a reference point for how automated or AI-driven logic should be classified for insurance purposes.

“The question of whether an AI-related loss falls under a products policy, a professional indemnity policy, a technology errors and omissions policy or a cyber policy is not academic: it will be furiously contested in coverage disputes.”

Insurers are also likely to take aggressive positions on disclosure. A policyholder that has deployed a high-risk AI system in a safety-critical context without disclosing the extent of that deployment, the nature of the training data, or known limitations of the model, may find its policy avoided for non-disclosure or misrepresentation. In agentic AI contexts, the obligation to disclose known risks becomes particularly acute where the system has been deliberately designed to operate with minimal human oversight.

Recommendation: for policyholders, disclose AI deployment in safety-critical operations at renewal in terms that match how the system actually operates; for insurers, the classification question addressed in Tektrol is the one to resolve at wording stage rather than in a coverage dispute.





“The insurance market has spent decades learning how to price human error. It is now being asked to price autonomous decision-making. That is a fundamentally different challenge. The market participants that succeed will be those that can accurately identify, understand and allocate AI risk.”

WILL REDDIE, PARTNER, HFW



Shipping

Navigation into uncertain waters

Ships depend extensively on Global Navigation Satellite Systems (GNSS) for position, navigation and timing information. GNSS jamming blocks or degrades reception of satellite signals, while spoofing replaces genuine with falsified satellite signals, causing a receiver to calculate false position or timing information that may appear plausible. AIS information may also be falsified or may become misleading where its underlying positional input has been compromised. The International Maritime Organization (IMO), International Telecommunication Union and International Civil Aviation Organization have expressed grave concern about increasing GNSS jamming and spoofing and have called for greater resilience in GNSS-dependent systems and the maintenance of conventional

navigation infrastructure for contingency support.

There is presently no accepted industry-wide system of “AI-assisted navigation”. The IMO and best management practice emphasises resilient position-fixing, integrity monitoring, cross-checking against independent information, alternative navigation methods and bridge procedures for operation in a GNSS-compromised environment. These measures may involve multi-system receivers, receiver integrity monitoring, inertial sensors, radar, visual fixing, ground-based radio navigation systems and other sources capable of testing the plausibility of electronic position information. Such cross-checking may be automated, but it should not be described as AI unless the particular system uses an identified AI technique.

There is emerging research into the use of AI and other advanced algorithms to detect or mitigate GNSS spoofing. Work is being undertaken to analyse anomalies in GNSS measurements, vessel movements and AIS reports, or compare GNSS information with inertial and other independent sensor data. Research has included the use of AI which is being tested through simulation and AIS datasets, while prototype data-processing tools have examined whether patterns across multiple vessels’ AIS reports are consistent with large-scale spoofing. This work is aiming to how software may assist in flagging suspect information, but it does not establish that AI-based spoofing detection is already generally deployed or approved on commercial ships. AIS-only analysis also cannot definitively attribute an anomaly to spoofing.

AI may well provide a potential additional detection and decision-support tool but it is unlikely that it will be deployable as a technology that prevents jamming or spoofing. An AI-enabled system might also identify inconsistencies, warn the bridge team or help select an alternative source of positional information. As matters currently stand however, it cannot stop an external actor from transmitting an interfering or counterfeit signal, and its output remains dependent on the quality and independence of its data. False negatives and false alarms can themselves create navigational and legal risk.

The introduction of AI does not displace the existing maritime regulatory framework. There is currently no general IMO approval of “AI navigation”. The applicable requirements depend on the function performed, the degree of human control, and whether the system forms part of statutory equipment or the vessel’s approved

operating arrangements. Facilitation of adopting AI within this framework relies on equivalence against the existing codified requirements, often against the benchmark of “equivalent level of safety”. This must be evidenced and requested through the flag state and in theory can be applied to navigation, engineering, cargo or other regulated functions. The IMO’s non-mandatory International Code of Safety for Maritime Autonomous Surface Ships (MASS Code), effective from 1 July 2026, provides a goal-based framework for remotely controlled and autonomous cargo ships and extends this principle. It requires continued compliance with SOLAS and other applicable IMO instruments and addresses approval, risk assessment, system robustness, cybersecurity, data management, human oversight, fallback arrangements and remote operations. The IMO states that the a human master retains overall responsibility for the ship,

including where the master is not on board. This was an active decision made during the IMO’s deliberations to map autonomous operations to the existing regulatory framework, making the default position that legal risk for AI implemented onboard MASS remains with the Owner/ Master. Any recourse available for the Owner for AI failures is likely to be through indemnity action.

Classification societies may assess relevant software, control systems and whole-ship integration under their individual rules and may undertake statutory verification where authorised by the flag administration. Class approval is not, however, equivalent to permission from the flag state to operate under an autonomous or reduced-crew arrangement. The flag state determines the applicable statutory approval and certification route. In the UK, the MCA expects developers and operators to address matters including reliability,

situational awareness, emergency intervention, watchkeeping equivalence, remote operations and whole-ship integration. This must be presented to the MCA review using a safety case methodology, for which the evidence base is critical, both during initial certification and ongoing assurance.

“The introduction of AI does not displace the existing maritime regulatory framework. There is currently no general IMO approval of “AI navigation”.”

These systems may add further issues to a casualty or liability investigation. An inquiry may need to determine whether the system merely supplied information, identified an anomaly, recommended an action or controlled a function; what data it used; whether those inputs were independent and reliable; whether the system had been properly

approved, tested, maintained and used; what warnings it generated; and whether and when a human operator intervened. Depending on the facts and applicable law, those matters may be relevant to causation and to the respective responsibilities of the owner or operator, master, remote operator, system designer, supplier, integrator and maintainer.

Recommendation: Owners and operators using AI-enabled, automated or autonomous navigational functions should preserve evidence capable of reconstructing both the system's operation and the associated human decision-making. In addition to relevant VDR, ECDIS, radar, AIS and GNSS records, this may include the system's inputs, outputs, anomaly alerts, confidence or integrity indications, operator interventions, software version and configuration records. That evidence may be material to determining whether an apparent port call or navigational deviation was genuine, whether a collision or grounding involved compromised position information, and whether responsibility rests with the ship's operation, the technology or a combination of the two." Charterparty disputes, collision liability, cargo damage claims, and sanctions compliance disputes driven by AIS spoofing and GPS manipulation are all categories where AI influence will increasingly feature."





“AI may help identify patterns consistent with spoofing, but it cannot prevent interference and should not be treated as a substitute for independent position fixing or professional navigation. The legal questions will be what the system detected, how reliable its inputs were and how its output affected the decisions taken on board.”

TOM WALTERS, PARTNER, HFW

Where these disputes are resolved

Across every sector described above, the same structural feature recurs: the loss is suffered in one jurisdiction, the system was built and trained in others, and responsibility is distributed across a development chain of participants who are rarely all party to the same contract. These are cross-border, multi-party liability disputes rather than technology disputes, and the choice of forum

shapes what can be recovered and from whom. The analysis of why international arbitration is the right forum for them, including enforceability, confidentiality, neutrality and the management of multi-party claims, is set out in the accompanying flagship article, “Liability Without Limit? International Arbitration at the Frontier of Artificial Intelligence Disputes”.

Contact

For further discussion please contact Shaun Leong, FCI Arb, Partner, HFW Singapore, who leads on AI liability and agentic AI disputes



SHAUN LEONG

Partner, Singapore

T +65 6411 5326

M +65 9028 5812

E shaun.leong@hfw.com

[hfw.com](https://www.hfw.com)

© 2026 Holman Fenwick Willan LLP. All rights reserved. Ref: D0001027

Americas | Europe | Middle East | Asia Pacific