



DOI: <https://doi.org/10.63056/jllsa.2.7.2026.249>

Journal of Language, Literature & Social Affairs

journal homepage: <https://scholarclub.org/index.php/jllsa>

Attributing Legal Liability for Autonomous AI Agents in Crypto-Asset Transactions: A Comparative Analysis of the European Union, England and Wales and Pakistan



Shahzeb Murad ^a

a. Texas A&M University School of Law

ARTICLE INFO

Received:

June 29, 2026

Revision Received:

July 12, 2026

Accepted:

July 27, 2026

Available Online:

August 11, 2026

Keywords:

artificial intelligence, crypto-assets, liability, autonomous agents, comparative law, European Union, England and Wales, Pakistan.

ABSTRACT

Autonomous artificial intelligence agents in the crypto-asset markets present a distinct attribution of legal liability issue. An autonomous agent can approve, arrange and conduct transactions without special user case consent and the relevant legal relationship is spread among various participants such as developers, AI providers, deployers, users, custodians, crypto-asset service providers and other infrastructure operators. Traditional legal systems usually do not recognize autonomous AI agents as independent legal entities and traditional private-law rules and specialized regimes do not offer a consistent process for determining the legal subject who is held liable for financial damage resulting from autonomous execution. This article has made a comparison of European Union with England and Wales and Pakistan. It takes into account the EU AI Act, Markets in Crypto-Assets Regulation, revised Product Liability Directive, and the proposed AI Liability Directive, as well as the Property (Digital Assets etc) Act 2025, the Financial Services and Markets Act 2000 (Crypto assets) Regulations 2026 in England and Wales and the National AI Policy 2025, the Regulation of Artificial Intelligence Bill 2024 and the Virtual Assets Act 2026 in Pakistan. The analysis shows that every jurisdiction has rules that potentially apply to AI governance, crypto-asset activity, custody, product liability and property rights/institutional responsibility, but no rules providing a methodology to assign legal responsibility when autonomous agents operate across various technical and institutional layers. The article is not a substantive cause of action but a normative attribution model that is the Control-Authorization-Duty-Causation-Rule framework. The framework provides guidance to adjudicators with respect to: determination of the technically and operationally competent actors; scope of authorization; identification of duties; breach or defect and causation; selection and application of the governing liability rule, including allocation of liability among multiple actors and available remedies. The framework distinguishes technical autonomy from legal responsibility: autonomy alone does not create or end liability and the AI system does not need to have legal personality for responsibility to be attributed to an existing legal person. The analysis also identifies “double opacity” – the interaction of algorithmic opacity and transactional opacity – as a distinctive evidentiary problem in autonomous crypto-asset disputes. A worked hypothetical illustrates the framework thru VASP-mediated self-custodied hybrid DeFi and institutional multi-agent architectures. The article argues that reforms to targeted attribution, evidence preservation and risk allocation are preferable to granting autonomous AI agents independent legal personality.

Correspondence to: Texas A&M University School of Law
E-mail address: Shahzebmurad4@gmail.com (S.Murad).



1. Introduction

1.1 The Problem Stated

AI-powered agents are becoming more common to observe portfolios, produce trading signals, rebalance portfolios, and even communicate directly with blockchain financial systems, such as decentralized exchanges, stable coin infrastructure, and more blockchain-based transaction systems. The legally relevant aspect is not only automation, but also the delegating of the functions of selecting transactions and executing them to systems that are capable of doing this without prior approval by a human being at the moment of execution. The concerns are especially urgent in the context of cryptocurrencies, given recent incidents involving AI agents failing in this space. In fact, the use of AI agents in cryptocurrency environments has recently raised a series of practical concerns, as evidenced by various incidents. Approximately 52.4 million LOBSTAR tokens, worth about US\$442,000 at the time, were reportedly transferred by the autonomous AI agent Lobstar in February 2026 due to a quantity-calculation error (Cointelegraph, 2026; Lobstar, 2026). The AI crypto wallet was reportedly hacked via an encoded prompt-injection attack, which transferred around US\$150,000 to US\$174,000 in May 2026 (Giskard, 2026). These incidents are not "dispositive" of the legal issue, which is still to be decided: who, among the human or institutional actors involved, is responsible for the loss?

The combination of AI and crypto-assets brings new attribution challenges, as selection of transactions, sequencing and execution can be delegated to systems that execute transactions without human consent at execution. Unlike standard automation that operates on set rules, AI-powered agents can adapt their behaviour to handle different inputs, communicate with external applications, and even generate transaction strategies that haven't been programmed into the system by the user. However, as far as the law is concerned, technical autonomy does not imply legal autonomy: the critical issue is whether and how the results of the agent's actions can be traced back to an existing natural or legal person. As of this writing, autonomous AI agents are not currently considered as legal persons by the jurisdictions analyzed in this article and are not subject to any direct legal liability. The question of the immediate attribution problem is therefore not that of the software being the liable party, but, which legally recognised natural or legal person is liable for the action accomplished by the use of the software system.

This is a development with implications for the limitations of liability models often based on a process of identifiable human action, institutional action and a relatively trackable chain of decision-making. Agency, contractual, negligence, product, statutory and enterprise-based doctrines can all be implicated in assigning liability for conduct, systems, relationships and/or risk allocation, not necessarily for the initial decision-making at the moment of harm. AI-powered agents are becoming more common to observe portfolios, produce trading signals, rebalance portfolios, and even communicate directly with blockchain financial systems, such as decentralized exchanges, stable coin infrastructure, and more blockchain-based transaction systems. The legally relevant aspect is not only automation, but also the delegating of the functions of selecting transactions and executing them to systems that are capable of doing this without prior approval by a human being at the moment of execution. The concerns are especially urgent in the context of cryptocurrencies, given recent incidents involving AI agents failing in this space. In fact, the use of AI agents in cryptocurrency environments has recently raised a series of practical concerns, as evidenced by various incidents. Approximately 52.4 million LOBSTAR tokens, worth about US\$442,000 at the time, were reportedly transferred by the autonomous AI agent Lobstar in February 2026 due to a quantity-calculation error (Cointelegraph, 2026; Lobstar, 2026). The AI crypto wallet was reportedly hacked via an encoded prompt-injection attack, which transferred around US\$150,000 to US\$174,000 in May 2026 (Giskard, 2026). These incidents are not "dispositive" of the legal issue, which is still to be decided: who, among the human or institutional actors involved, is responsible for the loss?

Existing scholarship has examined AI legal personality, algorithmic liability, autonomous systems, digital-asset property and crypto asset regulation largely as separate questions. A smaller body of scholarship considers AI-enabled financial activity, while cryptoasset scholarship has focused predominantly on regulatory classification, market integrity and custody. Less developed is a cross-jurisdictional framework for attributing legal responsibility where an autonomous AI system simultaneously separates decision-making, wallet control, authorization, execution and economic benefit among multiple actors. This article addresses that gap through a comparative analysis of the European Union, England and Wales, and Pakistan.

1.2 Research Questions and Scope

This paper addresses four interrelated questions. First, how do the EU, England and Wales, and Pakistan currently provide for or allocate legal responsibility for financial losses arising from autonomous AI-agent crypto transactions? Second, what doctrinal

and regulatory gaps arise at the intersection of AI autonomy, crypto-asset regulation, causation and legal liability? Third, how can legal responsibility for autonomous crypto-asset transactions be attributed among developers, providers, deployers, users, CASPs/VASPs and other legally relevant actors without conferring legal personality on AI systems? Fourth, to what extent can the Control-Authorization-Duty-Causation-Rule framework provide a jurisdiction-sensitive methodology for attribution in such disputes?

The scope is limited to legal liability for financial harm caused by autonomous AI agents in crypto-asset transactions. Criminal liability is excluded. The three jurisdictions were selected for their contrasting approaches: the EU represents a supranational, regulatory-intensive model; England and Wales represents a common-law, private-law adaptability model; and Pakistan represents an emerging-market, sectoral regulatory model.

1.3 Methodology

This paper uses a doctrinal and comparative legal approach, which is informed primarily by primary legal sources, and also draws on official policy documents, law reform materials, academic scholarship and selected practitioner or industry sources for factual context. The comparison does not make the assumption that there is an equivalent liability doctrine from jurisdiction to jurisdiction, instead it asks how each jurisdiction answers the following common questions concerning liability: legal status, legal relationship, duty, attribution, causation, liability standard, remedies and enforcement. Primary sources are legislation, regulations, delegated legislation, judicial decisions and official parliamentary sources. Secondary sources include law-reform reports and academic scholarship and institutional studies. Where to practitioner and industry sources are required, they are used only when authoritative judicial or governmental sources are not available to cover emerging technological incidents or developments. The comparison is functional, as the three jurisdictions have different constitutional, regulatory and private-law frameworks. The analysis assumes nothing about the role similarly named doctrines play in law. It rather compares the approaches each jurisdiction takes to defining the legal responsibilities of actors, the duties applicable to the relationship, what happens when things go “rogue”, and who is liable for whose losses in the transaction chain.

2. Autonomous AI Agents and Crypto-Asset Architectures

2.1 Defining Autonomous AI Agents

In this article, an autonomous AI agent refers to an AI-enabled software system that is able to select, sequence, or execute actions using connected tools or digital infrastructure to achieve a specific goal in an autonomous way, with relatively few human interventions during execution. This definition does not include rule-based automation or traditional deterministic smart contracts. The definition is functional and not technological, in that it does not require that the agent have consciousness, general intelligence or be totally autonomous—what is required is that the agent be able to select, sequence or execute transactions without case-specific human consent at the moment of transaction.

Decision autonomy, execution autonomy and economic autonomy are three different aspects of the concept of autonomy, which mean that the system can take the decision about what transaction to perform; the system can submit, sign or broadcast the transaction; the system is in control of assets or funds without human authorization for transaction level. The differences are important because an AI that offers a suggestion to trade is not the same as one that actually makes a trade, and one that manages the wallet is not the same who will actually make a trade. There are multiple aspects of autonomous AI agents that are pertinent to liability attribution:

Table 1: Features of Autonomous AI Agents Relevant to Liability Attribution

Feature	Liability Implication
Autonomy	Decisions made without direct human intervention at execution
Adaptability	Systems learn and change behaviour based on experience
Opacity	Decision-making processes may not be transparent
Non-determinism	Same inputs may produce different outputs over time
Complexity	Multiple actors involved in development, deployment and operation

2.2 Four AI-Crypto Transaction Architectures

Four distinct AI-crypto transaction architectures must be distinguished.

2.2.1 Architecture A AI Agent Acting Through a VASP/CASP

The agent operates through a regulated Virtual Asset Service Provider (VASP) or Crypto-Asset Service Provider (CASP), which provides custody, transaction execution, or other regulated services. Depending on the contractual and technical arrangement, the provider may retain operational authority over execution even though the customer or agent determines the substantive trading instruction. Here MiCA or PVARA obligations may become highly relevant.

2.2.2 Architecture B AI Agent Controlling a Self-Custodied Wallet

The wallet is self-custodied and controlled through credentials, smart-contract permissions, or other technical mechanisms operated on behalf of the user or organization. There may be no VASP/CASP, making this arguably the most difficult attribution scenario.

2.2.3 Architecture C Hybrid Custody/DeFi Execution

The agent uses a regulated custodian for custody while interacting autonomously with DeFi protocols for transaction execution, thereby separating asset custody from transaction decision-making and execution. Responsibility may involve developers, deployers, protocol developers, DAOs, and governance participants.

2.2.4 Architecture D Multi-Agent Institutional System

Multiple autonomous agents operate within or across organizations and may interact with one another, third-party systems, protocols, or counterparties. The agents may initiate transactions or contractual interactions on behalf of different human or institutional principals. This architecture creates distributed attribution problems because control, authorization, execution and economic benefit may be divided among several actors. Professional duties, regulatory obligations, contractual allocation and insurance may therefore become particularly important.

2.3 Double Opacity

A fundamental challenge in attributing liability to AI systems is what the UK Law Commission calls "opacity," also known as the "black box problem" (Law Commission, 2025b, para. 3.45). In the case of a crypto-asset transaction, where the harm occurs, the causation element is a complex one because it involves grasping the causal connection between the AI agent's decision-making

and the resulting harm. If the agent's decision making is materially opaque, it may be more difficult to establish the causal connection between an alleged defect or breach and the resulting transaction. Opacity can limit not only substantively but also the ability of the claimant to be able to gather and submit sufficient evidence to prove the relevant causal chain. But in the crypto world, the anonymity of many blockchain transactions adds to the opacity.

This paper draws on the concept of double opacity, which is a combination of algorithmic opacity and transactional opacity, for analysis purposes. Algorithmic opacity relates to the lack of clarity around how the AI made its decision. The problems of transactional opacity are related to pseudonymous identities, distributed infrastructure, fragmented records, and the dispersion of transaction participants in different jurisdictions. What is novel about the idea is not so much the fact that it names particular instances of opacity, but the fact that it shows how these two interact to produce a special evidentiary problem of legal causation and attribution. These can combine to present a challenge to a claimant in establishing the causal chain.

2.4 The Actor Ecosystem

Table 2: Actors in the AI-Crypto Transaction Chain

Actor	Role	Potential Legal Basis of Responsibility
Data Providers/Curators	Prepare training data	Defective data
Model Developers	Design AI algorithms	Design defects
AI Model/API Providers	Supply underlying models	System defects
Testers	Validate performance	Failure to identify risks
Software Developers	Integrate AI	Integration defects
Wallet/Infrastructure Providers	Provide execution infrastructure	System/contractual liability
Custodians	Hold assets	Fiduciary/custody obligations
VASP/CASP	Provide regulated services	Regulatory/contractual duties
Oracle/Data Providers	Supply external information	Defective data/misrepresentation
DEX/Protocol Operators	Execute transactions	Protocol/contractual liability
DAO/Governance Participants	Control protocol rules	Governance/attribution
Sellers	Distribute systems	Product liability
Deployers/Operators	Run AI in production	Negligent deployment
End Users	Utilize outputs	Misuse, unauthorized actions

Not all actors can be assumed as a liabilities in all architectures. The significance of each actor is determined by the particular architecture, jurisdiction and applicability of the liability ruling.

2.5 The Legal Personality Problem

None of the three legal frameworks analyzed, currently, establishes a general legal-personality regime for autonomous AI agents as independent legal persons with the ability to be held civilly liable. Accordingly, the question of who is liable for the actions of the agent is not just software versus some other party, but who is the legally recognized party or entity to be held responsible for those actions (Law Commission, 2025b, para. 5.1). The concept of conferring a legal personality on AI systems has been suggested as a “radical” solution (Law Commission, 2025b, para. 5.1). If there is no established legal framework for liability, the legal consequences of the use of AI tend to remain attached to an existing legal or natural person in relation to a relevant doctrine, including contract, agency, negligence, product liability and statutory liability.

3. Conceptual Foundations of Legal Attribution

3.1 Agency Law

When an autonomous software system is acting under the authority of a human or institutional principal, the concept of attribution, as it is understood in agency law, can help with the analysis. The analogy, however, should be used with caution since an AI system is not an agent in the doctrinal sense simply by virtue of the fact that it may do what a human agent does. It was required to differentiate the agency as a doctrinal relationship from an AI system's operational and effective execution of delegated tasks. It does not mean the same thing as legal agency and is not functional autonomy. A court could examine how the principal manifested, gave directions, set up the system, and delegated authority to see if the transaction took place within the scope of authority that is attributed to the principal. It is not clear whether agency concepts would be extended to autonomous AI systems (Chesterman, 2021).

3.2 Vicarious Liability and Enterprise Responsibility

An enterprise may be vicariously liable or have an enterprise liability only if the legal requirements of the jurisdiction where the AI system is used are met. That an AI system carries out the relevant act independently does not without more be sufficient for a legal relationship between the enterprise and the human or institutional actor to give rise to vicarious liability to be established. Direct negligence in deployment or supervision may be more straightforward than vicarious liability in the case of AI-agents. A business can be directly negligent in installing an insufficiently tested system, not setting up proper transaction caps or not supervising an agents activities.

3.3 Contractual Liability

There can be contractual liability if the AI agent's actions violate a contractual condition. This needs to involve a contractual arrangement between the claimant and a defendant and that the AI agent's actions are a breach to their defendant. Contractual relationships can be established in terms of use, service agreements, smart-contract relationships.

3.4 Negligence

Negligence may provide a basis for liability where a duty of care exists, the duty is breached, and the breach causes loss. The key difficulties in AI-agent contexts are establishing a duty of care owed by developers or deployers, and proving causation where AI decision-making is opaque. In England and Wales, proximity, foresee ability, assumption of responsibility, reliance, and policy considerations are particularly relevant. The developing case law on blockchain developer duties, particularly Tulip Trading, illustrates the difficulty of establishing such duties.

3.5 Product Liability

Product-liability regimes may provide a route where the relevant AI system or software falls within the applicable statutory definition of a product and the statutory requirements for compensable damage are satisfied. The revised EU Product Liability Directive expressly brings software, including AI systems, within its product-liability framework (Directive (EU) 2024/2853). However, its compensable-damage provisions do not establish a general strict-liability action for pure economic loss arising solely

from a financial transaction. The position under English law must be analyzed separately under the Consumer Protection Act 1987 and applicable common-law principles.

3.6 Fiduciary Duties

Fiduciary duties could exist when a VASP/CASP and/or custodian has crypto-assets in the custody of a client. The duties are not necessarily imposed solely because an actor is involved in the handling of crypto-assets, but rather on the basis of a relationship, the character of the duty, and the applicable law. It would depend on the nature of the relationship whether or not the relevant AI operation is considered to be within them.

3.7 Restitution and Proprietary Remedies

Where there is a case of crypto-assets being found and identified, restitution and proprietary remedy may be available. Proprietary remedies (claim to specific property) are different from personal restitutionary remedies (claim for value). The Property (Digital Assets etc) Act 2025 is to ensure that a thing (including something digital/electronic) shall not be prevented from being an object of personal property rights by being neither a thing in possession nor a thing in action. Nor does the Act establish a general legal liability regime for AI-agent transactions.

4. The European Union

4.1 The EU AI Act

The EU AI Act of 2024 (Regulation (EU) 2024/1689) is one of the most comprehensive regulations on AI worldwide, dating back to June 2024. It brings in a risk-based approach and requirement with obligations that are commensurate with the risk. Whereas for high-risk AI systems, the Act sets strict requirements such as for risk management systems, data governance with data quality and bias-related risks, documentation and transparency, human oversight, cybersecurity safeguards and pre-market and post-market controls. As per Article 26, systems that pose a high risk must have human oversight, have logs automatically generated for 6 months and report incidents. Article 99 establishes differentiated administrative-penalty ceilings: the highest ceiling up to €35 million or, for undertakings, 7% of worldwide annual turnover applies to infringement of the Article 5 prohibitions; lower ceilings apply to other categories of non-compliance (Regulation (EU) 2024/1689, art. 99).

The AI Act establishes regulatory obligations for providers and deployers but does not itself constitute a comprehensive EU-wide civil-damages regime for losses caused by autonomous AI systems. Questions of private-law liability therefore continue to depend substantially on other EU instruments and national law. It is important to note that AI-enabled crypto trading is not automatically high-risk merely because it involves financial transactions. The AI Act's financial-services high-risk use cases are much narrower, including consumer creditworthiness/credit scoring and life/health insurance risk assessment and pricing. Classification depends on whether the system falls within specified Annex III use cases.

4.2 MiCA and Crypto-Asset Regulation

Regulation (EU) 2023/1114 on Markets in Crypto-assets applies fully from 30 December 2024 (Regulation (EU) 2023/1114). MiCA regulates issuers and crypto-asset service providers through obligations concerning authorization, governance, conduct, custody, operational arrangements, and market integrity. It does not establish a general attribution rule for losses caused by autonomous AI agents operating outside or across those regulated relationships. MiCA Article 67 introduces prudential safeguards allowing firms to comply by maintaining own funds or obtaining professional indemnity insurance covering errors/omissions, regulatory breaches, business disruption/system failures and, where applicable, liability to clients under Article 75(8) (Regulation (EU) 2023/1114, art. 67). Article 75(8) provides that a CASP providing custody and administration is liable to clients for loss of crypto-assets or means of access where the loss results from an incident attributable to the CASP.

Importantly, Commission Delegated Regulation (EU) 2025/416 requires certain crypto-asset trading-platform records to identify the participant, the client, the person or computer algorithm responsible for the investment decision, and the person or algorithm responsible for execution (Commission Delegated Regulation (EU) 2025/416). Commission Delegated Regulation (EU) 2025/885 is principally a market-abuse technical standard. Within its requirements for persons professionally arranging or executing

crypto-asset transactions, it defines algorithmic trading as trading in which a computer algorithm automatically determines individual order parameters with limited or no human intervention (Commission Delegated Regulation (EU) 2025/885). Commission Delegated Regulation (EU) 2025/1140 establishes detailed record-keeping requirements covering crypto-asset services, activities, orders and transactions and requires records concerning investment decisions made by a person or computer algorithm within a CASP in specified circumstances (Commission Delegated Regulation (EU) 2025/1140).

These provisions illustrate that EU crypto-asset regulation can identify algorithmic systems as relevant components of regulated activity without treating those systems as independent legal persons. EU crypto regulation already provides a regulatory example of separating algorithmic decision-making from human/institutional accountability, but it does not transform that identification mechanism into a comprehensive civil-liability attribution rule. This regulatory architecture is therefore consistent with, and provides contextual support for, the Control-Authorization-Duty-Causation-Rule approach of locating legal responsibility in identifiable legal actors rather than in the autonomous system itself.

4.3 The Revised Product Liability Directive

Directive (EU) 2024/2853 on liability for defective products extends product liability to software and AI-integrated products (Directive (EU) 2024/2853). The Directive was adopted on 23 October 2024, published on 18 November 2024, and entered into force on 8 December 2024. It must be transposed by Member States by 9 December 2026. Following a May 2026 corrigendum, the Directive applies to products placed on the market or put into service after 8 December 2026 (Directive corrigendum 2026/90364; Directive (EU) 2024/2853, art. 2(1)). Key provisions include broad definition of product including software and digital content; broad definition of manufacturer including AI system providers; strict liability; defect judged by legitimate expectations; disclosure obligations; and liability for software updates. The Directive expressly recognizes that AI systems capable of developing unexpected behavior can still generate product-defect liability (Directive (EU) 2024/2853, recitals). However, compensable damage categories are limited. Article 6 covers death/personal injury, damage to/destruction of qualifying property, and destruction/corruption of non-professional data. The Directive's compensable-damage provisions do not establish a general strict-liability claim for pure economic loss arising solely from financial loss (Directive (EU) 2024/2853, arts. 6-7). As at 30 June 2026, the revised Product Liability Directive had been adopted but had not yet reached its general application date; consequently, its relevance to the present analysis is prospective rather than evidence of an already applicable EU strict-liability regime. The revised Product Liability Directive does not establish a general strict-liability remedy for mere financial loss resulting from an autonomous crypto transaction. Whether a given loss related to a crypto-asset is, in addition, a loss of property or data would depend on the nature of the loss in question and the applicable national law, and would need to be assessed separately under the applicable provisions of the Directive on damage.

4.4 The Withdrawn AI Liability Directive

The proposed AI Liability Directive would have been a significant step toward harmonisation of certain of its aspects, such as evidentiary disclosure, presumptions regarding causation, and more (European Commission, 2022). The Commission first announced its intention to withdraw the proposal in its work programme for 2025 (European Commission, 2025; Reuters, 2025) and the withdrawal came into effect on 6 October 2025 (Official Journal withdrawal notice C/2025/5423). The withdrawal is noteworthy in that it eliminates one proposed harmonization tool specifically targeted at certain evidentiary and causation issues that lie at the heart of autonomous-agent cases, while providing the EU with a more robust regulatory framework *ex ante*. It is important to note that the withdrawal leaves the EU with more robust regulation *ex ante*, but without a corresponding horizontal EU-wide framework to allocate legal responsibility for the conduct of autonomous agents. The importance of the proposed AI Liability Directive is not so much on the establishment of an effective liability framework, but on the approach to regulation it embodied. It would have added to national liability rules in certain AI-related claims, through its proposed mechanism of evidence and causation. Post-withdrawal EU framework is mainly made up of the AI Act, the amended Product Liability Directive at the Union level, MiCA for crypto-asset activities and the private-law regimes of the Member States. This yields significant convergence in the pre-existing regulatory obligations and product liability, though no specific EU measure on horizontal EU civil liability for autonomous AI-agent actions.

4.5 Assessment

The EU approach is very strong: comprehensive ex ante AI regulation, crypto-asset regulation with prudential safeguards, product liability for defective AI products. However, there are some limitations: Pure economic loss is not covered by the PLD; the AILD withdrawal does not unify fault-based liability under 27 national liability regimes; no specific liability framework exists for autonomous AI systems in financial activities; and no cross-regulatory legal liability framework exists for autonomous AI agents in crypto-asset transactions. The EU does not have an EU-wide legal liability framework for the autonomous conduct of AI-agents in crypto contexts, but has combined ex ante AI regulation, crypto-asset regulation and product-liability mechanisms.

5. England and Wales

5.2 The Law Commission's Discussion Paper

The Law Commission's discussion paper on "Artificial Intelligence and the Law" (Law Commission, 2025b) was published in July 2025, exploring issues arising from the use of AI. The paper considers consequences of AI "becoming increasingly autonomous and adaptive (especially with the rise of AI agents)" (Law Commission, 2025b, para. 1.5). Some of the issues are causation, mental elements, opacity, attribution, and possible liability gaps. Importantly, the paper does not include any proposals for law reform. It is important to emphasise that the Commission "is designed to raise awareness... and is part of a process of identifying those areas most in need of law reform" (Law Commission, 2025b, para. 1.7).

5.2 The Legal Personality Discussion

The Law Commission considers whether AI systems should be given legal personality, and asks whether the 'radical' approach of giving AI systems some form of legal personality is worthwhile (Law Commission, 2025b, para. 5.1). The advantages are that it can help address liability gaps and spur innovation. The arguments against are liability shields and functional complexity. The Commission has stated this is a radical approach that should be considered, but has not come up with any legislation adopting that approach.

5.3 Digital Asset Property Rights

The Law Commission also published Digital Assets: Final Report (Law Com No 412) in 2023 (Law Commission, 2023). The Property (Digital Assets etc) Act 2025 received Royal Assent on 2 December 2025, implementing aspects of the Law Commission's recommendations (Property (Digital Assets etc) Act 2025; Law Commission, 2025c). Section 1 provides that a thing, including something digital/electronic, is not prevented from being an object of personal property rights merely because it is neither a thing in possession nor a thing in action. The Act removes the statutory obstacle to treating digital/electronic things as objects of personal property rights and therefore strengthens the property-law foundation for disputes involving digital assets. The Act does not itself create a general legal liability regime for AI-agent transactions. Although the Act extends to Northern Ireland as well, the present analysis focuses exclusively on its implications for England and Wales.

5.4 Crypto-Asset Regulation

The Financial Services and Markets Act 2000 (Crypto assets) Regulations 2026 were made on 4 February 2026 as SI 2026/102, establishing statutory architecture for a new UK crypto asset regime (SI 2026/102; [GOV.UK](https://www.gov.uk/government/consultations/cryptocurrency-regulation), 2026). The Regulations were made on 4 February 2026 and establish the statutory architecture for the new regime, with the principal regulatory regime scheduled to commence on 25 October 2027. At the 30 June 2026 cut-off, the Regulations were enacted but not yet a fully operational liability regime. Certain provisions entered into force earlier to enable FCA preparatory activity and applications. The FCA published final crypto asset rules and guidance on 30 June 2026, which will apply to crypto asset firms authorized under FSMA from 25 October 2027 (FCA, 2026). The emerging UK regime increasingly regulates organizational control and operational resilience around crypto asset activities, even though it does not establish a dedicated legal liability attribution test for autonomous AI agents. The FCA's June 2026 operational-resilience guidance expressly connects crypto asset firms with systems/risk-control and outsourcing requirements.

5.5 Existing Private Law

English private law provides potential liability routes through negligence, fiduciary duties, contractual liability, unjust enrichment, proprietary remedies, and tracing. Relevant authorities include *AA v Persons Unknown* [2019] EWHC 3566 (Comm), where the High Court recognized Bitcoin as property for the purposes of proprietary relief, and *Tulip Trading Ltd v Van der Laan* [2023] EWCA Civ 83, which is relevant to the developing question whether blockchain developers may, in particular circumstances, owe legally recognized duties to digital-asset holders; it does not establish a general duty owed by blockchain developers to all users. Whether these doctrines can accommodate autonomous AI-agent conduct remains uncertain and fact-dependent.

5.6 Assessment

England and Wales combines an incremental private-law approach to AI with increasingly detailed statutory regulation of crypto assets. The Law Commission explores foundational AI-liability questions without proposing legislation, while the Property (Digital Assets etc) Act 2025 strengthens digital asset property status and the FSMA Crypto assets Regulations 2026 establish a future regulatory perimeter. Nevertheless, developments remain unintegrated with a dedicated legal liability framework for autonomous AI agents.

6. Pakistan

6.1 The National AI Policy 2025

Pakistan's National AI Policy 2025, approved by the Federal Cabinet on 30 July 2025, establishes a policy architecture for responsible AI development (Government of Pakistan, 2025; Ministry of IT & Telecommunication, 2025). But a policy instrument is not the same as a statute, regulatory rule, or private-law cause of action.

6.2 The Absence of AI-Specific Legislation

The Regulation of Artificial Intelligence Bill, 2024 was introduced in the Senate on 9th September 2024 (Senate of Pakistan, 2024). As at 30 June 2026, the Bill had not been enacted. The official record of the Senate states that the Bill was not passed nor rejected by the Committee and accordingly was withdrawn from the Committee and disposed of under the appropriate Standing Order. It therefore cannot be an operative source of statutory legal liability rules. The present analysis shall thus be directed to the legal consequences of the lack of enacted AI-specific legal liability legislation and the interaction of the existing private law with the virtual-asset regulatory framework.

6.3 The Virtual Assets Act 2026

The Virtual Assets Ordinance 2025 was promulgated on 8 July 2025 and was formalized as the Virtual Assets Act 2026, receiving Presidential assent on 4 March 2026 (Virtual Assets Act 2026, Act No. XIII of 2026; Pakistan Code, 2026). The Act provides for the establishment of the Pakistan Virtual Asset Regulatory Authority (PVARA) which is empowered to license, regulate and supervise the persons dealing in virtual assets (s. 3). The Act creates a distinct federal regulatory regime for virtual-asset activities. It addresses licensing requirements, criminal penalties of up to 5 years imprisonment or a PKR 50 million fine, Virtual Assets Appellate Tribunal, extraterritorial enforcement, and provisions relating to market manipulation, customer assets, cybersecurity, administrative sanctions, emergency intervention, restitution and corporate-officer liability. 55. Where a body corporate commits an offense with the consent, connivance or neglect of a director, manager, secretary or other similar officer, that officer is also liable for the offense. The Act imposes a fiduciary duty on licensees dealing with customer assets and requires the segregation of customer assets. It also provides a mechanism for customer compensation/protection. The Act defines a Virtual Asset as “a digital representation of value that can be digitally traded or transferred and used for payment or investment purposes” (s. 2). It covers VASPs and Issuers.

6.4 Existing Private Law

Furthermore, Pakistan's private law provides possible avenues of liability thru the Contract Act 1872 (agency, authority, breach of contract, indemnity, and unauthorized acts), tort/negligence principles, the Electronic Transactions Ordinance 2002 (which deals with electronic transactions and recognition of evidentiary value but does not establish an independent legal liability regime for an AI-agent), consumer protection, corporate liability, and rules of electronic evidence under the Qanun-e-Shahadat Order 1984. It is unclear, and fact dependent, whether these doctrines can accommodate the conduct of autonomous AI agents.

6.5 Assessment

The Virtual Assets Act 2026 of Pakistan sets out various actor-specific duties which could become relevant when considering AI-agent transactions, especially where an autonomous system transacts thru a licensed VASP or custodian. Its main gap is not the lack of duties but the lack of an explicit rule connecting autonomous AI behavior to those duties and assigning responsibility when control and execution are distributed among many actors. The National AI Policy signals recognition of AI's importance. However, significant limitations remain: no comprehensive AI governance framework exists; no framework exists for AI agents in cryptocurrency; reliance on existing laws for AI governance is inadequate. Pakistan's effort to enact dedicated AI legislation has not resulted in an enacted AI statute, while its virtual-asset sector is now governed by the Virtual Assets Act 2026. Pakistan has enacted substantial virtual-asset regulation, but the Act does not expressly provide a cross-doctrinal attribution methodology for autonomous AI systems whose control, authorization and execution functions are distributed among multiple actors.

7. Comparative Analysis

7.1 Common Challenges

All three jurisdictions face common challenges: opacity, attribution, legal personality, speed of development, cross-border nature, decentralization, and evidentiary asymmetry.

7.2 Different Approaches

Table 3: Comparative Analysis of Legal Frameworks

Dimension	European Union	England and Wales	Pakistan
AI governance	AI Act	Common law + Law Commission analysis + sectoral regulation	National AI Policy
Crypto regulation	MiCA	FSMA Cryptoassets Regs 2026 + FCA regime	Virtual Assets Act 2026
Digital asset property	National/Member State private law	Property Act 2025	General property/private law
AI legal liability statute	No	No	No
Product liability	Revised PLD	Consumer Protection Act framework	General law
Agent legal personality	No	Discussed, not recognized	No
Algorithm identification	Yes in specified MiCA records	Emerging regulatory framework	Not specifically developed

Dimension	European Union	England and Wales	Pakistan
AI Liability Directive	Withdrawn 2025	N/A	N/A
Custody responsibility	MiCA Article 75(8)	Contract/private law + future regulatory regime	VAA customer-asset/fiduciary provisions
Pure economic loss	Excluded from PLD	Determined by private law	Determined by private law

7.3 The Attribution Gap

The principal regulatory gap identified by this comparative analysis is not the complete absence of legal rules governing either artificial intelligence or crypto-assets. Each jurisdiction examined contains potentially relevant mechanisms through contract, negligence, agency, fiduciary obligations, product liability, proprietary remedies, regulatory duties and sector-specific legislation. The difficulty arises when those mechanisms must be applied to an autonomous agent whose technical operation separates decision-making, authorization, execution, custody, data provision and economic benefit among different actors. The issue becomes therefore one of architectural attribution. Conventionally, legally relevant action is considered to be traceable to a known natural or legal person with a known relationship. Autonomous crypto-agent systems can disrupt that assumption. A developer may control the model but not the wallet; a deployer may control the wallet but not the model; an oracle provider may control transaction inputs; a VASP/CASP may control custody or execution infrastructure; and a user may authorize general trading without authorizing the particular transaction that causes the loss.

The problem is therefore not regulatory absence but attribution fragmentation. The unresolved problem is the absence of a coherent methodology for determining how technical conduct by an autonomous agent should be connected to the legal duties and risk-bearing responsibilities of existing persons and entities. This attribution problem becomes particularly acute in self-custodied and decentralized architectures, where there may be no regulated intermediary and where control, authorization and execution can be technologically separated. This article addresses that gap through the Control-Authorization-Duty-Causation-Rule framework. The framework does not create a new cause of action. Instead, it provides a structured sequence through which existing substantive liability rules can be applied to autonomous-agent disputes while preserving jurisdiction-specific differences in duty, breach, causation, remedies and burden of proof.

8. The Control-Authorization-Duty-Causation-Rule (CADCR) Framework

8.1 Nature and Purpose of the Framework

The Control-Authorization-Duty-Causation-Rule framework is proposed as normative attribution architecture for disputes involving autonomous AI agents and crypto-asset transactions. It is not presented as an existing legal test, an autonomous cause of action, or a substitute for jurisdiction-specific rules of contract, tort, agency, product liability, restitution or statutory liability. The framework addresses a particular structural difficulty created by autonomous systems: the technical actor that performs the transaction may not correspond to the legal actor who designed the system, authorized its operation, controlled the assets, owed the relevant duty or benefited from the transaction. The framework therefore separates technical agency from legal responsibility.

What CADCR Adds Beyond Conventional Liability Analysis

CADCR is not novel because courts have never considered control, authorization, duty, causation or applicable law. It is novel because it integrates these factors specifically around the distributed technical architecture of autonomous crypto-agent systems, where the decision-maker may be an algorithm; wallet control may belong to another actor; authorization may be general rather

than transaction-specific; custody may be separated from execution; data may originate from another infrastructure provider; and causal evidence may be fragmented. CADCR therefore provides a structured methodology for applying existing legal doctrines to the distinctive factual configuration of autonomous crypto-agent disputes.

Table 4: CADCR Core Attribution Dimensions

Dimension	Central Question
Control	Who could design, configure, constrain, monitor, modify or stop the agent?
Authorization	Who legally or technically authorized the agent to act, and within what limits?
Duty	Which legally recognized actor owed a relevant duty to the claimant?
Causation	Did the relevant breach, defect or wrongful act cause the loss, and can that connection be established evidentially?
Rule	Which substantive liability rule governs the dispute and what consequences follow from its application?

The five dimensions should not be seen as five cumulative elements for a new cause of action. " Instead, they offer an ordered attribution architecture for linking autonomous technical conduct with the existing legal rules by the court or regulator. Breach/defect is not a separate CADCR element because it is evaluated within the applicable Rule. "rule" has no ingredient of liability. This is the last step of classification in which the court selects the substantive legal rule – contract, negligence, product liability, restitution, statutory liability – under which the established attribution findings will have legal consequences.

8.2 Control

The concept of control needs to be considered at multiple layers not as a binary one. Relevant forms include: authority over model architecture (design control); authority to place the agent into operation (deployment control); authority to modify or suspend the operation of the agent, to terminate the agent (operation control); authority to modify, suspend or terminate the agent, to modify its prompts, policies, tools or transaction limits (data-input control); authority over private keys, smart-contract permissions or custody arrangements (wallet control); and practical ability to modify, suspend or terminate the agent (intervention control). Control is thus diffused to the extent that various actors own various stages in the transaction chain. The framework does not presume that the actor that has the most technical control is the obligated one. Rather, control is an attribution factor which depends on the duty and liability rule of the dispute.

8.3 Authorization

Authorisation relates to the legal and technical limitations on the agent's authority. The key question is whether the human actually approved the exact transaction. Autonomous systems are systems that are especially engineered to operate without transaction-by-transaction approval. The inquiry should focus on the principal's instructions, the terms of the contracts, permission of the wallet, limits on spending and transactions, permission of the smart contracts, system policies, representations made by the relevant actors about it, whether the transaction was done within the risk parameters represented or reasonably established by the relevant actor. Thus, a legal authorization may be present, but not a technical authorization, or a legal authorization may be present in general, but not the delegated authorization for the individual transaction.

8.4 Duty

After identifying the actors and their control and authorization relationships, the focus is on duty analysis. The duty should be required by the law, not by the framework itself, of the relevant jurisdiction. These can arise from contractual obligations, negligence, duties in fiduciary and agency relationships, statutory duties, product-liability, regulatory duties, custody, or other

recognized private-law duties. This step helps to avoid the chilling of the framework and the automatic transformation of technical control into liability. An actor can have a significant degree of control without a legal obligation to the claimant.

8.5 Breach or Defect

As to breach or defect, it is evaluated under the applicable liability rule. The inquiry may be to whether a claim was negligently designed, whether there was inadequate testing, inadequate monitoring, failure to implement transaction limits, failure to provide warnings, failure to comply with regulations, failure to authorize execution, failure to perform, breach of contract, breach of fiduciary duty, or any other actionable conduct recognized by law. The framework does not, therefore, set a common global standard of care. The decision-maker must first determine a legal rule and then find whether a particular actor's conduct fits the elements of that legal rule.

8.6 Causation and Double Opacity

Causation should be considered apart from attribution. Control or authorization of an agent only establishes that the actor caused the claimant's loss if that control or authorization made a legal contribution to the claimant's loss. The causal inquiry should address the following questions: 1) Was there technical causation (was the agent's transaction the cause of the relevant blockchain event)? 2) Was there factual causation (would the loss have occurred without the relevant breach or defect)? and 3) Was there legal causation (is the loss sufficiently related to the defendant's transactional acts under the applicable jurisdiction's rules)? The double opacity is a unique problem of evidence in autonomous crypto transactions, as described here. Algorithmic opacity relates to an uncertainty of how the AI system arrived at its decision. In transactional opacity, the problem is with pseudonymous identities, distributed infrastructure, fragmented records and the separation of the participants of transactions across jurisdictions. These features, taken together, can make it challenging for a claimant to prove the causal chain. Thus, the preservation of evidence takes a pivotal role in proper attribution. Model versions, prompts, system logs, tool calls, wallet permissions, transaction policies, oracles, API records, smart-contract events, and governance logs and communications regarding deployment and interventions may be relevant evidence.

8.7 Rule

The final substantive phase is to determine the applicable liability rule. Contractual liability, negligence, agency, fiduciary liability, statutory liability, product liability, restitution, proprietary remedies, and sector-specific custody/regulatory liability are potential rules. The consequences of the same facts or circumstances can vary from jurisdiction to jurisdiction, depending on the different rules, duties, causation and remedies for liability.

8.8 Multi-Actor Allocation

Multiple actors may be involved in autonomous-agent disputes with each contributing to the others' contributions. It is therefore necessary to look at relative control, scope of authorisation, existence and content of duties, causal contribution, foreseeability, contractual allocation of risk, economic benefit, insurance, statutory allocation mechanisms and applicable rules governing contribution or joint liability separately in the context of the framework. Economic benefit is a factor, but does not create the liability itself. In the same manner, neither the role of the "user" nor that of the "developer" should be the determining factor in determining liability, but rather should be examined in light of the actor's actual role functions and/or legally relevant duties.

8.9 Architecture-Specific Application

Table 5: Architecture-Specific Application of CADCR

Architecture	Dominant CADCR Question
VASP/CASP-mediated agent	Which obligations arise from the provider's custody, execution and client relationship?

Architecture	Dominant CADCR Question
Self-custodied agent	Who controlled the wallet and agent, and who owed a duty to the injured party?
Hybrid custody/DeFi	How should responsibility be divided between custodian, agent operator, protocol and data providers?
Multi-agent institutional system	Which institution controlled, authorized and benefited from the interconnected agents?

The self-custodied architecture presents the most difficult case because the absence of a regulated intermediary may eliminate an important institutional source of duties, records and compensation mechanisms. By contrast, VASP/CASP-mediated systems may provide clearer attribution because regulated intermediaries are subject to identifiable contractual, regulatory, custody and record-keeping obligations.

8.10 Operationalization

The framework can be operationalized in nine stages. Stage 1 requires identifying all legally relevant actors. In Stage 2, the contractual, agency, fiduciary, regulatory and technical relationships will be mapped. At Stage 3, control of design, deployment, operation, wallet access, data and intervention are determined. The scope and source of authorization is the next step in Stage 4. In Stage 5, the legal obligation that applies to each and every prospective responsible party is identified. In stage 6, you need to decide on the requirement for the related liability rule or defect in the product or system was met or not. The requirement for Stage 7 is to prove factual and legal causation from the evidence available including technical and documentary evidence. In stage 8, the jurisdiction-specific liability rule will be applied and the issue of multiple actors' liability will be decided. In Stage 9, the remedies that could be obtained are established, such as damages, restitution, proprietary relief, contractual remedies, insurance or statutory compensation mechanisms.

8.11 Normative Principle

The core normative principle of the framework is thus the one which holds that the autonomy of execution must not equate with legal responsibility. Legal responsibility should be assigned to the relationship between the agent's technical conduct, control and authorization systems, legal responsibility and the causal connection and the substantive liability rule of the relevant actor. This approach avoids two extremes. The first is treating AI as legally responsible merely because it technically performed the transaction. The second is automatically imposing liability on the user merely because the user deployed the system. Instead, the framework distributes analytical attention across the actors who created, authorized, controlled, operated, supplied data to or benefited from the autonomous system.

9. Worked Comparative Hypothetical

Consider an AI agent with a self-custodied wallet that executes a series of decentralized exchange swaps based on manipulated oracle data, causing a \$2M loss for liquidity providers. The claimant is a natural-person liquidity provider acting outside a professional capacity.

Under EU law, the AI Act may impose regulatory obligations if the system falls within a high-risk category, but it does not itself create a private right of action. In the event that a regulated CASP is involved, MiCA may apply, but where the agent does not have a CASP, the custody and conduct provisions of MiCA do not directly apply. It is the pure economic loss that is not recoverable as a product liability claim under the PLD. National private law would therefore determine whether the developer, deployer or oracle provider owed a duty and breached it.

Under English law, the claimant would need to establish that the developer, deployer or oracle provider owed a duty of care, breached that duty, and caused the loss. Proximity, foreseeability and assumption of responsibility would be central. Proprietary remedies may be available if assets can be traced. The Property (Digital Assets etc) Act 2025 would strengthen the property-law

foundation. *AA v Persons Unknown* provides authority for treating crypto-assets as property for proprietary relief. *Tulip Trading* illustrates the developing nature of developer duties.

Under Pakistan law, the Virtual Assets Act 2026 may impose duties on licensed VASPs if the agent operated through such a provider. If the agent was self-custodied, the Act's customer protection and fiduciary provisions may not directly apply. The Contract Act 1872 and tort/negligence principles would determine whether the developer, deployer or oracle provider owed a duty and breached it. The Electronic Transactions Ordinance 2002 and Qanun-e-Shahadat Order 1984 may assist with evidentiary issues.

With the framework in mind, the actors to be identified in Stage 1 are the developer, the AI provider, user/principal, deployer, the oracle provider, the DEX protocol, and the liquidity providers. Stage 2 involves creating legal connections between the developer and the user, as well as the oracle provider and the user, and the governance/operator relationship between the DEX protocol and the user. In stage 3, control and authorization are established: the developer has control of the code deployed, and the training data; the user has control of their wallet and has a right to general trading but not specific trading, and the oracle provider has control over data inputs. Not having express permission to perform the specific transactions is not enough to address attribution. The court, instead, would have to decide whether the agent's actions were within a general power vested as part of the programming, deployment configuration, contract or other forms of manifestation of the authority of the principal. Stage 4 involves establishing breach/defect: if the agent was vulnerable to an oracle manipulation, it is likely that the developer has a design defect; if the oracle provider supplied defective data, then the provider has a defect; if the deployer had not adequately monitored, then the deployer has a defect. The chain in Stage 5 is designed from developer design to agent decision to oracle data to DEX execution to loss with evidential support. If the evidence shows the relevant loss would not have occurred, but for the transactions made by the agent, the claimant may establish the factual causation. But this discovery does not satisfy the "cause" element of a cause of action: the court must also reach the issue of whether the loss is sufficiently related to the defendant's misconduct and whether the intervening oracle manipulation is a superseding cause. The challenges with evidence are algorithmic opacity and blockchain pseudonymity. The court would have to consider, in view of the rules of the relevant jurisdiction concerning proximity, foresee ability, assumption of responsibility, reliance, and policy, whether the developer had a duty of care owed to that particular claimant. At stage 6, the liability rule to be applied is to be settled, whether it be the potential strict liability (product liability), the negligence (duty of care) rule or the contractual liability rule. Stage 7 involves distributing to several liable actors based on the applicable national law. In Stage 8, remedies such as insurance, compensation mechanisms, restitution and damages or residual loss allocation are to be determined in accordance with the applicable law.

10. Reform Implications

The framework suggests that future reform should not be about assigning an artificial legal personhood to autonomous agents but be about enhancing the ability to attribute and preserve evidence and allocate risks.

Pillar 1 Attribution

More concrete regulatory expectations regarding the allocation of legal liability across the financial services value chain in relation to the use of AI-enabled financial services, including the instances where providers, deployers, and users are liable for the autonomous-agent's actions. Distinct responsibility between providers and deployers should be defined, including by contractual means as appropriate. Statutory definition of responsibility according to control/authorization/duty.

Pillar 2 Evidence

Implications of mandatory preservation of agent decision/execution records to address double opacity and evidentiary asymmetry: model versions, prompts, system instructions, tool calls, wallet permissions, transaction limits, oracles, API requests, execution logs, blockchain transaction hashes, model updates, human intervention, failed transactions, security alerts. Structures to deal with the evidentiary imbalance in autonomous-system conflicts should be put in place.

Pillar 3 Risk Allocation

Self-custodied agents should have transaction level authorization controls such as spending limits, frequency control, and whitelist/blacklist controls. There is a need to build insurance or compensation mechanisms for regulated custodial activity, and also for possibly high risk autonomous-agent deployments. Innovation through safe-harbour/risk-allocation rules, with accountability. If a regulated AI-crypto system is being operated, it may be necessary to have a documented control map to ensure control reform. Machine-readable transaction-authority limits may be needed for authorization reform. Minimum duties for professional deployers could be included in duty reform. The preservation of model/tool/wallet/oracle traces may be important for causation reform.

11. Limitations

There are a number of limitations to this analysis. First, it is doctrinal as opposed to empirical: it doesn't look at real litigation outcomes. Second, as at 30 June 2026, there were no reported judicial decisions in the three jurisdictions which are directly relevant to determining legal liability arising from an autonomous AI agent making a crypto-asset transaction. Thirdly, all three jurisdictions are subject to a very dynamic regulatory landscape, especially in the areas of AI and crypto-asset regulation. Fourth, the analysis for England and Wales is more forward-looking in terms of liability issues with respect to AI-specific liability, as the Law Commission's 2025 discussion paper notes that there are potential issues up for debate for future reform but does not propose any changes to the law. Fifth, liability allocation may change due to contractual private agreements. Sixth, decentralized systems can include actors offshore, which can also result in jurisdictional complexity.

12. Conclusion

The structural weakness of modern liability law is highlighted by the fact that AI agents making crypto-asset transactions are autonomous. The current liability regime has an important structural defect: AI agents making crypto-asset transactions are autonomous. It's not because current law has no way to cover AI or crypto-assets. On the contrary, it is due to the fact that the technical decision-making can be differentiated from the legally relevant relationships by which the responsibility is normally established in an autonomous system. The comparative analysis shows three possible configurations of regulation. The EU has the most comprehensive regulation of AI and Crypto-assets. The AI Act introduces ex-ante obligations for certain AI systems, MiCA regulates CASP and the revised PL Directive introduces strict product liability for software and to exclude PEL for pure economic loss. However, the proposed AI Liability Directive does not leave behind any horizontal EU-wide legal instrument with a clear focus on addressing attribution and causation issues in the context of autonomous AI conduct. Important for this analysis, unquestionably, are the new legal provisions of the EU crypto-asset legislation that explicitly refer to the identification of computer algorithms used in investment decisions and transaction execution, while not viewing the algorithms as legal persons themselves.

In England and Wales there is a more incremental approach. In its discussion paper for 2025, the Law Commission notes that the challenges of autonomy, attribution, causation, opacity and legal personality are "growing issues" but doesn't make a call for law reform. Meanwhile, the Property (Digital Assets etc) Act 2025 enhances the property law framework for digital assets, and the Financial Services and Markets Act 2000 (Crypto assets) Regulations 2026 introduces a future regulatory framework for certain crypto asset activities. Private law is thus still the core of the conflicts between AI agents and their users, especially in contractual and negligent claims, restitution and proprietary rights.

In the case of Pakistan, things are different. While the National AI Policy 2025 is intended to set a national policy framework for responsible and ethical use of AI, it doesn't provide a general statutory liability regime for AI. The Regulation of Artificial Intelligence Bill 2024 was not passed into law. In contrast, the Virtual Assets Act 2026 provides the foreseeable sectoral framework for virtual assets, such as licensing, obligations of the asset service provider towards their customers, fiduciary obligations, cyber-security, enforcement and compensation mechanisms. The lack of legal obligations in the virtual asset space is thus not the main shortcoming, but the lack of a clear link between the legal obligations and autonomous conduct of AI when multiple players are involved in decision making, authorisation, custody and execution.

The comparative evidence is a justification for a more limited understanding of the regulatory gap. No jurisdiction analyzed offers a specific and detailed mechanism to assign legal responsibility to an existing legal entity if an autonomous AI agent conducts an adverse crypto-asset transaction over several layers of technology and institutions. This is particularly a problem in a self-custodied and decentralized architecture where there is no regulated intermediary that can make control, authorization, evidence

and compensation much more difficult. The Control-Authorization-Duty-Causation-Rule framework attempts to address this issue by framing the issue as a structured inquiry into the factors of Control, Authorization, Duty, Causation and Rule. It is intentionally offered as a normative analytical structure, and not as a new cause of action. It is its main attribute to link technical architecture to the current legal frame of reference without having to grant legal personality to autonomous AI systems. There is also a distinction between 'technical causality' and 'legal causality' and the 'evidentiary implications' of 'double opacity' is an interaction between 'algorithm opacity' and 'transaction opacity'.

The overarching norm is that autonomy should not be a legal "detour". Having been deployed should not automatically be a liability either. Rather, responsibility ought to run along the parameters of control, authorization, duty, cause and the substantive liability principle between the actors involved that has legal relevance. With the greater integration of autonomous AI agents in the crypto-asset markets, the courts and regulators will more and more find themselves in disputes where no one has full control of the transaction chain. What is not appropriate, therefore, is to give software a legal personality simply because software acts without any human or institutional control or because software has no relation to the human and institutional actors who design, authorize, control, operate and benefit from those systems. One such approach is through the Control-Authority-Duty-Causation-Rule framework.

References

- AA v. Persons Unknown, [2019] EWHC 3566 (Comm).
- Chesterman, S. (2021). Artificial intelligence and the limits of legal personality. *International & Comparative Law Quarterly*, 70(4), 837-867.
- Cointelegraph. (2026). AI agent Lobstar Wilde accidentally sends \$442K to beggar. <https://cointelegraph.com/news/openai-employee-s-ai-agent-accidentally-sent-442k-to-beggar>
- Commission Delegated Regulation (EU) 2025/416 of 30 January 2025 supplementing Regulation (EU) 2023/1114 of the European Parliament and of the Council with regard to regulatory technical standards specifying the records to be kept of all crypto-asset services, activities, orders and transactions undertaken, 2025 O.J. (L 2025/416).
- Commission Delegated Regulation (EU) 2025/885 of 29 April 2025 supplementing Regulation (EU) 2023/1114 of the European Parliament and of the Council with regard to regulatory technical standards specifying the arrangements, systems and procedures to prevent, detect and report market abuse, the templates to be used for reporting suspected market abuse, and the coordination procedures between the competent authorities for the detection and sanctioning of market abuse in cross-border market abuse situations, 2025 O.J. (L 2025/885).
- Commission Delegated Regulation (EU) 2025/1140 of 27 February 2025 supplementing Regulation (EU) 2023/1114 of the European Parliament and of the Council with regard to regulatory technical standards specifying records to be kept of all crypto-asset services, activities, orders and transactions undertaken, 2025 O.J. (L 2025/1140).
- Consumer Protection Act 1987, c. 43 (UK).
- Contract Act, 1872 (Pakistan).
- Corrigendum to Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC, 2026 O.J. (L 2026/90364).
- Electronic Transactions Ordinance, 2002 (Pakistan).
- European Commission. (2022). *Proposal for a directive of the European Parliament and of the Council on adapting non-contractual civil liability rules to artificial intelligence* (COM/2022/496 final).
- European Commission. (2025). *Commission work programme 2025* (COM/2025/100 final).
- Financial Conduct Authority. (2026). *Cryptoasset operational resilience: Guidance for firms* (FG26/6). <https://www.fca.org.uk/publication/finalised-guidance/fg26-6.pdf>
- Financial Services and Markets Act 2000 (Cryptoassets) Regulations 2026, SI 2026/102 (UK).
- Giskard. (2026). *How Grok got prompt-injected: An X user drained \$150,000 from an AI wallet*. <https://www.giskard.ai/knowledge/how-grok-got-prompt-injected-an-x-user-drained-150-000-from-an-ai-wallet>
- Government of Pakistan. (2025). *National artificial intelligence policy 2025*. Ministry of Information Technology and Telecommunication. <https://moitt.gov.pk/SliderDetail/NWIyMzEwYzktYTIwYyooZTk0LWI5ODUtYjg5NjZmYTYYMTgz>
- GOV.UK. (2026). *Policy note: Draft statutory instrument amending the Cryptoasset Regulations*. <https://www.gov.uk/government/publications/policy-note-draft-statutory-instrument-amending-the-cryptoasset-regulations>
- Law Commission. (2023). *Digital assets: Final report* (Law Com No. 412).

- Law Commission. (2025b). *Artificial intelligence and the law: A discussion paper*. <https://lawcom.gov.uk/publication/artificial-intelligence-and-the-law-a-discussion-paper/>
- Law Commission. (2025c). *The Property (Digital Assets etc) Act 2025 has received Royal Assent*. <https://lawcom.gov.uk/news/the-property-digital-assets-etc-act-2025-has-received-royal-assent/>
- Lobstar. (2026). *Lobstar Wilde*. <https://www.lobstar.wtf/>
- Official Journal of the European Union. (2025, October 6). *Withdrawal notice C/2025/5423*.
- Pakistan Code. (2026). *Virtual Assets Act 2026* (Act No. XIII of 2026). <https://pakistancode.gov.pk/english/sHyuRiF.php>
- Property (Digital Assets etc.) Act 2025, c. 29 (UK).
- Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on markets in crypto-assets, and amending Regulations (EU) No. 1093/2010 and (EU) No. 1095/2010 and Directives 2013/36/EU and (EU) 2019/1937, 2023 O.J. (L 150).
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No. 300/2008, (EU) No. 167/2013, (EU) No. 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828, 2024 O.J. (L 2024/1689).
- Reuters. (2025, February 12). EU ditches plans to regulate tech patents, AI liability, online privacy. *Reuters*. <https://www.reuters.com/technology/eu-ditches-plans-regulate-tech-patents-ai-liability-online-privacy-2025-02-12/>
- Senate of Pakistan. (2024). *Bill summary: Regulation of Artificial Intelligence Bill 2024*. <https://www.senate.gov.pk/en/billsummary.php?bid=1544>
- Tulip Trading Ltd v. Van der Laan, [2023] EWCA Civ 83.
- Virtual Assets Act, 2026, Act No. XIII of 2026 (Pakistan).