




Thought Leadership

The Convergence of AI and Distributed Ledger Technology: Opportunities and Risks

September 2025



Contents

1	Foreword	03
---	----------	----

2	Executive summary	04
---	-------------------	----

3	Convergence of AI and DLT: use cases	07
3.1	Use case overview	07
3.2	How might AI enhance DLT platforms and processes and make smart contracts smarter? Use case #1: AI-augmented smart contract development Use case #2: AI-powered blockchain oracles	08
3.3	How might DLT improve AI training and development and enhance AI capability? Use case #3: Using DLT to unlock data for AI training Use case #4: Financial services-enabled agentic AI	11

4	Legal and regulatory landscape	15
4.1	AI laws and guidance	15
4.2	DLT and cryptoasset-specific laws and guidance	18
4.3	Privacy and data protection	19
4.4	Liability considerations	22
4.5	Financial services regulatory overlay	24
4.6	Market abuse	27

5	AIxDLT implementation: key considerations	28
---	---	----

1

Foreword



Diego Ballon Ossio
Partner, Clifford Chance



Sabih Behzad
Head of Digital Assets &
Currencies Transformation,
Deutsche Bank

Emerging technologies are reshaping the way we live, work, and engage with the world. This paper explores two of the most transformative innovations currently gaining prominence –artificial intelligence (AI) and distributed ledger technology (DLT). While each of these technologies is often examined in isolation, our aim is to explore their potential convergence (AIxDLT), identify promising use cases and consider some of the key legal and regulatory implications that may arise as these technologies evolve and converge.

By embracing innovation, we can unlock new pathways to enhance products, services and operations for both consumers and businesses. The potential benefits of deploying these technologies – including speed, transparency, and scalable growth – are significant. Yet such benefits must be balanced with a commitment to consumer, user and investor protection, business resilience and regulatory alignment to ensure a responsible and inclusive transition to new models, processes and service offerings.

This paper represents a collaborative effort between Deutsche Bank and Clifford Chance. By combining our resources and insights from financial services and legal perspectives, we aim to spotlight ideas that provoke thought and conversation and provide a valuable reference for others navigating this evolving landscape. Transformational change of this magnitude cannot be driven by isolated actors; it requires a collective willingness across the ecosystem to invest in the exploration of responsible innovation. This may include engagement across industry and with policymakers and regulators as these technologies and the regulatory frameworks governing them continue to evolve.

Our hope is that, by contributing to the growing discussion on the potential of these technologies, this paper will help further define the issues for exploration and innovation – whether by building upon existing models or challenging the status quo. These conversations are essential to shaping the future of multiple sectors.

We invite readers to engage with our teams to discuss our work in this space and explore opportunities for collaboration. Asking questions and sharing perspectives and knowledge are vital to advancing the future of AIxDLT. As this field continues to evolve, we remain committed to collaborative engagement and responsible leadership in harnessing these technologies to enhance business outcomes and societal impact.

2

Executive summary

The transformational potential of AlxDLT

AI and DLT are transforming industries. While each of these technologies in isolation can offer benefits, we believe that further value lies in their combination, particularly where the inherent strengths of one technology help address some of the challenges faced in relation to use of the other.

Where the predictive and analytical capabilities of AI are combined with the decentralised and resilient infrastructure of DLT, this synergy of 'AlxDLT' has the potential to accelerate the transformation of a range of products, services and industries by enabling smarter, automated systems that operate with improved transparency, efficiency and trust.

There are a range of use cases emerging or being explored in the finance sector with differing levels of maturity. Imagine a world where a fully autonomous agent could use a person's digital identity and preferences to book and pay for a weekend in Paris across different service providers, or a supply chain system where AI analyses vast troves of data to predict liquidity needs and then automatically triggers smart contracts to optimise working capital. Some AlxDLT use cases are already being tested or rolled out, for example in areas such as treasury management where products such as Ant International's Whale platform are being tested and launched to support 24/7 global treasury operations through AI-powered systems optimising cashflows across multiple currencies and jurisdictions in real time and use of DLT to ensure that all transactions are recorded securely and immutably.

We envisage further use cases for AlxDLT along a continuum from enhancement to convergence. In this paper, we focus on four illustrative use cases that have the potential to transform a range of industries:

- improvements to smart contract development and performance through the ability to test and interrogate their functionality;
- AI-powered blockchain oracles for enhanced reliability when connecting distributed ledgers to real-world data;
- controlled access to private datasets to further AI development; and
- AI agents using blockchain wallets to support payment processes and commerce.

While additional promising use cases are emerging, these examples reflect trends we are observing that explore how the convergence of AI and DLT could: (1) make smart contracts better and more accessible; (2) make blockchain processes 'smarter' in their use of real-world data; (3) improve and enhance AI training and development with DLT-based controls for data holders offering immutable records of data provenance; and (4) leverage digital identity, for example by taking the next steps towards frictionless, payment-enabled agentic AI.

Navigating the rapidly evolving legal landscape

While the development of AIxDLT use cases presents significant potential opportunities, there are many legal, commercial, ethical and technical challenges to navigate, not all of which can be predicted at this stage. The convergence of AI and DLT necessitates consideration of a complex patchwork of legal and regulatory requirements, including evolving technology-specific laws, data protection requirements, contract and private law considerations, as well as sector-specific regulations that could impact use cases for financial services, healthcare, energy, transportation and more. These parallel and overlapping regimes and laws demand careful, case-by-case legal analysis to ensure compliance and to mitigate liability and other risks. In some cases, possibilities for innovation may be in tension with frameworks for accountability and liability, particularly where any use cases move towards true decentralisation.

The trend towards more technology-focused regulation will continue apace as more jurisdictions adopt specific frameworks. Regulatory approaches will likely become more sophisticated as use cases develop. In parallel, courts across the globe are starting to grapple with novel liability issues in relation to both AI and DLT technologies, which will inform how AIxDLT projects are structured and documented.

The last few years have seen significant changes to the legal landscape for AI and DLT, including introduction of AI-focused legislation, such as the EU's AI Act, US state laws, such as Colorado's Concerning Consumer Protections in Interactions with Artificial Intelligence Systems Act (Colorado AI Act) and the Texas Responsible Artificial Intelligence Governance Act, and the Provisional Administrative Measures for Generative Artificial Intelligence Services applicable in China. These laws have developed alongside new comprehensive regulatory regimes governing markets in cryptoassets and cryptoasset service providers. This trend towards more technology-focused regulation will continue apace as more jurisdictions adopt specific frameworks and regulatory approaches will likely become more sophisticated as use cases develop. In parallel, courts across the globe are starting to grapple with novel liability issues in relation to both AI and DLT technologies, which will inform how AIxDLT projects are structured and documented. In addition, AIxDLT use cases can give rise to issues around contractual interpretation and technical errors in the context of smart contracts – and associated liability risks and potential disputes – with considerations differing depending on whether a smart contract complements a conventional natural language contract with 'off-chain' provisions or a code-only smart contract is used.

In this paper, we spotlight some examples of key legal considerations relevant to AlxDLT, noting some of the wider range of relevant legal risks and challenges that may also apply. Our aim is to further develop the vital conversations within businesses, across industries and with policymakers that could help further explore and develop potentially transformative AlxDLT use cases.

While there are numerous hurdles and complexities, for many organisations the potentially transformative opportunities arising from the combination of AI and DLT may significantly outweigh the challenges. The groundwork for this transformation is already being laid and, for many use cases, their implementation will not be in the far-distant future but a shorter-term reality.

Key practical considerations for AlxDLT implementation

Successful and responsible exploration of AlxDLT requires internal and external engagement, a holistic approach and significant investment in infrastructure and talent to fully realise the potential of these technologies. In section 5, we share five practical insights for organisations developing or considering exploring AlxDLT use cases:

- **Take a holistic and collaborative approach to use case exploration and implementation:**
Organisations must consider the impact of proposed projects and business models on their daily operations and adopt a holistic and collaborative approach to implementation. This involves understanding the implications for the entire business and fostering collaboration across various departments, including legal, technology, risk, compliance, sales, HR, finance, and tax.
- **Build internal capabilities and technology literacy to support the changes ahead:**
Building internal capabilities and literacy in AI and technology is crucial to leverage these opportunities fully. Staff and stakeholders need to understand the capabilities and limitations of the technologies being used to ensure that use is appropriate and that there are checks and guardrails.

- **Consider wider impact and alignment with firm policies and culture:**
Organisations must also align new technologies with their policies and culture, considering factors such as customer care, transparency, accountability, energy consumption and sustainability goals.
- **Strategically assess and navigate legal obligations and risk:**
Navigating the complex legal and regulatory landscape requires early due diligence, strategic risk management (including through the contractual framework) and collaboration with trusted advisers.
- **Consider engagement with policymakers, regulators and industry stakeholders, and anticipate new legislation:**
Engaging with policymakers, regulators and industry stakeholders can also be crucial to staying ahead of regulatory changes, understanding sectoral norms and regulatory expectations and shaping emerging policies.

3

Convergence of AI and DLT: use cases

3.1 Use case overview

Distributed ledger technology (DLT)/ blockchain

DLT enables recording and storage of data across multiple ledgers simultaneously. It transforms record-keeping across networks by enabling the use of a shared ledger and uses cryptographic techniques to ensure security and provide immutability of data – the inability to modify data recorded on the blockchain. DLT therefore provides a framework for enhancing trust, efficiency and automation that offers many benefits to various business processes, from financial transactions to supply chain management.

Artificial Intelligence (AI)

According to the EU AI Act, an 'AI system' means "a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs, such as predictions, content, recommendations, or decisions that can influence physical or virtual environments". The multiple capabilities displayed by AI offer many benefits to business processes, such as pattern recognition, increased automation, rapid analysis of vast datasets, and enabling humans to interact with machines in natural human language.

AI and DLT provide the potential to reshape entire industries, not just through their individual capabilities but also in some instances even more significantly through their convergence. While AI brings intelligence, automation and analytical capabilities, DLT offers decentralisation, transparency and data provenance. When combined, the technologies can mutually reinforce each other's capabilities.

DLT can complement AI by potentially playing a role in driving forward data sharing for AI development and by helping address certain key limitations. For example, the distributed ledger's digital record can offer enhanced transparency as to the private datasets used to develop AI models, providing a potentially tamper-proof record relating to key aspects of the models' development and training. This can go some way towards addressing the challenge of explainable AI and any applicable record-keeping requirements.

Conversely, AI can complement DLT in respect of areas of perceived weakness, such as efficiency and real-time data analysis. For example, AI models can enhance smart contracts to make dynamic, real-time decisions based on predictive analytics, market conditions and complex data inputs, enhancing intelligence beyond preset conditions.

Beyond mutual reinforcement, this intersection could also enable entirely new use cases that are 'AI and DLT-native', meaning systems that fundamentally depend on both technologies for their operation. These integrated use cases point to a future where trust, intelligence and autonomy are embedded directly into digital infrastructure.

Theoretical and emerging use cases for AIxDLT include:

- **AI for DLT security and anomaly detection:** Using AI to identify problems that might have occurred on-chain.
- **DLT to unlock private datasets for AI training*:** Advancing AI development through controlled access to private datasets and even, potentially, also allowing for collaborative training of AI models in ways that do not necessitate the sharing of raw data.

- **AI-augmented smart contract development***: Using AI to improve development and accessibility of blockchain-based smart contracts.
- **AI-powered blockchain oracles***: Using AI to help intelligently connect live datasets to blockchain processes.
- **Payment-enabled agentic AI***: Enabling AI agents to use blockchain wallets to support payment processes and agentic commerce.
- **Context-aware digital identities**: Enabling entities to use DLT to securely verify identity and control disclosure for specific interactions with AI, continuously analysing vast amounts of real-time data for context and risk assessment. This could enable data sharing on a need-to-know basis, driven by context, reducing default broad or persistent access.
- **Cross-platform identity verification for AI agents using DLT**: In the future, we expect to see more AI agents operating across different platforms (vs. single platform AI agents available today). Verifiable identity will be required to enable this and DLT could provide a tamperproof ledger for AI agents, ensuring that their identity is not tied to any single vendor or platform.

The integration of AI and DLT can be conceptualised along a continuum from enhancement, where one technology augments the capabilities of the other, to convergence, where their combined implementation enables capabilities that are not possible in isolation. We have outlined a selection of these use cases below to highlight the interplay between these two technologies.

We now consider four of the AIxDLT use cases mentioned above in greater depth, looking at some of the potential applications across industries. Examples of some of the key legal and regulatory considerations for AIxDLT use cases are set out in section 4, with our key practical takeaways for organisations implementing AIxDLT in section 5.




Figure 1: AI and DLT convergence spectrum - examples.

3.2

How might AI enhance DLT platforms and processes and make smart contracts smarter?

Use case #1:

AI-augmented smart contract development

AI is used to improve the development and accessibility of smart contracts

Smart contracts

Smart contracts are software code that is designed to execute automatically upon the occurrence of predefined conditions, for example making a payment or transferring ownership of assets. It is deployed within a DLT environment and may be self-contained or executed within the context of a separate 'wrapper' or written legal contract between the counterparties of a transaction.

Smart contracts hold promise for numerous applications across different industries; however, their adoption has been constrained. Programmability – the ability to embed conditional logic into digital processes using smart contracts or similar mechanisms on a blockchain – and composability – when different smart contracts are able to interact without obstacles – are important features in supporting improvements to smart contracts. One challenge for implementation is that non-technical business users cannot easily understand and verify whether the intended business logic has been captured correctly. This issue is magnified as specialised blockchain development languages, such as Solidity, are commonly used for smart contracts, and even technical staff familiar with more common programming languages may not fully understand the development needs. In addition, there are limited solutions for checking the security and robustness of smart contract code and technical staff development could require significant resources.

AI tools and techniques may help with solving some of these challenges. For example, large language models (LLMs) will make it possible to specify the required business logic within smart contracts using natural language and with limited specialist technical knowledge or skills. LLMs already demonstrate remarkable abilities to translate between natural language and programming code as well as between different programming languages, and they are rapidly improving in capability. AI also has strengths in application testing, where LLMs can generate test cases automatically to validate that different functions work as expected. Many organisations are already using AI to drive developer productivity, albeit focused largely on more common programming languages to date.

As an example, imagine a logistics and shipping platform aiming to launch a new smart contract for supply chain execution. In the initial phase, a developer uses an AI-powered development environment to input high-level natural language specifications (for example, to automate delivery payments for perishable goods based on verified arrival temperature). The AI platform then generates an initial draft of the smart contract code that would operate via a DLT network to offer greater transparency through real-time monitoring and automated payment functionality. As code is generated, the AI platform performs real-time analysis for security vulnerabilities, suggests optimisations, and identifies logical inconsistencies that support the developer in reviewing the code.

While the shipping process is underway, the developer continues to interact with the AI to add additional specifications and even simulate shipping scenarios and the final test cases. The use of AI in this scenario helps optimise the operation of the DLT-based logistics and shipping platform by augmenting the human developer's capabilities and allowing for expedited development.

There are already commercial offerings that use AI to code smart contracts and serve as a user interface for blockchain applications: for example, [Third Web's Nebula](#) and [Nethermind's Audit Agent](#). Prototypes have also demonstrated other exciting possibilities, including a recent [Deutsche Bank partnership with AI firm finaXai to transform tokenised funds servicing with cutting-edge AI](#).

Not only can AI accelerate new blockchain-based products, but it can also lower the barrier to entry for non-technical users, signalling a transformative shift in how blockchain applications can be adopted.

However, while these systems can significantly enhance product development and improve accessibility, they typically still rely on a user's domain awareness and specific knowledge to formulate useful smart contracts. AI can help address known cybersecurity weaknesses, but it cannot solve low-quality prompting or anticipate the additional functionality that a user wishes to achieve. Therefore, it may not deliver expected results if the instructions are too general, for example. The AI system – even if it is domain-specific fine-tuned – would only perform at its peak performance level if there is sufficient awareness of what is to be achieved to guide it.

It must also be acknowledged that certain LLM translations still have issues with errors and misinterpreting user intent, but they are improving rapidly. In time, we expect that continued investment in AI training and fine-tuning as well as increased sophistication in AI use (including improved prompting) will help address these issues. The potential rewards from doing so – enabling automation within all types of code development and translation, as well as wider applications – will be significant.

Use case #2: AI-powered blockchain oracles

An AI-enabled 'hybrid smart contract' is used as a way of connecting real-world data to a blockchain oracle

Oracles

A mechanism that connects blockchains to external systems, enabling smart contracts to execute based on real-world inputs and outputs.

A key limitation of blockchains is that they are self-contained and therefore closed systems, i.e. they do not have access to external data. This is so regardless of whether a blockchain is public or private. This aspect means that smart contracts within a blockchain lack the native ability to initiate outbound connections to the internet, read from external application programming interfaces (APIs), interact with traditional databases or fetch real-world information (such as weather data or stock prices).

Most real-world agreements and applications depend on a level of access to external information. In the absence of outside (off-chain) data, smart contracts are limited to only executing based on on-chain data; for example, tasks, such as cryptocurrency transfers or internal token balance retrieval.

This inherent data isolation is solved using oracles. Oracles are services that act as a bridge, connecting smart contracts to external data sources, such as banking data, real-time price data, and data that exists on other blockchains.

However, one issue that oracles raise is that of centralisation. Centralised oracles can pose a risk in relation to data accuracy and availability, representing a single point of failure. Oracles may also raise potential risks of data tampering and data sufficiency. To overcome these limitations, decentralised oracle networks have emerged. Instead of consuming data from a single source, these oracles leverage numerous independent nodes that collectively retrieve, verify and aggregate data from diverse sources.

While these decentralised oracles significantly enhance reliability, the volume and complexity of real-world data present new challenges. AI can play a role here in acting as an 'intelligent' layer on oracle networks. AI algorithms can be used to perform advanced data verification, anomaly detection and pattern recognition for manipulation that might evade regular aggregation and consensus mechanisms on blockchains.

AI can also enable complex off-chain computations, pushing oracles beyond merely fetching data to becoming an intelligent bridge between blockchains and the real world.

One of the potential use cases for these AI-enhanced smart contracts involves optimisation of trading strategies, where AI is initially used to analyse real-time and historical price data to predict spikes or drops in asset prices and subsequently trigger smart contracts to adjust positions to minimise exposure or halt trades to avoid losses. Similar uses could exist for counterparty default, interest rate movements, liquidity risk and decentralised finance (DeFi) loans.




Figure 2: How AI may help connect off-chain data to a blockchain oracle.

AI-powered blockchain oracles may also be able to support standardisation of input data from external sources, leading to cost efficiencies. This has been explored in the context of issuers providing corporate actions data in a range of unstructured formats, which causes inefficiencies and risks. Inefficient corporate action processes cost regional investors, brokers and custodian businesses in the region of \$3 million to \$5 million annually, and 75% of firms revalidate custodian and exchange data manually, according to research by [TheValueExchange](#). To tackle this challenge, [Chainlink](#), [Euroclear](#), [Swift](#) and participating financial institutions participated in an initiative using AI models, oracles and blockchains to generate a unified central source of corporate action data, which is then synchronised among various blockchains.

Further development of use cases where AI enhances DLT and smart contracts might involve the use of AI agents first to negotiate the commercial terms of contracts and subsequently to write smart contracts, incorporating any agreed AI-enabled oracles.

A notable risk with AI is the potential for nonsensical or erroneous output. This means that the use of AI for dynamic smart contracts needs to be carefully reviewed. Cybersecurity is another key consideration. In the words of the Ethereum co-creator, [Vitalik Buterin](#), "it is important to be careful: if someone builds e.g., a prediction market or a stablecoin that uses an AI oracle, and it turns out that the oracle is attackable, that's a huge amount of money that could disappear in an instant."

One area being explored in relation to AI reliability and safety is the use of Byzantine Fault Tolerance (BFT) systems. This can help ensure that AI systems operate using redundant, cooperating modules to ensure data validation comes from many sources that can support a consensus algorithm and thereby help reduce AI hallucination. (See, for example, "[A Byzantine Fault Tolerance Approach towards AI Safety\[1\]](#)" by Dr. Matthias Artz and John deVadoss).

3.3

How might DLT improve AI training and development and enhance AI capability?

Use case #3: Using DLT to unlock data for AI training

Advancing AI development through controlled access to private datasets

Data marketplaces built on DLT could help make privately held data available for AI training in a controlled and traceable manner. These marketplaces would be decentralised platforms where participating data holders (who may be organisations or individuals) can allow access to datasets by consumers of data (such as AI developers and trainers) using blockchain infrastructure.

Such marketplaces typically involve representing data assets with tokens, which can be bought, sold or exchanged on a platform. These tokens grant the right to access and process or compute over the data under predefined conditions. The actual data would remain stored off-chain. Data sharing agreements are encoded into smart contracts that automate permissions, pricing and compliance obligations (such as, for example, consent withdrawal or usage limits). Each transaction (data access) is logged immutably on the DLT platform, assisting with traceability and accountability. This could support greater visibility and control for data holders when their data is used for AI training as well as facilitate some aspects of accountability in AI development.

The marketplace and underlying DLT network could be public or private and/or permissioned with access limited to specified participants. Such marketplaces could provide access to datasets themselves, helping to spur AI development simply through incentivising and facilitating data access. Alternatively, such transactions might involve allowing an algorithm to run on a dataset with the outputs or results of this training shared without sharing a copy of the underlying data.




Figure 3: Examples of different approaches to AI training using data marketplaces.

Such results could be shared in a bilateral transaction or multiple bilateral transactions with an AI developer in relation to a traditional, centralised model of AI development. In addition, an area of growing interest is the development of federated AI learning using DLT. This would mean that individual devices (nodes) participating in a decentralised network would together be responsible for training AI models, distributing computation and data storage. Each node would train the AI model on its local server and share only the results (e.g., the updated parameters of the AI model following the training), without transferring raw data, which is only processed locally on the relevant device. The DLT would record all of the local nodes' updates to the global AI model's code, helping with transparency and ensuring no single point of failure. A token-based incentive system could be used to help incentivise the sharing of results and ensure that protocols for sharing quality data and other safeguards are followed.




Figure 4: Example of potential federated learning for AI development.

Federated learning may excel in areas where access to the relevant raw datasets is tightly controlled, and where data is distributed and cannot, or should not, be centralised. To illustrate this capability, consider the healthcare sector. While medical data can be immensely valuable for advancing diagnostics, drug discovery and personalised medicine, it is highly sensitive and, in many jurisdictions, benefits from specific incremental legal protections. This sensitivity often leads to data silos across hospitals and research institutions and hinders AI model training. However, healthcare providers and research institutions may be able to leverage data marketplaces and participate in federated learning initiatives as data holders. AI developers can provide AI models to the healthcare providers (the data holders) who may then train models locally on these diverse datasets without the raw data leaving the secure local environment. This process could be used to ensure that only certain agreed outputs (e.g., insights or model parameters) are returned to the AI developers' global AI model. The associated legal and commercial challenges are likely to remain significant for both data holders (depending on factors, such as the particular data involved and agreed outputs) and AI developers (depending on the laws applicable to them, including, for example, the degree of diligence they are required to exercise over data used in AI model development). The construct also raises questions relating to allocation of responsibilities and liability. However, to the extent that such systems could be implemented in a manner that complies with relevant laws, they could help further unlock the collective intelligence of distributed datasets.

In addition to the legal considerations discussed in section 4 below, practical and commercial questions remain, including how to ensure fair compensation for valuable contributions without penalising protective steps that may be taken to preserve rights or protect privacy in these types of AI marketplaces. However, while the creation of these systems (particularly decentralised federated AI learning) may be technically and legally complex, potentially requiring policy interventions, research in this area is raising interesting possibilities for secure collaborative machine learning and use of AI. With today's increasing policy focus on the widening of certain forms of data access for the promotion of innovation, we may well see further investigation as to how to enable these types of use cases.

Use case #4: Financial services-enabled agentic AI

AI agents use blockchain wallets to support payment processes and agentic commerce

While there are varying definitions of 'agentic AI' and a range of underlying architectures, the term 'AI agent' is often used to refer to a system or program, for example packaged LLMs, that can, without human intervention, interact with its environment, collect data and use it to execute specific tasks or actions to meet predetermined goals based on a set of instructions. As AI improves in capability, AI agents are becoming increasingly autonomous and efficient at handling more complex tasks. However, there are limitations to the autonomy and convenience that agents can provide, including – currently – due to their limited financial capabilities, e.g., the ability to execute payments. An example of this is a customer-facing chatbot that can take simple details and answer basic questions but ultimately cannot provide support in real time when a payment-related function, such as processing a refund or making a purchase on behalf of an individual or a business, is needed.

Agents today generally rely on users (a human-in-the-loop) to complete transactions, which creates a disruption to the workflow and reduces efficiency. The use of traditional payment methods by AI agents poses challenges, as the use of credit or debit cards would require an AI model developer to store sensitive card information, and common banking security methods generally restrict AI agents from accessing users' bank accounts. Meanwhile, market expectations are being redefined, with a projected \$50 billion AI agent sector and the \$36.75 trillion digital payments industry ready to create a paradigm shift, according to the Payments Association.

A potential solution is an AI agent designed with its own unique blockchain wallet, accessible by the AI agent. Concurrently, a multi-signature smart contract wallet would be established, requiring the deposit of funds into the multi-signature wallet, and the approval of both an approving individual and the AI agent for any transaction to be executed (subject to any agreed exceptions).

Technology solutions are coming to market to support agentic AI payments. Some proofs of concept and emerging examples use digital assets, such as stablecoins, for payment between autonomous AI agents.

Threshold-based approvals would allow the configuration of different approval requirements based on transaction size. For example, this could be set to require both signatures for large transactions but allow smaller transactions to be approved automatically. While there might be potential for friction – humans may need to support some of the AI agent's actions in a chain of tasks – some actions may not require human engagement where predefined limits and parameters are built in.

Enabling agents to handle payments directly would enhance their functionality, making end-to-end task completion possible. Agents could handle tasks, such as purchasing supplies, paying bills, booking services and managing subscriptions. This increased level of autonomy further enhances agents' utility and efficiency in business operations for users.

This development provides the possibility of AI-to-AI transactions and the provision of better pay-as-you-go services. Agents could automatically pay small amounts for services, such as access to information, computational resources and specialised services from other AI agents.

An initial use case could be in the insurance industry, where a payment-enabled AI agent could help automate the claims process. Consider, for example, areas of insurance with high volume, but low overall claim value, such as windshield damage or lost luggage. Today, it is expensive to have special investigations teams working on these claims – an AI agent in this scenario could minimise the investigative costs, while blockchain wallets could support automated payment of successful claims.

Technology solutions are coming to market to support agentic AI payments. Some proofs of concept and emerging examples use digital assets, such as stablecoins, for payment between autonomous AI agents. Another example of innovation in DLT-based solutions is the use of multi-party computation wallets to support AI agents to send and receive payments autonomously using an open standard. At the same time, new solutions that use non-DLT approaches are also being explored. For example, some payment services providers are introducing tokens to allow AI agents to initiate secure transactions through conversational interfaces such as chat or voice.

Looking ahead, as digital identity frameworks mature and are integrated with agentic, payment-enabled AI, there is an opportunity for dynamic adjustment of security, access and personalisation based on situational factors. A futuristic fully autonomous agent could use a person's digital identity and preferences to accomplish a range of tasks. For example, if an individual would like to book a weekend in Paris, the AI agent could theoretically work across flight, restaurant, tour and accommodation comparison sites to optimise package options using the individual's profile for those sites. If the person has set parameters on cost, a payment-enabled AI agent could then purchase the products on their behalf – even without human intervention, if desired. The agent would then present all relevant details to the human user in a preferred format. Ultimately, we may even see humans manage teams of AI agents to ensure their efficacy. All of this is important in the context that 61% of consumers are willing to spend more with companies that offer customised experiences, according to [research by Medallia](#).

4

Legal and regulatory landscape

Organisations exploring use cases at the intersection of AI and DLT will need to consider and navigate a complex legal and regulatory landscape, including a wide array of existing laws (such as privacy, cyber, intellectual property, antitrust, consumer protection, disability and employment laws, as well as sector-specific and technology-targeting legislation) in addition to rapidly evolving technology-specific and data laws, contract and private law considerations alongside sector-specific regulations. In this section, we provide examples of some of the key legal issues, spotlighting some of the relevant legal regimes and considerations. We also highlight some of the additional sectoral requirements that may apply to financial services sector deployments of AIxDLT. Specialist legal considerations will also apply to other regulated sectors, including transportation, energy, healthcare and life sciences.

For some of the more innovative AIxDLT use cases, new legislation and/or discussion and collaboration with regulators is likely to be needed. A fundamental issue to be considered is the application of existing, or the creation of new, frameworks for responsibility, liability and accountability for AI agents operating on decentralised models.

4.1

AI laws and guidance

The global policy and regulatory landscapes are evolving rapidly in response to AI's growing capabilities and societal impact. Over 1,000 AI-related policy initiatives have reportedly been introduced in over 70 jurisdictions, according to the Organisation for Economic Cooperation and Development (OECD), reflecting a shared recognition of AI's transformative potential.

A range of existing laws and regulations already apply to the development and use of AI, including privacy, cyber, intellectual property, antitrust, consumer protection and employment laws, as well as sector-specific and technology-targeting legislation. Some countries, such as the UK and Singapore, have currently opted for decentralised and, often, sector-led approaches with a focus on engagement of existing regulators and production of guidance as to how existing laws and requirements apply to AI. However, increasingly, many countries have begun passing laws and regulations that specifically focus on regulating AI.

China was among the countries rolling out some of the earliest AI-specific rules targeting specific types of AI use, including the Provisional Administrative Measures for Generative Artificial Intelligence Services, rules on 'deep synthesis' or deepfakes (the Administrative Provisions on Applying Deep Synthesis Technology in Provision of Internet Information Service) and algorithmic recommendations (the Administrative Provisions on Applying Algorithm Technologies in Provision of Internet Information Service). Other jurisdictions enacting AI-specific legislation include the EU, South Korea and certain US states. The entry into force of the EU's AI Act was a significant milestone for AI legislation globally. This comprehensive AI law has global reach, impacting entities based outside the EU in certain circumstances (such as where an AI system is put into service in the EU or placed on the EU market, or where the AI system's output is used in the EU), in addition to EU-based entities. It imposes obligations on operators across the entire AI value chain, and the AI Act's rules and compliance requirements are defined following a tiered approach dependent on the AI risk. These laws often have strong enforcement regimes – for example, penalties under the AI Act can reach up to 7% of an undertaking's global annual turnover.

In addition, a range of regulators globally, including sectoral regulators, data protection authorities and antitrust regulators, have issued AI-specific regulatory guidance, helping clarify regulatory views on the application of existing legislation and regulatory requirements to AI.

Despite a vast divergence in the approach and scope of these laws and frameworks, global trends in AI-specific legislation include:

- **Extraterritoriality:**

A country or region's AI regulation often has extraterritorial effect. In many cases, factors such as the market into which an AI system is placed, where an AI user or deployer is based and/or where the AI output is used can dictate which laws apply. Before developing or deploying AI in any scenario, it is important to map the applicability of AI-specific legislation alongside other laws, including potential overlaps with sectoral legislation (for example, financial regulation as discussed in section 4.5 below). This may not always be straightforward, particularly in the early stages of AI development (which may be the case for use case #3 (using DLT to unlock data for AI training), for example).

- **Risk-based regulation:**

Some jurisdictions are adopting frameworks that classify AI systems by risk level, with stricter obligations for high-risk applications, or which apply only to certain activities identified as higher-risk (e.g., the making of employment or credit-related decisions or other decisions with significant impact for individuals). Therefore, applicable requirements and restrictions vary not only according to geographic footprint, but also the specific context in which AI is deployed.

For example, under the EU AI Act, some AI practices would be prohibited (e.g., subliminal / manipulative techniques or exploiting vulnerabilities and social scoring), and AI that falls within the definition of a 'high-risk' AI system under the EU AI Act will be subject to additional stricter requirements, subject to certain exceptions. High-risk AI systems include, for example, particular uses of AI in certain HR or recruitment processes or in the context of credit assessments, or where an AI system is used as a safety component in the management and operation of critical digital infrastructure. The relevant requirements differ according to an organisation's role. For providers of high-risk AI systems, for example, requirements relate to risk and quality management; data governance (including obligations to prevent bias); record-keeping; transparency; registration; human oversight; accuracy, robustness and cybersecurity; and conformity assessments. Developers of general-purpose AI models – AI models displaying significant generality, which can perform a wide variety of tasks across multiple contexts and be integrated into a variety of downstream systems or applications – are also subject to specific obligations, including transparency, risk management and copyright compliance. Additional safety and security requirements apply to providers of general-purpose AI models that present high-impact capabilities and qualify as 'models with systemic risks'.

In relation to the use cases discussed in this paper, when overlaying applicable laws, particular consideration should be given to the context in which the use case arises. For example, while some instances of use case #3 (e.g., tokenised data marketplaces) may help incentivise greater data access for AI model development, they might also present challenges in relation to, e.g., requirements under the EU AI Act relating to diligence, validation and testing of datasets used in high-risk AI systems where the datasets are not accessible for quality control and other assessments, including, e.g., to assess suitability and detect and mitigate possible biases or quality issues.

While some forms of tokenised data marketplaces may help incentivise greater data access for AI model development, versions of this use case that do not make underlying datasets accessible for quality control and other assessments might present challenges under laws that require diligence, validation and testing of datasets used in certain AI systems.

- **Roles and responsibilities:**

The requirements and restrictions applicable to an organisation, and potential associated liability, will typically differ according to its role in the 'AI value chain'. The EU AI Act, the Provisional Administrative Measures for Generative Artificial Intelligence Services applicable in China and South Korea's AI Act place responsibilities on players across the value chain. In the US, few state laws do so comprehensively, and most legislation focuses, instead, on specific use cases, roles and/or harms. The concept of decentralised AI development raises questions regarding allocation of responsibility and liability under AI laws that are primarily aimed at specific AI developers or providers.

- **AI literacy:**

Ensuring that persons using AI are equipped with the requisite knowledge and skills is a goal under various regimes, although the regimes' targets vary. For example, under the EU AI Act, a key focus is on the providers and deployers of AI systems, to ensure the literacy of their staff and others dealing with the operation and use of AI systems on their behalf. In the US, President Trump's AI Action Plan and a wide range of state laws call for various versions of government education on and use of AI, articulating various goals and calling for the creation of committees, commencement of studies, etc. The significance of these requirements is evident for use cases #1 (smart contract development) and #2 (AI oracles), where a lack of awareness of AI's limitations and risks – for example, in relation to errors and hallucinations – could lead to over-reliance on such technology without appropriate checks and safeguards.

- **Transparency and accountability:**

Arguably building on some of the early work of the OECD, there is a growing emphasis on explainability, human oversight, and documentation of AI decision-making processes. Many jurisdictions also focus on transparency in the development and/or use of AI, e.g., a need to inform users that they are interacting with AI or that AI is involved in certain decision-making, the labelling of AI-generated content and/or not overstating the role or abilities of AI in an organisation's business or product ('AI washing').

Where DLT can assist with understanding provenance of data used in AI training (as may, for example, be the case in certain versions of use case #3 (using DLT to unlock data for AI training), depending on how they are implemented), such use could help support transparency and accountability in AI development due to DLT's quasi-immutable nature. In theory, this can enable reliable recording of steps along the AI training pipeline from data ingestion for AI development to fine-tuning and updates. Such practice may support compliance with transparency obligations under certain laws, such as the EU AI Act, which requires clear documentation and user notification of certain AI systems. The practice may well have applications beyond adherence to AI-specific laws or principles – for example, it may be relevant to issues regarding intellectual property rights in data and materials used for AI training.

For use cases involving agentic AI, particularly where AI agents interact with a person (e.g., an AI chatbot), transparency obligations may include ensuring that the person is made aware that they are interacting with AI.

- **Fairness and bias-mitigation:**

Various AI-specific laws contain provisions relating to preventing or mitigating unfair bias in AI use and outputs. For example, the EU AI Act includes requirements to subject the training, validation and testing datasets of high-risk AI systems to examination in view of possible biases as well as appropriate measures to detect, prevent and mitigate biases that may have been identified. AI deployers have certain obligations related to the relevance of input data, where they have control over it. Similarly, the Colorado AI Act, the first comprehensive US state AI law, focuses on preventing algorithmic discrimination of consumers, and introduces a standard of 'reasonable care' for developers and deployers of AI together with distinct obligations.

Blockchain-based data marketplaces may allow diverse contributors to share datasets under transparent terms. This could broaden the range of data available for AI training, including potentially encouraging the inclusion of underrepresented groups. Frameworks that allow validators to review DLT-logged data flows could potentially enable the flagging of biased sources or help assess the overall suitability of training sets. However, depending on how such systems are implemented, they could also create challenges if AI developers cannot appropriately diligence datasets, or in circumstances where dispersed roles undermine accountability.

4.2

DLT and cryptoasset-specific laws and guidance

Until recently, in the EU there were limited regulatory regimes directly impacting the use of DLT and cryptoassets beyond anti-money laundering (AML) related requirements that implemented the international Financial Action Task Force (FATF) recommendations. However, that changed in 2024, when the EU's Markets in Cryptoassets Regulation (MiCA) entered into force, creating an EU regulatory framework for the issuance of, and intermediating and transacting in, cryptoassets, alongside a market abuse regime for cryptoasset markets. Similar regimes are taking shape across key financial centres. A UK regulatory regime that has a similar scope to MiCA is being developed, with HM Treasury issuing a draft statutory instrument to extend the regulatory perimeter to include certain cryptoasset-specific activities. Additionally, the UK's Financial Conduct Authority (FCA) has issued relevant consultation papers on specific aspects of the future regulatory regime applicable to cryptoassets, which is expected to take effect during 2026. The US is also actively developing a legislative and regulatory framework related to cryptoassets, including the recently enacted Guiding and Establishing National Innovation for U.S. Stablecoins Act (GENIUS Act) with respect to payment stablecoins, broader crypto market structure legislation, such as the Digital Asset Market Clarity Act of 2025 (CLARITY Act) moving through Congress, as well as a liberalisation of the treatment of crypto-related assets and services by prominent securities and banking regulatory bodies, including the US Securities and Exchange Commission (US SEC), the Federal Reserve, and the Office of the Comptroller of the Currency. Alongside emerging regulation, Courts are being asked to opine on novel private law legal issues associated with DLT and cryptoassets, such as jurisdictional and governing law matters, and proprietary rights and remedies.

MiCA asset scope

MiCA applies with respect to "crypto-assets", which are defined very broadly to include cryptoassets without any backing or reference assets, such as Bitcoin and Eth, but also stablecoins.

MiCA activity scope

MiCA broadly applies to persons (i) offering cryptoassets to the public and/or (ii) seeking to admit cryptoassets to trading on a cryptoasset trading venue in the EU, as well as those undertaking specific cryptoasset services (cryptoasset service providers or CASPs) within the EU. A cryptoasset market abuse regime applies where cryptoassets are admitted to trading on an authorised platform.

Application of MiCA to AlxDLT

In line with other financial services regulatory frameworks, MiCA does not regulate the investment asset or the technological infrastructure generally, but it defines certain regulated activities, the performance of which carries regulatory consequences (authorisation, prohibitions, penalties, etc.). As such, MiCA will only apply to AlxDLT use cases to the extent that they result in an offer of cryptoassets to the public, a person seeking to admit cryptoassets to trading, where a service provider is carrying on one of the regulated activities or where market abuse provisions are triggered. For the use cases described in this paper, use case #3 (using DLT to unlock data for AI training) is most likely to fall within the scope of MiCA, where the tokenisation of data creates a cryptoasset and a DLT-based 'data marketplace' or trading platform is set up to trade such data/cryptoassets. However, as with all deployments of new technology, it is critical that an analysis of any use case incorporates suitable case-specific legal advice.

4.3

Privacy and data protection

Data protection and privacy regimes around the world are diverse and evolving. Application of these laws to each of AI and DLT is the subject of considerable scrutiny and debate, with various areas of untested law. In this section, we set out a high-level introductory overview and a few examples of key considerations. With many international organisations using the GDPR as a data protection 'high bar' or base for default compliance approaches, we have used this data protection law as an example for many of the issues discussed in this section. In some cases, the points raised may also apply to other data protection regimes, but we anticipate that significant nuance and complexity will be encountered as these issues are further explored at a more granular level and with more specific consideration of other data protection frameworks.

DLT and data protection

Application of data protection laws to blockchain technology has been a point of regulatory and academic debate. Points of tension largely relate to two core features on DLT – its decentralised nature and its 'immutable' architecture. Challenges can include:

- assigning accountability in a decentralised public network, including identification of data controllers, data processors and equivalent roles under privacy laws, and putting in place agreements and protective measures required by many data protection laws (including, for example, for certain international data transfers); and
- application of certain common data protection requirements to a distributed ledger's 'immutable' architecture; for example, how to adhere to GDPR principles of time-limited storage of personal data and data accuracy, and how to give effect to rights to data rectification and erasure.

To the extent that any of the use cases involve processing personal data through blockchain, some of the key points for consideration include:

- **Allocating data protection responsibilities:**

Data protection laws typically allocate differing responsibilities according to the nature of the role that a person or an entity plays using personal data, often depending on the level of control that person or entity has over how and why the data is processed. The premise that for each data point or dataset there exists at least one data controller responsible for ensuring compliance with data protection laws can conflict with the principle of decentralisation, making it challenging in certain circumstances to assign responsibilities, such as ensuring that certain contracts are in place between the controller and those processing the data, controlling international transfer of the data, and giving effect to data subject rights. This can be particularly challenging where public permissionless blockchain networks are used. There are incremental questions and challenges for use cases involving federated learning.

Points of tension between data protection laws and blockchain technology largely relate to two core features of DLT – its decentralised nature and its 'immutable' architecture.

- **Deletion and correction of data:**

Under many privacy laws, including the GDPR, individuals have the right to have their personal data made accurate (if it is incorrect) and permanently erased (in certain circumstances). There are also typically broader requirements, for example to ensure that data is not stored for longer than it needs to be for the specific purposes for which it is processed. Blockchain's resistance to data modification complicates compliance with these requirements.

Several data protection authorities have identified issues in respect of blockchain. Recently, the European Data Protection Board (EDPB) published [draft guidelines on processing personal data through blockchain](#), which consider that data stored directly on-chain should not allow the direct identification of the data subject (i.e. the individual to whom the data relates) and that, instead, blockchain participants need to consider how to make it possible for any data on the blockchain to be rendered anonymous once it is no longer needed (or, potentially, to address rights to data erasure). The EDPB suggests that this could involve ensuring that data linked to an identifiable individual or which can be used to 'single out' an individual is stored off-chain in a format that is capable of erasure, with a link or reference to the data stored on-chain. This means that it may be possible to prevent the future identification of a data subject through erasure of off-chain data, depending on the exact method chosen and the specific facts in question. The broad definition of 'personal data' under laws, such as the GDPR (including data that allows for indirect identification of an individual) makes this challenging, and record-keeping obligations under other laws can add further complexity for certain blockchain participants.

The broad definition of 'personal data' under laws such as the GDPR (and equivalent concepts under other laws) is a crucial point to consider in all cases, including when considering data marketplaces. Where data can be used to identify or single out an individual using reasonable means, it will remain personal data even if the data does not directly identify an individual. The application of privacy laws and definitions in the context of AI (e.g., in relation to training data, model parameters and output data) is an area of significant regulatory discussion.

- **International data transfer:**

The cross-border nature of DLT typically involves cross-border data flows. However, many privacy and other laws have rules regarding how data can be shared across international borders. This, for example, can include requirements for documented assessments and the putting in place of specific contractual terms between data exporter and data recipient or, under some laws and for certain data, may require approvals from authorities.

- **Automated decision-making:**

The execution of a smart contract may, in some cases, constitute an automated decision, which can be subject to specific requirements and restrictions under data protection laws. For example, where the EU GDPR applies, there are strict rules for decisions made on a solely automated basis (without meaningful human involvement) and which have a legal or similarly significant impact on an individual. An organisation that is deemed a data controller would need to ensure that the processing is only carried out for the limited permitted purposes ('legal bases') set out in the GDPR, and that certain safeguards are satisfied, including the possibility of human intervention, and allow the data subject to contest the decision, even if the smart contract has already been performed and regardless of what is registered on the blockchain.

AI and data protection

Questions regarding compliance with data protection requirements in relation to AI development and use have been an area of focus for data protection authorities across the world, resulting in various guidance, frameworks and enforcement activities. Considerations include: (1) the appropriate legal basis for processing personal data to train AI, particularly where data is repurposed or scraped from the internet; (2) giving effect to individuals' rights under data privacy laws, which may include rights to correction or deletion of data, rights to certain information regarding its processing and in relation to certain types of automated decisions; (3) compliance with data transfer requirements for cross-border transfers; and (4) adherence with privacy principles, such as fairness and transparency, accuracy, confidentiality and the time-limited storage of data.

Some of the requirements that organisations may need to consider in connection with these use cases include:

- **Legal basis for processing:**

A data controller (i.e. the person or entity that determines how and why personal data is being processed) is typically required to ensure that the processing of personal data for the development of or use by AI is covered by a purpose or condition ('legal basis') set out in applicable data protection legislation. Under many privacy laws, a key legal basis is the consent of the data subject (the person to whom the data relates). However, there is often a high bar for establishing that valid consent has been given, including the information provided on how the data will be used and by whom. Consent is also usually revocable at any time by the data subject, raising questions as to how this can be achieved in practice where an AI model has 'learned' that information.

Under the EU GDPR, several other legal bases are available, including the basis that the processing is necessary in the legitimate interests of the controller or a third party. However, this requires a factual assessment to establish that the rights and freedoms of the data subject do not, in the particular circumstances, override those legitimate interests, and such assessments can be challenged by data subjects.

- **Transparency and data subject access:**

Controllers are required to provide data subjects with certain information regarding the processing of their data. Under the GDPR, this includes information about any automated decision-making and profiling, as well as the logic involved, method of processing and its effects for data subjects. Rights to access copies of personal data also exist under the GDPR and other laws. Depending on how data markets discussed in use case #3 are implemented, the visibility of data provenance that this could afford may help with meeting requirements regarding transparency and the provision of information.

- **Data correction and erasure:**

As discussed above in relation to data protection and blockchain, many privacy laws include requirements regarding the correction and deletion of personal data. In the context of AI development, compliance with these requirements can be challenging, raising questions regarding what it means to give effect to these rights in respect of an AI system that has already 'learned' from the data and whether 'unlearning' can be achieved. In relation to federated AI model development using DLT, these questions become more complex legally and technically.

Data security and confidentiality:

Data protection laws and cybersecurity laws require, among other things, that appropriate safeguards are in place to ensure the security of personal data. Additional requirements can apply under cybersecurity and operational resilience laws, including in relation to non-personal data, with the nature and applicability of such requirements often differing based on context, such as the nature of an organisation and the sector within which it operates. Cybersecurity risks need particular consideration in the context of AI systems. For example, in relation to protection of personal data, various kinds of external attacks on AI systems (in particular, membership inference and model inversion attacks) might lead to the possibility of using their models to reproduce some or all of the training data used in their development, or at least to identify the individuals on whose data the systems were trained. This is a crucial consideration for use case #3 (using DLT to unlock data for AI training) in relation to decisions regarding allowing the training of AI on locally held private data sets, even where the raw data is not shared. More broadly, it is also crucial to consider the degree to which an AI system is relied upon for a critical process or might otherwise carry particular risks should it become compromised or malfunction, and what checks and safeguards would be in place to prevent, detect and mitigate the effects of any cybersecurity issue. This is a key consideration for AI agents as their degree of autonomy grows (use case #4), and also relevant to use cases #1 (smart contract development) and #2 (smart oracles), where consideration needs to be given to the appropriate degree of reliance on AI and how malfunctions, hallucinations, errors and other undesirable operations would be detected and addressed.

In the context of AI development, compliance with data correction and deletion requirements can be challenging, raising questions regarding what it means to give effect to these rights in respect of an AI system that has already ‘learned’ from the data and whether ‘unlearning’ can be achieved.

4.4

Liability considerations

Liability for AI use

Liability for the actions of AI systems is a complex and emerging area of law, with a range of legislative and judicial approaches across jurisdictions. While specific AI liability frameworks are being considered and developed in some jurisdictions, it is important to assess potential liability considerations under a wide range of existing legal frameworks, including intellectual property, product liability, consumer protection, employment and disability laws. For example, while the EU has recently withdrawn its planned AI liability directive, several other regimes relating to liability remain, including product liability laws, such as the recently revised EU Product Liability Directive, which includes updates made with AI in mind. There is also an ongoing debate in England as to whether common law principles will need to be adapted to accommodate the increasing use of AI.

Notwithstanding the variation across jurisdictions, in general AI is treated in the same way as other kinds of technology and software and does not have a separate legal personality; rather, it is a tool. Alongside liability under statutory regimes (such as the product liability laws referenced above), providers and/or deployers of AI will have some measure of responsibility for the actions of an AI system under other legal frameworks, including contract and tort. The fact that AI is increasingly 'agentive' and independent is unlikely to affect this in the short term – it would generally be wrong to assume that an AI 'agent' has any independent status that would free its developer or deployer from liability (unless and until AI is given separate legal personality in any relevant jurisdiction). AI developers and deployers should therefore take steps to identify and mitigate liability risks, including practical steps, such as controls on AI output and actions, and sensible legal steps, such as agreeing contractual terms to clarify the allocation of liability between the parties. Contracts are an incredibly useful mechanism for allocating liability risks and responsibilities as well as ensuring assistance from a counterparty in complying with legal obligations (often by way of flow-down). However, it is important to remember that contracts cannot always override or reverse responsibility under mandatory laws or applicable regulation. Therefore, it is crucial to ensure AI safety 'by design', i.e. to consider potential risks and set out appropriate safeguards and guardrails at the earliest stages of a project, with regular monitoring and auditing of all AI models and tools.

Contractual interpretation (smart contracts)

While autonomous AI systems have contracted with one another for years in the context of algorithmic trading, this typically takes place on an exchange where AI systems exchange a limited range of fixed electronic instructions defined by the exchange's trading systems. By contrast, smart contracts are written by users with bespoke code. This allows greater flexibility in how they operate but also enables greater scope for disputes over whether the code properly reflects and executes the parties' agreement.

It is helpful to distinguish between situations where a smart contract complements a conventional natural language contract, for example, where the smart contract is used for execution purposes, and code-only smart contracts with no natural language counterpart. AI could be used to draft smart contracts in both scenarios.

Where parties are dealing with one another under the framework of a natural language contract, there is potential for disputes if a party claims that the result of the smart contract and the terms of the natural language counterpart are not consistent (or for example in the case of secondary transactions where the natural language contract is not available to a subsequent party to the on-chain contract). There may be complex issues as to how their terms should interact and which terms should prevail, which may require an investigation into how the smart contract was created. Different considerations may arise where one party has relied on another to create the smart contract, or where a third party has been used to do so. The use of AI does not fundamentally change this analysis, but may give rise to novel evidential issues, such as the relevance of the prompts given to the AI.

Parties in such situations can achieve greater certainty by agreeing express terms that specify the hierarchy between the smart contract and the natural language contract and setting out a mechanism for determining whether there is any error in the smart contract and how it should be addressed.

It might be assumed that such conflicts could be avoided by using code-only smart contracts. While dealing on a code-only basis is likely to be attractive to many in the industry, it may give rise to significant legal uncertainty in a dispute. Counterparties may seek to rely on the wider commercial negotiations to argue that the smart contract was mistaken or defective in some way, or that it was qualified by side-agreements. Natural language contracts typically contain express provisions that they represent the entirety of the parties' agreement, but code-only smart contracts are unlikely to contain such provisions. The use of AI to generate a smart contract would not avoid this risk.

This may be anathema to some smart contract users. The phrase 'code is law' exemplifies the view of some market participants that in contracts executed through code, the code represents the entirety of the parties' agreement. However, parties that wish to contract on the basis that 'code is law' may be best served by entering into a natural language contract to that effect, for example by expressly agreeing that any smart contracts between them are final and binding and represent the entirety of the parties' agreements, and that the parties warrant that they understand their terms and effect. It is also critical to consider and specify the governing law and dispute resolution mechanisms.

Challenges to smart contract

Bugs can lead to a smart contract malfunctioning, for example by locking users' funds in certain circumstances. Vulnerabilities can be exploited by hackers, as has occurred on several cryptocurrency exchanges. This can mean that smart contracts operate in ways that were not intended by their users.

In many jurisdictions, contracts can be challenged and transactions reversed (or rescinded) on several grounds; for example, if they were based on mistake, misrepresentation, or under duress. However, smart contracts typically by design do not allow transactions to be reversed and attempts to do so are highly controversial among crypto users, even in cases of large-scale hacks. Further, because it is often impossible to identify other blockchain users, it can be difficult to identify the counterparty against whom a claim should be brought, unless the parties are already known to each other. Courts are being asked to enforce parties' rights and remedies in such circumstances, meaning that the evolving jurisprudence is being watched with interest by the market.

Where AI is used to generate code-only smart contracts, particularly with anonymous counterparties, particular care is needed to ensure that the smart contracts are robust, free of vulnerabilities and consistent with the users' intentions. Where possible, parties would be well-advised to put in place contractual provisions setting out their rights and obligations in case of a smart contract malfunction and mechanisms for resolving any disputes about how this should be remediated.

4.5

Financial services regulatory overlay

For regulated firms in many jurisdictions, additional requirements will apply under financial services frameworks. The increasing use of new technologies such as AI and DLT in financial services has led to heightened regulatory scrutiny internationally. While regulators in some jurisdictions may impose rules that are specific to a novel technology, in others, financial regulation will be designed to be technology-neutral, based instead on broader principles and/or outcomes. This may be supplemented by regulatory guidance that confirms how existing rules will apply to new technologies, for example. Some key themes can be identified that raise important principles and questions for regulated firms to consider when deploying AIxDLT use cases. In some cases and jurisdictions, there will be overlap with new requirements under AI-specific laws, e.g., requirements for AI to meet fairness and transparency requirements.

Strong governance and clear accountability

Boards and senior management will be expected by regulators to oversee the deployment of AIxDLT use cases within their firms, with clear lines of accountability and strong governance.

For example, in the UK, the Financial Conduct Authority (FCA) has emphasised that firms remain fully accountable for the outcomes of their use of AI and stressed the importance of ensuring that any use of AI by regulated firms complies with existing "high-level overarching requirements", including the Principles for Business, in addition to "more specific rules and guidance relating to systems and controls under the Senior Management Arrangements, Systems and Controls (SYSC) sourcebook". Relevant provisions include Principles for Business 2 and 3, which require firms to conduct their business with due skill, care and diligence and organise and control their affairs responsibly and effectively, with adequate risk management systems, respectively. SYSC requirements relevant to governance of new technologies include organisational and governance requirements under SYSC 4 and SYSC 7, which require management bodies to approve and review policies for managing, monitoring and mitigating relevant risks.

The European Securities and Markets Authority (ESMA) issued a [public statement](#) concerning the use of AI in retail investment services in May 2024, which emphasised that any integration of AI in retail investment services will need to comply with the existing rules under MiFID II, in addition to the broader framework of the EU AI Act. The role of management bodies is singled out as crucial to this process, in particular the introduction of "robust governance structures" and fostering a "culture of risk ownership, transparency and accountability".

Adequate testing of systems and controls

Financial services firms will generally be required by regulators to ensure that their risk management frameworks are effective in covering any AI and DLT-specific risks. These frameworks must be fit for purpose and will need to subject the firms' systems and controls to sufficiently rigorous testing. Broadly, we have seen regulators worldwide increasing their involvement in the testing of AI risk frameworks.

For example, in April 2025, the FCA published a proposal for [AI Live Testing](#). This would involve the launch of a live testing environment for AI systems used in UK financial markets products that have completed the proof-of-concept stage and are mature enough to be rolled out. The FCA would collaborate in real time with firms to explore methods of evaluating the impact of AI on customers, including output-driven validation. In June 2025, the FCA also announced its plan to launch a [Supercharged Sandbox](#) alongside NVIDIA to help firms in the discovery and experiment stage explore AI safely.

In many jurisdictions, there are also specific cybersecurity and operational resilience laws and regulations that apply to financial services firms or more broadly across sectors that may impact deployment of AIxDLT projects and related contracts. For example, the EU's Digital Operational Resilience Act (DORA). While these requirements often apply broadly to rollouts of new technology and could impact projects using only AI or DLT, some aspects of compliance are likely to become more challenging due to the combination of AI and DLT, for example the allocation of responsibilities for decentralised systems, including for decentralised federated learning for AI and potentially, increased points of vulnerability for agentic AI.

In some cases, guidance and standards developed by wider agencies will also be relevant for financial services firms. In the US, the National Institute of Standards and Technology's (NIST) voluntary [AI risk management framework](#) was developed through a public consultation process and is intended to "incorporate trustworthiness considerations into the...evaluation of AI products, services and systems".

Financial regulators will expect boards and senior management to oversee the deployment of AIxDLT use cases within their firms, with clear lines of accountability and strong governance.

Fairness and anti-discrimination

Regulators are also concerned about the potential for AI making discriminatory decisions that result in unfair customer outcomes. There is likely to be a particular focus in the context of the use of agentic AI by regulated firms and the use of AI-powered blockchain oracles that could inadvertently embed bias.

In the UK, the FCA has emphasised that "AI systems should not undermine the legal rights of individuals or organisations, discriminate unfairly against individuals or create unfair market outcomes". In addition, the FCA's Consumer Duty requires firms to, among other things, act in good faith, avoid foreseeable harm and support retail customers' financial objectives in relation to any deployment of new technology, such as AI and/or DLT.

In the EU, ESMA has referenced algorithmic biases as one of AI's key inherent risks, noting that they can lead to unfair outcomes for customers as a result of "AI systems... inadvertently perpetuat[ing], amplify[ing], or introduc[ing] discrimination against certain individuals or groups". The European Commission has also stressed the importance of "the protection of EU values and fundamental rights such as non-discrimination".

Given that this is a focus area for regulators, firms will need to ensure that any implementation of AI products is accompanied by sufficient discrimination safeguards effectively built into governance frameworks.

Transparency of processes

Regulators in many jurisdictions expect that firms will proactively provide information on any AI products or systems they use and will be able to explain how these products make decisions; this especially applies to areas affecting customers.

In the UK, for example, while FCA rules do not mandate specific transparency requirements for AI or DLT use, broader high-level consumer protection obligations will apply. These include the cross-cutting obligation under the Consumer Duty to act in good faith (including being honest, fair and dealing openly with consumers) and a broader requirement for firms to communicate in a way that is clear, fair and not misleading. Firms should be able to explain how AI systems work and how decisions are made, in a way that is appropriate for the intended audience, as well as maintaining appropriate documentation and audit trails.

Firms will therefore need to consider their compliance with relevant transparency obligations, including the implementation of customer-facing explanations of their use of AI in AlxDLT use cases.

As outlined above, in some AlxDLT use cases, the use of DLT may assist with compliance with transparency obligations.

Contestability and human oversight

Regulators have made clear that there must be sufficient human oversight of AI systems used by regulated firms with the ability to challenge decisions made by AI.

In the UK, the FCA has emphasised that firms must give individuals the right to contest decisions made by AI systems and that appropriate human oversight is maintained for any use of AI through the product life cycle.

Firms will need to consider their existing oversight frameworks in advance of AI integration, and whether these need to be expanded or strengthened to allow for human intervention into automated processes, as well as considering customer complaints and redress processes.

Risk-based and proportionate approach

In some jurisdictions, different levels of safeguards will apply depending on the risk of a product, its intended use, and possible outcomes. For example, use cases of AI and DLT that may directly impact retail customers are likely to be subject to higher regulatory expectations and requirements. This is consistent with the approach taken in EU law under the EU AI Act, as outlined above, with tiered compliance obligations based on four levels of risk.

In the UK, the FCA has emphasised that "the principle of proportionality...informs [the FCA's] thinking and approach to AI, including any potential future regulatory interventions".

When developing any new AlxDLT use cases, firms will need to consider carefully the level of risk associated and ensure that their governance approach is proportionate and defensible. Internal risk assessments of use cases will be critical for firms to supplement compliance with baseline regulatory requirements and oversight. Firms should also evaluate the extent to which any such internal assessments are aligned with the approach of relevant regulators and whether the systems and metrics used enable a sufficiently thorough examination.

Market abuse

The increasing use of AI in financial markets as a tool to analyse data, identify patterns and rapidly adjust pricing and trading strategies poses several potential risks in the context of competition law and market abuse frameworks.

Tacit collusion of AI systems

Anti-competitive and price-fixing agreements are widely prohibited under market abuse legislation, including in the EU, the UK and the US. However, where AI systems are employed by separate companies to set pricing and trading strategies (such as in algorithmic trading), there is a risk that such systems may learn to co-ordinate their adjustments with competitors as a way to maximise profit. In the UK, the Competition and Markets Authority (CMA) has warned that the use of complex pricing algorithms could still lead to breaches of competition law where the result is effectively collusion, even without an explicit price-fixing agreement. This might be seen where two competitor firms use the same piece of software that results in confidential information, such as in relation to pricing or marketing, being shared. This would need particular attention for federated learning AI models and the use of AI-powered blockchain oracles. The US SEC recently hosted a roundtable discussion focused on the risks, benefits and governance of AI in the financial services industry.

Market manipulation by AI systems

There is also a risk that AI systems may be used to facilitate market manipulation. This could take the form of deliberate manipulation by using AI systems for the placement or cancellation of large numbers of orders to influence prices, unintentional manipulation through AI systems adopting abusive practices at their own initiative through iterative learning (where effective guardrails have not been imposed by deployers), or with firms inadvertently publishing misleading information as a result of AI-generated content. In Europe, the Markets in Financial Instruments Directive (MiFID) framework has specific rules for algorithmic trading, including that algorithms must not be used for purposes contrary to the Market Abuse Regulation. As well as considering their own use of technology, firms need to have systems in place to detect and prevent such forms of market misconduct by their clients. It will be particularly critical when developing or deploying AI agents supported by DLT and AI-powered blockchain oracles that effective control mechanisms are built in to avoid inadvertent market manipulation.

5

AlxDLT implementation: key considerations

As we explore the potential capabilities of, and opportunities arising from, AlxDLT use cases, it is crucial to acknowledge that any path forward is not without its challenges. Regulatory frameworks are evolving rapidly to keep pace with technological advancements. Legal requirements, ethical considerations and risk implications regarding certain uses of AI and DLT will need to be carefully navigated. Engagement across industry, with policymakers and with regulatory stakeholders will be critical for successful exploration of innovative use cases. Moreover, industries will need to invest significantly in infrastructure and talent to realise fully the potential of these technologies.

Whether considering one of the specific use cases detailed in this paper or another application of AlxDLT, it is important for organisations to consider carefully how the development of these technologies may impact them and their daily operations. Ensuring that the positive benefits can be fully achieved means working partnerships internally and externally. The market is shifting and each exploration or implementation of AlxDLT will differ, so we set out some key practical insights below.

○ Take a holistic and collaborative approach to use case exploration and implementation:

While AlxDLT proposals might be led by one area, it is crucial to understand what a use case or development might mean for your business, organisation or group as a whole and for structures to exist that allow collaboration and oversight across teams to prevent resistance, manage risk and take advantage of opportunities. Many areas of your organisation – legal, technology, risk and compliance, sales, HR, finance and tax – will have to come together to consider the potential of AlxDLT use cases for new and existing business models and the impact on the business or organisation overall.

- Build internal capabilities and literacy to support the changes ahead:

To take full advantage of the opportunities ahead, many areas of your organisation will need to understand the basics of this technology shift and how it aligns to the work they do on a daily basis. It is crucial to prioritise AI and technology literacy so that staff and stakeholders understand a system's capabilities and limitations (including, for example, the risk of errors) to ensure that use is appropriate and that there are checks and guardrails. Creating the right teams and processes to implement a comprehensive testing strategy is an important part of enabling AlxDLT development and use. The hiring of new and training of existing staff and support from vendors and partners are also likely to be crucial parts of the journey, alongside considering potential employment implications of AlxDLT use cases.

Testing AlxDLT systems

A strong, end-to-end testing strategy is essential to make sure AlxDLT systems work safely, consistently, and in line with legal and business expectations. Testing must be built into every stage – from design to operations – to ensure compliance is evidence-based across critical activities. Success should be measured using clear metrics such as accuracy, safety, and privacy. It is important to use curated trusted test datasets as well as scenario walkthroughs that reflect real business workflows on and off the ledger. Layered defences can include security tests that mimic real attacks on AI models – such as prompt injection, jailbreaks, and data/private key/configuration exfiltration attempts – which are backed up by safeguards such as policy checks and human oversight. In production environments, gated development pipelines with pre-release evaluations, gradual or 'canary' rollouts, and continuous monitoring tied to automated rollback and audit trails can help mitigate unintended safety regressions and drifts.

AI models are inherently non-deterministic: even in response to the same prompts they can give different answers. This unpredictability can affect how different models work together. Techniques such as semantic and trigger-aware design, automated reviews within tolerance bands, and other variance-reduction strategies can help stabilise outcomes.

Testing AlxDLT systems goes beyond AI guardrails as blockchain and smart contracts bring their own vulnerabilities to consider, mitigate and address in the AI learning loop. These need to be tested and monitored too. Regular security reviews, updated tools and "red team" exercises – where experts try to break the system – can help keep outputs reliable and trusted.

- Consider wider impact and alignment with firm policies and culture:

It will be important for senior leaders to assess risk appetite and the organisation's priorities across many vectors in rolling out new technologies. Alongside increased regulatory attention, there is also growing scrutiny from the press and within certain workforces, making the consideration of any PR and HR risks a critical aspect of most projects. Alignment with an organisation's approach to wider issues is important. For example, energy sources required to run AI and DLT platforms and services may need to be considered by businesses as part of any rollout, factoring in alignment with any wider ESG goals or policies. Some models and platforms perform better than others in efficiency and energy consumption, and some AlxDLT use cases may even support sustainability goals. Many organisations are also considering the role of alternative energy sources in meeting the energy needs of newer technologies. Appropriate governance can assist organisations in taking an aligned approach to projects, which can be supported by appropriate governance structures, policies and effective internal systems.

- **Strategically assess and navigate legal obligations and risk:**


The current legal and regulatory landscape is multi-layered and rapidly shifting, with many laws having extra-territorial effect and some taking a risk-based approach, which means that obligations differ according to the nature of the use case. In regulated industries, AlxDLT projects will also be subject to industry-specific rules, which will need to be layered into any compliance strategy. Early legal due diligence, mapping of relevant frameworks and rules and working with trusted advisers who can offer a global perspective and sector-specific insights is critical in navigating the patchwork of legal and technical requirements that might apply in different jurisdictions, as well as in managing wider risk, such as potential legal liability. Prioritising compliance with requirements relating to data governance, security and operational resilience, transparency, accountability and fairness can also mitigate reputational and operational risk and earn customer trust.

- **Consider engagement with policymakers, regulators and industry stakeholders, and anticipate new legislation:**

In contrast to some more settled areas of law, new technology-focused regulations are disparate and may rapidly evolve. This dynamic is important to factor into any AlxDLT project or acquisition at the outset. For cutting-edge AlxDLT use cases that push into uncharted territory, organisations should be prepared to work with regulators and even help shape emerging policy. Some innovative projects may raise novel questions about responsibility and liability and organisations may want to make use of regulatory sandboxes or other opportunities to engage with regulators and policymakers. More broadly, regulators and policymakers are often looking for input from the industry as they develop frameworks and policies. Consider whether your organisation will be proactive in sharing insights and perspectives, for example through trade associations or direct participation in industry consultations, standards bodies, or sandboxes. In an area where it is crucial to stay ahead of the regulatory curve, active dialogue, advocacy and collaboration can give your organisation early visibility into likely rule changes as well as help to ensure that future regulations (or new legislation) are practical.

Contacts

Clifford Chance



Diego Ballon Ossio
Partner

diego.ballonossio@cliffordchance.com
+44 207006 3425



Laura Nixon
Knowledge Director – Fintech

laura.nixon@cliffordchance.com
+44 207006 8385



Marc Benzler
Partner

marc.benzler@cliffordchance.com
+49 69 7199 3304



Steve Gatti
Partner

steven.gatti@cliffordchance.com
+1 202 912 5095




Devika Kornbacher
Partner and Co-Chair, Global Tech Group

devika.kornbacher@cliffordchance.com
+1 7138212818



Holger Lutz
Partner

holger.lutz@cliffordchance.com
+49 69 7199 1670



Herbert Swaniker
Partner

herbert.swaniker@cliffordchance.com
+44 207006 6215




Rita Flakoll
Knowledge Director – Tech

rita.flakoll@cliffordchance.com
+44 207006 1826




Alexandre Balducci
Lawyer

alexandre.balducci@cliffordchance.com
+33 1 4405 5137



Jane Chen
Senior Associate

jane.chen@cliffordchance.com
+86 10 6535 2216



Jack Harris
Senior Associate

jack.harris@cliffordchance.com
+44 207006 1614



Kimi Liu
International Partner, Shanghai He Ping Law Firm

kimi.liu@hepinglaw.com
+86 10 6535 2263



Kate Scott
Partner

kate.scott@cliffordchance.com
+44 207006 4442



Charlotte Walker-Osborn
Knowledge Director – Tech

charlotte.walker-osborn@cliffordchance.com
+44 207006 2662




In partnership with
Deutsche Bank



Joy Adams
COO, Digital Assets & Currencies Transformation
joy.adams@db.com



Thomas Brophy
Digital Assets & Currencies Transformation
thomas.brophy@db.com



Sabih Behzad
Head of Digital Assets & Currencies Transformation
sabih.behzad@db.com



Tim Mason
Head of Product Innovation and AI Programme
tim.mason@db.com



References in this paper to 'China' and 'PRC' stand for the People's Republic of China. To the extent our reference is made to the PRC law, we have not accounted for laws that are applicable to each of the Hong Kong Special Administrative Region, Macau Special Administrative Region and Taiwan Region respectively. The geographic jurisdiction scope shall be interpreted accordingly.

As is the case for all international law firms with representative offices in the PRC, while Clifford Chance is authorised to provide information concerning the effect of the Chinese legal environment, we are not permitted to engage in Chinese legal affairs.

Clifford Chance LLP and Shanghai He Ping Law Firm (FTZ) Joint Operation Office is a joint operation established in the China (Shanghai) Pilot Free Trade Zone with the approval of the Shanghai Bureau of Justice. Shanghai He Ping Law Firm is a partnership established under the laws of the PRC and is licensed to practise PRC law. Legal advice in relation to the laws of the PRC is provided in the name of the joint operation by Shanghai He Ping Law Firm.

This publication does not necessarily deal with every important topic or cover every aspect of the topics with which it deals. It is not designed to provide legal or other advice.

cliffordchance.com

Clifford Chance, 10 Upper Bank Street, London, E14 5JJ

© Clifford Chance 2025

Clifford Chance LLP is a limited liability partnership registered in England and Wales under number OC323571

Registered office: 10 Upper Bank Street, London, E14 5JJ

We use the word 'partner' to refer to a member of Clifford Chance LLP, or an employee or consultant with equivalent standing and qualifications

If you do not wish to receive further information from Clifford Chance about events or legal developments which we believe may be of interest to you, please either send an email to nomorecontact@cliffordchance.com or by post at Clifford Chance LLP, 10 Upper Bank Street, Canary Wharf, London E14 5JJ

Abu Dhabi • Amsterdam • Barcelona • Beijing • Brussels • Bucharest • Casablanca • Delhi • Dubai • Düsseldorf • Frankfurt • Hong Kong • Houston • Istanbul • London • Luxembourg • Madrid • Milan • Munich • Newcastle • New York • Paris • Perth • Prague • Rome • São Paulo • Shanghai • Singapore • Sydney • Tokyo • Warsaw • Washington, D.C.

*AS&H Clifford Chance, a joint venture entered into by Clifford Chance LLP.

**Clifford Chance has entered into association agreements with Clifford Chance Prague Association SRO in Prague and Clifford Chance Badea SPRL in Bucharest.

Clifford Chance has a best friends relationship with Redcliffe Partners in Ukraine.